

Relating Strong Behavioral Equivalences for Processes with Nondeterminism and Probabilities

Marco Bernardo^a, Rocco De Nicola^b, Michele Loreti^c

^a*Dipartimento di Scienze di Base e Fondamenti – Università di Urbino – Italy*

^b*IMT – Institute for Advanced Studies Lucca – Italy*

^c*Dipartimento di Statistica, Informatica, Applicazioni – Università di Firenze – Italy*

Abstract

We present a comparison of behavioral equivalences for nondeterministic and probabilistic processes whose activities are all observable. In particular, we consider trace-based, testing, and bisimulation-based equivalences. For each of them, we examine the discriminating power of three variants stemming from three approaches that differ for the way probabilities of events are compared when nondeterministic choices are resolved via schedulers. The first approach compares two resolutions with respect to the probability distributions of all considered events. The second approach requires that the probabilities of the set of events of a resolution be individually matched by the probabilities of the same events in possibly different resolutions. The third approach only compares the extremal probabilities of each event stemming from the different resolutions. The three approaches have very reasonable motivations and, when applied to fully nondeterministic processes or fully probabilistic processes, give rise to the classical well studied relations. We shall see that, for processes with nondeterminism and probability, they instead give rise to a much wider variety of behavioral relations, whose discriminating power is thoroughly investigated here in the case of deterministic schedulers.

Keywords: bisimulation equivalence, testing equivalence, failure equivalence, trace equivalence, nondeterminism, probability

1. Introduction

Process algebras (see [3] and the references therein) are mathematically rigorous languages that have been widely used to model and analyze the behavior of interacting systems. Their structural operational semantics associates with each process term a labeled transition system (LTS) [27], whose states are the terms themselves and whose labels are the actions that each term can perform. In order to abstract from unwanted details, the operational semantics is often coupled with observational mechanisms that permit equating those systems that cannot be distinguished by external entities. The resulting behavioral equivalences heavily depend on how the specified systems are expected to be used. Indeed, there is still disagreement on which are the “reasonable” observations and how their outcomes can be used to distinguish or identify systems. Thus, many equivalences have been proposed and much work has been done to assess their discriminating power and mutual relationships.

The first study in this direction was presented in [13]. There, most of the then known equivalences over LTS models were “ordered” and it was shown that *trace equivalence* (equating systems performing the same sequences of actions) is strictly coarser than *failure equivalence* (equating systems performing the same sequences of actions and refusing the same sets of actions after them), which in turn is strictly coarser than *bisimulation equivalence* (equating systems performing the same sequences of actions and recursively exhibiting the same behavior). It was also shown that the equivalence obtained by *testing* processes with external observers was coincident with failure equivalence. Afterwards, [39] built the first spectrum that relates twelve different equivalences defined in the literature over LTS models and set up a general testing scenario that could be used to generate many more equivalences.

When process algebras have been enriched with additional dimensions to deal with probabilistic, stochastic, and timed systems, new behavioral equivalences have been defined and possible classifications have been proposed. Here, we would like to concentrate on equivalences for probabilistic systems. For this class of systems, comparative results have been obtained only for so-called fully probabilistic systems [26, 22, 2] or only for bisimulation-based and testing-based relations [2, 29, 35, 43].

In this paper, we aim at a systematic account of the known probabilistic equivalences for nondeterministic *and* probabilistic systems. We shall consider an extension of the LTS model combining nondeterminism and probability that we call Nondeterministic and Probabilistic LTS (NPLTS), 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 [28, 31]. Like in [31], we shall resort to the notion of scheduler (or adversary) to resolve nondeterminism. A *scheduler* 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 system, a fully probabilistic model, called *resolution*, is obtained. Actions will be assumed to be visible (i.e., τ -actions will not be considered) and, when defining the various strong equivalences, resolutions of nondeterminism will be obtained by applying memoryless deterministic schedulers.

Even once one has decided to have only visible actions and only certain classes of schedulers, the number of possibilities for defining behavioral equivalences over NPLTS models is very high if compared with the one for LTS models. These possibilities are determined by the combination of:

- the equivalence-specific events that have to be measured and
- the paths within resolutions that have to be compared.

Indeed, when checking probabilistic systems for equivalence, the mainly used measure is the probability that *equivalence-specific events* take place. Some of the options exploited in the literature are:

- performing specific sequences of actions, to obtain trace semantics;
- exhibiting specific traces decorated with additional information, to obtain failure or testing semantics;
- reaching certain sets of equivalent states via given actions, to obtain bisimulation semantics.

Instead, the alternatives one has in establishing a correspondence between paths of resolutions are: to look for a direct correspondence between all the paths of the considered resolutions; or rather to permit that sets of paths of one resolution could be matched by putting together paths of different resolutions. In this respect, three main approaches can be singled out:

1. The typical approach followed in the literature (see, e.g., [34, 32, 33]) consists of comparing the *probability distributions of all equivalence-specific events* of two resolutions. Two processes are considered 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* equivalence-specific event is the same in the two resolutions (*fully matching resolutions*). For the known relations based on this approach, we have that the probabilistic bisimilarity in [34] implies the probabilistic failure equivalence in [33] that in turn implies the probabilistic trace equivalence in [32]. All these relations are conservative extensions of the corresponding behavioral relations defined over fully nondeterministic models [21, 9] and fully probabilistic models [19, 26, 22], but in many situations they turn out to have a high discriminating power.
2. A different approach has been followed in the literature for defining testing equivalences (see, e.g., [44, 25, 33, 15]). Instead of comparing individual resolutions of the parallel composition of processes and tests, the comparison is performed between the *extremal probabilities* of reaching success *over all resolutions* generated by the experiments on processes under test (*max-min-matching resolution sets*). In this case, it holds that the resulting probabilistic testing equivalence is implied by the probabilistic bisimilarity in [34], but it is related neither to the probabilistic failure equivalence in [33] nor to the probabilistic trace equivalence in [32] when attention is restricted to deterministic schedulers. Moreover, the resulting probabilistic testing equivalence is neither a conservative extension of testing equivalence for fully nondeterministic processes [14] nor a conservative extension of testing equivalence for fully probabilistic processes [10].




Figure 1: Two NPLTS models distinguished by the first approach and identified by the third approach

3. Recently, in [12, 38, 37, 4, 8] a further approach has been proposed that compares resolutions on the basis of the *probabilities of individual equivalence-specific events*. A resolution of any of the two processes can be matched, with respect to *different equivalence-specific events*, by *different resolutions* of the other process (*partially matching resolutions*). For the behavioral relations resulting from this approach, which weakens the impact of schedulers, we have that probabilistic bisimilarity implies probabilistic failure equivalence, which in turn implies probabilistic testing equivalence, which finally implies probabilistic trace equivalence. This approach has contributed to the development of new probabilistic bisimilarities in [12], [38, 8], and [37] that, unlike the one in [34], are characterized by standard probabilistic logics such as quantitative μ -calculus, PML, and PCTL/PCTL*, respectively. Moreover, in the case of testing equivalence this approach has the advantage of being conservative for fully nondeterministic models and fully probabilistic models, while in the case of trace equivalence it surprisingly results in a congruence with respect to parallel composition (see the full version of [4]).

In order to appreciate the differences among the three approaches outlined above, let us consider the NPLTS models in Figs. 1 and 2 and analyze the impact on them of probabilistic bisimilarities (similar considerations do apply to the other probabilistic equivalences). The two models in Fig. 1 describe two scenarios representing the offer to Player1 and Player2 of three differently biased dice. The game is conceived in such a way that if the outcome of a throw gives 1 or 2 then Player1 wins, while if the outcome is 5 or 6 then Player2 wins. In case of 3 or 4, the result is a draw. Differently, in the first scenario of Fig. 2 the two players are offered a choice among a fair coin and two biased ones, while in the second scenario the players can simply choose between the two biased coins of the former scenario. In both scenarios of Fig. 2, Player1 wins with head while Player2 wins with tail.

The basic idea behind the first approach is deeming equivalent two processes if and only if for each resolution of one process (the challenger) there exists a resolution of the other process (the defender) such that the two resolutions are probabilistic bisimilar in the sense of [19]. This leads to distinguishing both the pair of processes in Fig. 1 and that in Fig. 2, as ensured by the probabilistic bisimilarity (based on deterministic schedulers) in [34].

In some cases, the first approach might end up being too demanding. Indeed, if one is interested in the set of probabilities of winning/drawing/losing, which is $\{0.6, 0.4, 0\}$ for both players, it is conceivable to consider equivalent the two processes in Fig. 1. In fact, these two processes are identified by equivalences defined according to the third approach, as ensured by one of the two probabilistic bisimilarities in [8]. Similarly, if what matters is the extremal – i.e., minimal and maximal – probabilities of winning (0.3 and 0.7), it is conceivable that the two processes in Fig. 2 be identified. This is indeed the case for equivalences defined according to the second approach, like the other probabilistic bisimilarity in [8].

To obtain equivalences that identify the processes in Fig. 1 and those in Fig. 2, it is necessary to weaken the role of schedulers. While in [34] the challenger and the defender must stepwise behave the same along two matching resolutions of nondeterminism, one might offer the defender the possibility of choosing different resolutions in response to different directions taken by the challenger. In other words, instead of requiring, as in [34], that for each resolution of the challenger there is a fully matching resolution of the defender, one could consider bisimulation games with partially matching resolutions as in [12, 38, 37, 8]. The new equivalences can then be obtained by comparing processes according to all sets of probabilities or only to the extremal probabilities.




Figure 2: Two NPLTS models distinguished by the third approach and identified by the second approach

In our view, the motivations behind the three approaches are all very reasonable. Indeed, when applied to fully nondeterministic processes or fully probabilistic processes, they give rise to well-studied relations that for the fully nondeterministic setting fit into the spectra in [13, 39] and for the fully probabilistic setting fit into the spectra in [26, 22]. The situation is significantly different when the three approaches are instantiated for nondeterministic *and* probabilistic processes; in this case, they give rise to a much wider variety of relations.

In the paper, for each of the three approaches we shall first consider trace, failure, testing, and bisimulation equivalences over NPLTS models, then, to complete the picture, we shall discuss also other variants such as completed-trace equivalence, other decorated-trace equivalences (failure trace, readiness, and ready trace), and the kernels of simulation-based preorders. In the case of deterministic schedulers, we shall see that the family of equivalences that assign a central role to schedulers by requiring that the result of a specific choice in one process be fully matched by the other one (*fully matching resolutions*), yields a hierarchy that is in accordance with the one for fully probabilistic processes in [26, 22]. Conversely, the family of equivalences that assign a weaker role to schedulers in resolving nondeterminism (*partially matching resolutions*), gives rise to relations that are coarser than the previous ones and yields a hierarchy that is in accordance with the one for fully nondeterministic processes in [13, 39]. Finally, the family of equivalences that only consider extremal probabilities (*max-min-matching resolution sets*), has again several analogies with the fully nondeterministic spectrum and yields even coarser relations. There are however some noticeable anomalies in the last two families, evidenced by a few equivalences that are incomparable with most of the others.

The rest of the paper, which is a revised and extended version of [6], is organized as follows. In Sect. 2, we introduce some basic notions about the NPLTS model. In Sects. 3 to 6, we define and compare the trace, failure, testing, and bisimulation equivalences, respectively, that arise from the three approaches outlined above in the case of deterministic schedulers. In Sect. 7, we graphically summarize the results by depicting the spectrum of the considered equivalences, somehow in analogy with the reduced spectrum in [13]. In the same section, additional equivalences are considered and a fuller spectrum, analogous to that in [39], is presented. Finally, in Sect. 8 we draw some conclusions, discuss how the spectrum changes in the case of randomized schedulers, and indicate directions for future work.

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 [16] and are representative of a number of slightly different probabilistic computational models including internal nondeterminism such as, e.g., concurrent Markov chains [42], alternating probabilistic models [20, 44, 30], probabilistic automata in the sense of [31], and the denotational probabilistic models in [23] (see [36] for an overview). We formalize them as a variant of simple probabilistic automata [31].

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.




Figure 3: Graphical representation of NPLTS models: two examples

- A is a countable set of transition-labeling actions.
- $\longrightarrow \subseteq S \times A \times \text{Distr}(S)$ is a transition relation, where $\text{Distr}(S)$ is the set of discrete probability distributions over S . \blacksquare

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., $\text{supp}(\mathcal{D}) = \{s' \in S \mid \mathcal{D}(s') > 0\}$. We write $s \xrightarrow{a}$ to indicate that s has an a -transition. The choice among all the transitions departing from s is external and nondeterministic, while the choice of the target state for a specific transition is internal and probabilistic. An NPLTS represents (i) a *fully nondeterministic process* when every transition leads to a distribution that concentrates all the probability mass into a single target state or (ii) a *fully probabilistic process* when every state has at most one outgoing transition.

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 representing state s to a set of vertices linked by a dashed line, each of which represents a state $s' \in \text{supp}(\mathcal{D})$ and is labeled with $\mathcal{D}(s')$ – label omitted if $\mathcal{D}(s') = 1$. Figure 3 shows two NPLTS models: the one on the left mixes internal nondeterminism and probability, while the one on the right does not.

In this setting, a computation is a sequence of state-to-state steps, each denoted by $s \xrightarrow{a} s'$ 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:

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

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 $s_i \in \text{supp}(\mathcal{D}_i)$, with $\mathcal{D}_i(s_i)$ being the execution probability of step $s_{i-1} \xrightarrow{a_i} s_i$ 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. We denote by $\text{first}(c)$ and $\text{last}(c)$ the initial state and the final state of c , respectively, and by $\mathcal{C}_{\text{fin}}(s)$ the set of finite-length computations from s . \blacksquare

A resolution of a state s of an NPLTS \mathcal{L} is the result of a 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 convex combination of equally labeled transitions of \mathcal{L} (*randomized scheduler*) among all the outgoing transitions of that state. Below, we introduce the notion of resolution arising from a deterministic scheduler as a fully probabilistic NPLTS (randomized schedulers are discussed in Sect. 8). Notice that, when \mathcal{L} is fully nondeterministic, resolutions boil down to computations.

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 $\text{corr}_{\mathcal{Z}} : Z \rightarrow S$ such that $s = \text{corr}_{\mathcal{Z}}(z_s)$, for some $z_s \in Z$, and for all $z \in Z$ it holds that:

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

We say that 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 $\text{Res}(s)$ the set of resolutions of s obtained via a deterministic scheduler and by $\text{Res}_{\max}(s)$ the set of maximal resolutions of s obtained via a deterministic scheduler. ■

Since $Z \in \text{Res}(s)$ is fully probabilistic, the probability $\text{prob}(c)$ 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 $\text{prob}(c)$ being always equal to 1 if \mathcal{L} is fully nondeterministic. This notion is lifted to $C \subseteq \mathcal{C}_{\text{fin}}(z_s)$ by letting $\text{prob}(C) = \sum_{c \in C} \text{prob}(c)$ whenever none of the computations in C is a proper prefix of one of the others.

We finally introduce a notion of fully synchronous parallel composition for NPLTS models, borrowed from [25], that is instrumental to the definition of testing equivalences.

Definition 2.4. Let $\mathcal{L}_i = (S_i, A, \longrightarrow_i)$ be an NPLTS for $i = 1, 2$. The *parallel composition* of \mathcal{L}_1 and \mathcal{L}_2 is the NPLTS $\mathcal{L}_1 \parallel \mathcal{L}_2 = (S_1 \times S_2, A, \longrightarrow)$ where $\longrightarrow \subseteq (S_1 \times S_2) \times A \times \text{Distr}(S_1 \times S_2)$ is such that $(s_1, s_2) \xrightarrow{a} \mathcal{D}$ iff $s_1 \xrightarrow{a}_1 \mathcal{D}_1$ and $s_2 \xrightarrow{a}_2 \mathcal{D}_2$ with $\mathcal{D}(s'_1, s'_2) = \mathcal{D}_1(s'_1) \cdot \mathcal{D}_2(s'_2)$ for each $(s'_1, s'_2) \in S_1 \times S_2$. ■

3. Trace Equivalences for NPLTS Models

Trace equivalences examine the probability with which two states perform computations labeled with the same traces for each possible way of resolving nondeterminism. As outlined in Sect. 1, there are three different approaches to defining them. The first approach is to match resolutions according to *trace distributions*, which means that for each resolution of one of the two states there must exist a resolution of the other state such that, *for every trace*, the two resolutions have the same probability of performing a computation labeled with that trace. In other words, matching resolutions of the two states are related by the fully probabilistic version of the trace equivalence (fully matching resolutions). The second approach is to consider *a single trace at a time*, i.e., to anticipate the quantification over traces with respect to the quantification over resolutions. In this way, differently labeled computations of a resolution of one of the two states are allowed to be matched by computations of several different resolutions of the other state (partially matching resolutions). The third approach is to compare only the *extremal probabilities* of performing each trace over the various resolutions (max-min-matching resolution sets).

Given an NPLTS $\mathcal{L} = (S, A, \longrightarrow)$, we call *trace* a finite sequence of actions of A . We say that a finite-length computation c is *compatible* with a trace $\alpha \in A^*$, or equivalently that c is an α -computation, iff the sequence of actions labeling the steps of c is equal to α . Given $s \in S$ and $Z \in \text{Res}(s)$, we denote by $\text{CC}(z_s, \alpha)$ the set of α -compatible computations from z_s and by $\text{Res}_\alpha(s)$ the set of resolutions in $\text{Res}(s)$ having no maximal computations corresponding to proper prefixes of α -computations of \mathcal{L} .

In the following equivalence definitions, as well as in those of the next sections, we assume $s_1, s_2 \in S$ and we explicitly add a reference whenever the defined equivalence has already appeared in the literature. Moreover, \sqcup/\sqcap denotes the supremum/infimum of a set of numbers in $\mathbb{R}_{[0,1]}$, which is assumed to be 0 when the set is empty.

Definition 3.1. (*Probabilistic trace-distribution equivalence* – $\sim_{\text{PTr}, \text{dis}}$ – [32])

$s_1 \sim_{\text{PTr}, \text{dis}} s_2$ iff for each $Z_1 \in \text{Res}(s_1)$ there exists $Z_2 \in \text{Res}(s_2)$ such that for all $\alpha \in A^*$:

$$\text{prob}(\text{CC}(z_{s_1}, \alpha)) = \text{prob}(\text{CC}(z_{s_2}, \alpha))$$

and symmetrically for each $Z_2 \in \text{Res}(s_2)$. ■

Definition 3.2. (*Probabilistic trace equivalence* – \sim_{PTr} – [4])

$s_1 \sim_{\text{PTr}} s_2$ iff for all $\alpha \in A^*$ it holds that for each $Z_1 \in \text{Res}(s_1)$ there exists $Z_2 \in \text{Res}(s_2)$ such that:

$$\text{prob}(\text{CC}(z_{s_1}, \alpha)) = \text{prob}(\text{CC}(z_{s_2}, \alpha))$$

and symmetrically for each $Z_2 \in \text{Res}(s_2)$. ■




Figure 4: Two NPLTS models distinguished by $\sim_{\text{PTr}, \sqcup \sqcap}$ if only resolutions having α -computations were considered

Definition 3.3. (*Probabilistic $\sqcup \sqcap$ -trace equivalence* – $\sim_{\text{PTr}, \sqcup \sqcap}$)

$s_1 \sim_{\text{PTr}, \sqcup \sqcap} s_2$ iff for all $\alpha \in A^*$:

$$\begin{aligned} \bigsqcup_{Z_1 \in \text{Res}_\alpha(s_1)} \text{prob}(\text{CC}(z_{s_1}, \alpha)) &= \bigsqcup_{Z_2 \in \text{Res}_\alpha(s_2)} \text{prob}(\text{CC}(z_{s_2}, \alpha)) \\ \bigsqcap_{Z_1 \in \text{Res}_\alpha(s_1)} \text{prob}(\text{CC}(z_{s_1}, \alpha)) &= \bigsqcap_{Z_2 \in \text{Res}_\alpha(s_2)} \text{prob}(\text{CC}(z_{s_2}, \alpha)) \end{aligned}$$

■

The three trace equivalences defined above are all backward compatible with the trace equivalences respectively defined in [9] for fully nondeterministic processes – which we denote by $\sim_{\text{Tr}, \text{fnd}}$ – and in [26, 22] for fully probabilistic processes – which we denote by $\sim_{\text{Tr}, \text{fpr}}$.

Theorem 3.4. It holds that:

1. $\sim_{\text{PTr}, \text{dis}} = \sim_{\text{PTr}} = \sim_{\text{PTr}, \sqcup \sqcap} = \sim_{\text{Tr}, \text{fnd}}$ over fully nondeterministic NPLTS models.
2. $\sim_{\text{PTr}, \text{dis}} = \sim_{\text{PTr}} = \sim_{\text{PTr}, \sqcup \sqcap} = \sim_{\text{Tr}, \text{fpr}}$ over fully probabilistic NPLTS models.

Proof

1. The result over fully nondeterministic NPLTS models is a straightforward consequence of the fact that the resolutions of these models correspond to the computations of the models themselves, hence the probability of performing within a resolution of one of these models a computation compatible with a trace can only be either 1 or 0.
2. The result over fully probabilistic NPLTS models is a straightforward consequence of the fact that each of these models has a single maximal resolution, which corresponds to the model itself. ■

It is worth observing that in Def. 3.3 we consider only resolutions having no maximal computations corresponding to proper prefixes of computations that are labeled with the trace under examination in the two processes. While this restriction is negligible for the computation of suprema, it avoids infima to be trivially equal to 0. In fact, apart from the empty trace ε that does not have any proper prefix and results in a maximum probability and a minimum probability both equal to 1, for an arbitrary trace $\alpha \neq \varepsilon$ it holds that the probability of performing a computation compatible with α is 0 along every resolution in which all computations have length less than the length of α , and hence the infimum of such probabilities is 0. Thus, for a meaningful comparison of infima, we have at least to restrict attention to resolutions having no maximal computations corresponding to proper prefixes of α -computations of the two processes, so that the infimum on either side is 0 only if there is a sufficiently extended resolution along which α cannot be performed.

This restriction is weaker than considering only resolutions having at least one α -computation. Such a stronger restriction leads however to counterintuitive facts, such as getting an infimum equal to 0 only if α cannot be performed along any resolution. As a consequence, the resulting $\sim_{\text{PTr}, \sqcup \sqcap}$ would distinguish the two processes in Fig. 4 because, for $\alpha = ab$, in s_1 only the leftmost maximal resolution and the central maximal resolution would be considered – with minimum probability equal to 0.5 – and in s_2 only the leftmost maximal resolution would be considered – with minimum probability equal to 1. The rightmost maximal resolution of s_1 and that of s_2 would not be considered as they have no α -computation, thus excluding a minimum probability of performing α equal to 0 for both processes.




Figure 5: Two NPLTS models distinguished by $\sim_{\text{PTr}, \sqcup \cap}$ if only maximal resolutions were considered

One may be tempted to impose a restriction independent from a specific trace such as considering only maximal resolutions, which works well for testing equivalences as we shall see in Sect. 5. With this restriction, the two fully nondeterministic processes depicted in Fig. 5 would be distinguished by $\sim_{\text{PTr}, \sqcup \cap}$ because, for trace $\alpha = ab$, the minimum probability of performing an α -computation is 0 over the two maximal resolutions of s_1 and 1 over the only maximal resolution of s_2 . But then $\sim_{\text{PTr}, \sqcup \cap}$ would not be backward compatible with $\sim_{\text{Tr}, \text{fnd}}$ because the two processes in Fig. 5 are identified by $\sim_{\text{Tr}, \text{fnd}}$.

We now investigate the relationships among the three trace equivalences. As expected, the equivalence relying on trace distributions is finer than the equivalence considering a single trace at a time, which in turn is finer than the equivalence based on extremal probabilities of traces.

Theorem 3.5. It holds that $\sim_{\text{PTr}, \text{dis}} \subseteq \sim_{\text{PTr}} \subseteq \sim_{\text{PTr}, \sqcup \cap}$.

Proof Let (S, A, \longrightarrow) be an NPLTS and $s_1, s_2 \in S$. The fact that $s_1 \sim_{\text{PTr}, \text{dis}} s_2$ implies $s_1 \sim_{\text{PTr}} s_2$ is easily seen by taking the same fully matching resolutions considered in $\sim_{\text{PTr}, \text{dis}}$.

Suppose now that $s_1 \sim_{\text{PTr}} s_2$. This implies that for all $\alpha \in A^*$ it holds that:

- For each $Z_1 \in \text{Res}_\alpha(s_1)$ there exists $Z_2 \in \text{Res}_\alpha(s_2)$ such that:
 $\text{prob}(\text{CC}(z_{s_1}, \alpha)) = \text{prob}(\text{CC}(z_{s_2}, \alpha))$
- For each $Z_2 \in \text{Res}_\alpha(s_2)$ there exists $Z_1 \in \text{Res}_\alpha(s_1)$ such that:
 $\text{prob}(\text{CC}(z_{s_2}, \alpha)) = \text{prob}(\text{CC}(z_{s_1}, \alpha))$

This is to say that:

$$\begin{aligned} \bigcup_{Z_1 \in \text{Res}_\alpha(s_1)} \{\text{prob}(\text{CC}(z_{s_1}, \alpha))\} &\subseteq \bigcup_{Z_2 \in \text{Res}_\alpha(s_2)} \{\text{prob}(\text{CC}(z_{s_2}, \alpha))\} \\ \bigcup_{Z_2 \in \text{Res}_\alpha(s_2)} \{\text{prob}(\text{CC}(z_{s_2}, \alpha))\} &\subseteq \bigcup_{Z_1 \in \text{Res}_\alpha(s_1)} \{\text{prob}(\text{CC}(z_{s_1}, \alpha))\} \end{aligned}$$

Equivalently:

$$\bigcup_{Z_1 \in \text{Res}_\alpha(s_1)} \{\text{prob}(\text{CC}(z_{s_1}, \alpha))\} = \bigcup_{Z_2 \in \text{Res}_\alpha(s_2)} \{\text{prob}(\text{CC}(z_{s_2}, \alpha))\}$$

which implies:

$$\begin{aligned} \bigsqcup_{Z_1 \in \text{Res}_\alpha(s_1)} \text{prob}(\text{CC}(z_{s_1}, \alpha)) &= \bigsqcup_{Z_2 \in \text{Res}_\alpha(s_2)} \text{prob}(\text{CC}(z_{s_2}, \alpha)) \\ \prod_{Z_1 \in \text{Res}_\alpha(s_1)} \text{prob}(\text{CC}(z_{s_1}, \alpha)) &= \prod_{Z_2 \in \text{Res}_\alpha(s_2)} \text{prob}(\text{CC}(z_{s_2}, \alpha)) \end{aligned}$$

This means that $s_1 \sim_{\text{PTr}, \sqcup \cap} s_2$. ■

Both inclusions in Thm. 3.5 are strict:

- Figure 1 shows that $\sim_{\text{PTr}, \text{dis}}$ is strictly finer than \sim_{PTr} . It holds that $s_1 \not\sim_{\text{PTr}, \text{dis}} s_2$ because for instance the trace distribution of the leftmost maximal resolution of s_1 – which assigns probability 1 to trace ε and trace *offer*, probability 0.4 to trace *offer draw*, probability 0.6 to trace *offer win*₁, and probability 0 to any other trace – is not matched by the trace distribution of any of the three maximal resolutions of s_2 . In contrast, $s_1 \sim_{\text{PTr}} s_2$ because, given an arbitrary trace α , for each resolution of s_1

(resp. s_2) there exists a resolution of s_2 (resp. s_1) such that an α -compatible computation has the same probability of being executed in both resolutions. For example, the leftmost maximal resolution of s_1 is matched by the central maximal resolution of s_2 with respect to trace *offer draw* and by the rightmost maximal resolution of s_2 with respect to trace *offer win₁*.

- Figure 2 shows that \sim_{PTr} is strictly finer than $\sim_{\text{PTr}, \sqcup \sqcap}$. It holds that $s_1 \not\sim_{\text{PTr}} s_2$ because for instance the probability 0.5 of executing the trace *offer win₁* in the central maximal resolution of s_1 is not matched by the probability of executing the same trace in any of the two maximal resolutions of s_2 . In contrast, $s_1 \sim_{\text{PTr}, \sqcup \sqcap} s_2$ because, given an arbitrary trace α , the maximum probability and the minimum probability of performing α over all resolutions having no maximal computations corresponding to proper prefixes of α -computations of the two processes are respectively the same in both processes. For example, trace *offer win₁* has maximum probability 0.7 and minimum probability 0.3 in both processes.

4. Failure Equivalences for NPLTS Models

Failure equivalences generalize trace equivalences by considering the actions that can be refused after performing a trace. Given an NPLTS $\mathcal{L} = (S, A, \longrightarrow)$, we call *failure pair* an element φ of $A^* \times 2^A$ formed by a trace α and a decoration F called *failure set*. Given $s \in S$ and $\mathcal{Z} \in \text{Res}(s)$, we say that $c \in \mathcal{C}_{\text{fin}}(z_s)$ is *compatible* with φ iff $c \in \mathcal{CC}(z_s, \alpha)$ and $\text{corr}_{\mathcal{Z}}(\text{last}(c))$ 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 4.1. (*Probabilistic failure-distribution equivalence* – $\sim_{\text{PF}, \text{dis}}$ – [33])

$s_1 \sim_{\text{PF}, \text{dis}} s_2$ iff for each $\mathcal{Z}_1 \in \text{Res}(s_1)$ there exists $\mathcal{Z}_2 \in \text{Res}(s_2)$ such that for all $\varphi \in A^* \times 2^A$:

$$\text{prob}(\mathcal{FCC}(z_{s_1}, \varphi)) = \text{prob}(\mathcal{FCC}(z_{s_2}, \varphi))$$

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

Definition 4.2. (*Probabilistic failure equivalence* – \sim_{PF} – [4])

$s_1 \sim_{\text{PF}} s_2$ iff for all $\varphi \in A^* \times 2^A$ it holds that for each $\mathcal{Z}_1 \in \text{Res}(s_1)$ there exists $\mathcal{Z}_2 \in \text{Res}(s_2)$ such that:

$$\text{prob}(\mathcal{FCC}(z_{s_1}, \varphi)) = \text{prob}(\mathcal{FCC}(z_{s_2}, \varphi))$$

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

Definition 4.3. (*Probabilistic $\sqcup \sqcap$ -failure equivalence* – $\sim_{\text{PF}, \sqcup \sqcap}$)

$s_1 \sim_{\text{PF}, \sqcup \sqcap} s_2$ iff for all $\varphi = (\alpha, F) \in A^* \times 2^A$:

$$\begin{aligned} \bigsqcup_{\mathcal{Z}_1 \in \text{Res}_{\alpha}(s_1)} \text{prob}(\mathcal{FCC}(z_{s_1}, \varphi)) &= \bigsqcup_{\mathcal{Z}_2 \in \text{Res}_{\alpha}(s_2)} \text{prob}(\mathcal{FCC}(z_{s_2}, \varphi)) \\ \bigsqcap_{\mathcal{Z}_1 \in \text{Res}_{\alpha}(s_1)} \text{prob}(\mathcal{FCC}(z_{s_1}, \varphi)) &= \bigsqcap_{\mathcal{Z}_2 \in \text{Res}_{\alpha}(s_2)} \text{prob}(\mathcal{FCC}(z_{s_2}, \varphi)) \end{aligned}$$

■

The three failure equivalences defined above are all backward compatible with the failure equivalences respectively defined in [9] for fully nondeterministic processes – which we denote by $\sim_{\text{F}, \text{fnd}}$ – and in [26, 22] for fully probabilistic processes – which we denote by $\sim_{\text{F}, \text{fpr}}$.

Theorem 4.4. It holds that:

1. $\sim_{\text{PF}, \text{dis}} = \sim_{\text{PF}} = \sim_{\text{PF}, \sqcup \sqcap} = \sim_{\text{F}, \text{fnd}}$ over fully nondeterministic NPLTS models.
2. $\sim_{\text{PF}, \text{dis}} = \sim_{\text{PF}} = \sim_{\text{PF}, \sqcup \sqcap} = \sim_{\text{F}, \text{fpr}}$ over fully probabilistic NPLTS models.

Proof Similar to the proof of Thm. 3.4. ■

We point out that in Def. 4.3 the considered resolutions are the same as those considered in Def. 3.3. Note that the failure set F is not taken into account when selecting resolutions. For example, this is not necessary to distinguish the two fully nondeterministic processes depicted in Fig. 5. Given the failure pair $\varphi = (a, A)$, state s_1 has three resolutions each having no maximal computations corresponding to proper

prefixes of a -computations in the original process, and along the one formed by the rightmost a -transition a state is reached after performing a that refuses all actions (in the original process). In contrast, state s_2 has two such resolutions, but in both of them the state reached after performing a cannot refuse all actions (in the original process). Therefore, the two processes have different maximum probabilities of getting to a deadlock state after performing a .

We now investigate the relationships of the three failure equivalences among themselves (first property below) and with the three trace equivalences defined in Sect. 3 (last three properties below). As expected, each of the three failure equivalences is finer than the corresponding trace equivalence.

Theorem 4.5. It holds that:

1. $\sim_{\text{PF}, \text{dis}} \subseteq \sim_{\text{PF}} \subseteq \sim_{\text{PF}, \sqcup \sqcap}$.
2. $\sim_{\text{PF}, \text{dis}} \subseteq \sim_{\text{PTr}, \text{dis}}$.
3. $\sim_{\text{PF}} \subseteq \sim_{\text{PTr}}$.
4. $\sim_{\text{PF}, \sqcup \sqcap} \subseteq \sim_{\text{PTr}, \sqcup \sqcap}$.

Proof Let (S, A, \longrightarrow) be an NPLTS and $s_1, s_2 \in S$:

1. Similar to the proof of Thm. 3.5.
2. Suppose that $s_1 \sim_{\text{PF}, \text{dis}} s_2$. Then we immediately derive that:
 - For each $\mathcal{Z}_1 \in \text{Res}(s_1)$ there exists $\mathcal{Z}_2 \in \text{Res}(s_2)$ such that for all $\alpha \in A^*$:
 $\text{prob}(\text{CC}(z_{s_1}, \alpha)) = \text{prob}(\text{FCC}(z_{s_1}, (\alpha, \emptyset))) = \text{prob}(\text{FCC}(z_{s_2}, (\alpha, \emptyset))) = \text{prob}(\text{CC}(z_{s_2}, \alpha))$
 - Symmetrically for each $\mathcal{Z}_2 \in \text{Res}(s_2)$.

This means that $s_1 \sim_{\text{PTr}, \text{dis}} s_2$.

3. Suppose that $s_1 \sim_{\text{PF}} s_2$. Then we immediately derive that for all $\alpha \in A^*$:
 - For each $\mathcal{Z}_1 \in \text{Res}(s_1)$ there exist $\mathcal{Z}_2 \in \text{Res}(s_2)$ such that:
 $\text{prob}(\text{CC}(z_{s_1}, \alpha)) = \text{prob}(\text{FCC}(z_{s_1}, (\alpha, \emptyset))) = \text{prob}(\text{FCC}(z_{s_2}, (\alpha, \emptyset))) = \text{prob}(\text{CC}(z_{s_2}, \alpha))$
 - Symmetrically for each $\mathcal{Z}_2 \in \text{Res}(s_2)$.

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

4. Suppose that $s_1 \sim_{\text{PF}, \sqcup \sqcap} s_2$. Then we immediately derive that for all $\alpha \in A^*$:

$$\begin{aligned}
\bigcup_{\mathcal{Z}_1 \in \text{Res}_\alpha(s_1)} \text{prob}(\text{CC}(z_{s_1}, \alpha)) &= \bigcup_{\mathcal{Z}_1 \in \text{Res}_\alpha(s_1)} \text{prob}(\text{FCC}(z_{s_1}, (\alpha, \emptyset))) = \\
&= \bigcup_{\mathcal{Z}_2 \in \text{Res}_\alpha(s_2)} \text{prob}(\text{FCC}(z_{s_2}, (\alpha, \emptyset))) = \bigcup_{\mathcal{Z}_2 \in \text{Res}_\alpha(s_2)} \text{prob}(\text{CC}(z_{s_2}, \alpha)) \\
\bigcap_{\mathcal{Z}_1 \in \text{Res}_\alpha(s_1)} \text{prob}(\text{CC}(z_{s_1}, \alpha)) &= \bigcap_{\mathcal{Z}_1 \in \text{Res}_\alpha(s_1)} \text{prob}(\text{FCC}(z_{s_1}, (\alpha, \emptyset))) = \\
&= \bigcap_{\mathcal{Z}_2 \in \text{Res}_\alpha(s_2)} \text{prob}(\text{FCC}(z_{s_2}, (\alpha, \emptyset))) = \bigcap_{\mathcal{Z}_2 \in \text{Res}_\alpha(s_2)} \text{prob}(\text{CC}(z_{s_2}, \alpha))
\end{aligned}$$

This means that $s_1 \sim_{\text{PTr}, \sqcup \sqcap} s_2$. ■

All the inclusions in Thm. 4.5 are strict:

- Figures 1 and 2 respectively show that $\sim_{\text{PF}, \text{dis}}$ is strictly finer than \sim_{PF} and \sim_{PF} is strictly finer than $\sim_{\text{PF}, \sqcup \sqcap}$.
- Figure 6 shows that $\sim_{\text{PF}, \text{dis}}$, \sim_{PF} , and $\sim_{\text{PF}, \sqcup \sqcap}$ are strictly finer than $\sim_{\text{PTr}, \text{dis}}$, \sim_{PTr} , and $\sim_{\text{PTr}, \sqcup \sqcap}$, respectively. Indeed, for each resolution of s_1 (resp. s_2) there exists a resolution of s_2 (resp. s_1) such that both resolutions have precisely the same trace distribution, thus s_1 and s_2 are identified by $\sim_{\text{PTr}, \text{dis}}$ (and hence by \sim_{PTr} and $\sim_{\text{PTr}, \sqcup \sqcap}$). In contrast, the leftmost a -computation of s_1 is compatible with the failure pair $(a, \{c\})$ while s_2 has no computation compatible with that failure pair, hence s_1 and s_2 are distinguished by $\sim_{\text{PF}, \sqcup \sqcap}$ (and hence by \sim_{PF} and $\sim_{\text{PF}, \text{dis}}$).




Figure 6: Two NPLTS models distinguished by $\sim_{\text{PF},\text{dis}}/\sim_{\text{PF}}/\sim_{\text{PF},\sqcup\cap}$ and identified by $\sim_{\text{PTr},\text{dis}}/\sim_{\text{PTr}}/\sim_{\text{PTr},\sqcup\cap}$

Moreover:

- \sim_{PF} and $\sim_{\text{PF},\sqcup\cap}$ are incomparable with $\sim_{\text{PTr},\text{dis}}$, because in Fig. 1 it holds that $s_1 \sim_{\text{PF}} s_2$ (and hence $s_1 \sim_{\text{PF},\sqcup\cap} s_2$) and $s_1 \not\sim_{\text{PTr},\text{dis}} s_2$, while in Fig. 6 it holds that $s_1 \not\sim_{\text{PF},\sqcup\cap} s_2$ (and hence $s_1 \not\sim_{\text{PF}} s_2$) and $s_1 \sim_{\text{PTr},\text{dis}} s_2$.
- $\sim_{\text{PF},\sqcup\cap}$ is incomparable with \sim_{PTr} , because in Fig. 2 it holds that $s_1 \sim_{\text{PF},\sqcup\cap} s_2$ and $s_1 \not\sim_{\text{PTr}} s_2$, while in Fig. 6 it holds that $s_1 \not\sim_{\text{PF},\sqcup\cap} s_2$ and $s_1 \sim_{\text{PTr}} s_2$.

5. Testing Equivalences for NPLTS Models

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

Definition 5.1. 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 that has no outgoing transitions. We say that a computation of \mathcal{T} is successful iff its last state is ω . ■

Definition 5.2. 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 $s \in S$, $o \in O$, and $\mathcal{Z} \in \text{Res}(s, o)$, 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 $\mathcal{I}(\mathcal{L}, \mathcal{T})$. ■

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 a test 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 resolution of the interaction system. We can restrict our attention to maximal resolutions because they contain all successful computations. One option is comparing, for the two configurations, *only the extremal values of these success probabilities* over all maximal resolutions of the interaction system – considering non-maximal resolutions would lead to obtain always 0 as infimum. An alternative option is comparing *all the success probabilities* and requiring that for each maximal resolution of either configuration there is a matching maximal resolution of the other configuration.




Figure 7: Two NPLTS models distinguished by $\sim_{\text{PTe-}\sqcup\cap}/\sim_{\text{PTe-}\forall\exists}$ and identified by classical testing equivalence

Definition 5.3. (*Probabilistic $\sqcup\cap$ -testing equivalence* – $\sim_{\text{PTe-}\sqcup\cap}$ – [44, 25, 33, 15])

$s_1 \sim_{\text{PTe-}\sqcup\cap} s_2$ iff for every NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$:

$$\begin{aligned} \bigsqcup_{Z_1 \in \text{Res}_{\max}(s_1, o)} \text{prob}(\text{SC}(z_{s_1, o})) &= \bigsqcup_{Z_2 \in \text{Res}_{\max}(s_2, o)} \text{prob}(\text{SC}(z_{s_2, o})) \\ \prod_{Z_1 \in \text{Res}_{\max}(s_1, o)} \text{prob}(\text{SC}(z_{s_1, o})) &= \prod_{Z_2 \in \text{Res}_{\max}(s_2, o)} \text{prob}(\text{SC}(z_{s_2, o})) \end{aligned}$$

■

Definition 5.4. (*Probabilistic $\forall\exists$ -testing equivalence* – $\sim_{\text{PTe-}\forall\exists}$ – [4])

$s_1 \sim_{\text{PTe-}\forall\exists} s_2$ iff for every NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ it holds that for each $Z_1 \in \text{Res}_{\max}(s_1, o)$ there exists $Z_2 \in \text{Res}_{\max}(s_2, o)$ such that:

$$\text{prob}(\text{SC}(z_{s_1, o})) = \text{prob}(\text{SC}(z_{s_2, o}))$$

and symmetrically for each $Z_2 \in \text{Res}_{\max}(s_2, o)$.

■

Neither $\sim_{\text{PTe-}\sqcup\cap}$ nor $\sim_{\text{PTe-}\forall\exists}$ is backward compatible with the testing equivalence defined in [14] for fully nondeterministic processes. For instance, Fig. 7 shows two such processes identified by the classical testing equivalence, which are distinguished by $\sim_{\text{PTe-}\sqcup\cap}$ and $\sim_{\text{PTe-}\forall\exists}$ as can be seen by taking the test in the same figure. Indeed, the two maximal resolutions of the interaction system of the second process and the test reach success with probability p_1 and p_2 , respectively. In contrast, the interaction system of the first process and the test results in four maximal resolutions in which success is respectively reached with probability 1, p_1 , p_2 , and 0. Likewise, $\sim_{\text{PTe-}\sqcup\cap}$ and $\sim_{\text{PTe-}\forall\exists}$ are not backward compatible with the testing equivalence defined in [10] for fully probabilistic processes (see the full version of [4] for a counterexample).

The reason of the higher discriminating power of $\sim_{\text{PTe-}\sqcup\cap}$ and $\sim_{\text{PTe-}\forall\exists}$ arises from the presence of probabilistic choices within tests, which results in the capability of making copies of intermediate states of the process under test [1] and hence in a questionable estimation of success probabilities [18]. In order to counterbalance this strong discriminating power, as illustrated in [4] the idea is to consider *success probabilities in a trace-by-trace fashion* rather than on entire resolutions. Since traces come again into play, the idea can be implemented in three different ways according to the three approaches used in Sects. 3 and 4.

In the following, given an NPLTS $\mathcal{L} = (S, A, \longrightarrow)$, $s \in S$, an NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$, $Z \in \text{Res}(s, o)$, and $\alpha \in A^*$, we denote by $\text{CCC}(z_{s, o}, \alpha)$ the set of α -compatible computations from $z_{s, o}$ that are *completed*, i.e., that do not correspond to proper prefixes of computations from (s, o) . Moreover, we denote by $\text{Res}_{\max, \mathcal{C}, \alpha}(s, o)$ the set of resolutions $Z \in \text{Res}_{\max}(s, o)$ such that $\text{CCC}(z_{s, o}, \alpha) \neq \emptyset$ and, for each such resolution, we denote by $\text{SCC}(z_{s, o}, \alpha)$ the set of successful α -compatible computations from $z_{s, o}$.

Definition 5.5. (*Probabilistic trace-by-trace-distribution testing equivalence* – $\sim_{\text{PTe-tbt,dis}}$)

$s_1 \sim_{\text{PTe-tbt,dis}} s_2$ iff for every NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ it holds that for each $Z_1 \in \text{Res}_{\max}(s_1, o)$ there exists $Z_2 \in \text{Res}_{\max}(s_2, o)$ such that for all $\alpha \in A^*$ it holds that $\text{CCC}(z_{s_1, o}, \alpha) \neq \emptyset$ implies $\text{CCC}(z_{s_2, o}, \alpha) \neq \emptyset$ and:

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

and symmetrically for each $Z_2 \in Res_{\max}(s_2, o)$. ■

Definition 5.6. (*Probabilistic trace-by-trace testing equivalence* – $\sim_{\text{PTe-tbt}}$ – [4])

$s_1 \sim_{\text{PTe-tbt}} s_2$ iff for every NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ and for all $\alpha \in A^*$ it holds that for each $Z_1 \in Res_{\max, C, \alpha}(s_1, o)$ there exists $Z_2 \in Res_{\max, C, \alpha}(s_2, o)$ such that:

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

and symmetrically for each $Z_2 \in Res_{\max, C, \alpha}(s_2, o)$. ■

Definition 5.7. (*Probabilistic $\sqcup\sqcap$ -trace-by-trace testing equivalence* – $\sim_{\text{PTe-tbt}, \sqcup\sqcap}$)

$s_1 \sim_{\text{PTe-tbt}, \sqcup\sqcap} s_2$ 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 $Res_{\max, C, \alpha}(s_1, o) \neq \emptyset$ iff $Res_{\max, C, \alpha}(s_2, o) \neq \emptyset$ and:

$$\begin{aligned} \bigsqcup_{Z_1 \in Res_{\max, C, \alpha}(s_1, o)} prob(SCC(z_{s_1, o}, \alpha)) &= \bigsqcup_{Z_2 \in Res_{\max, C, \alpha}(s_2, o)} prob(SCC(z_{s_2, o}, \alpha)) \\ \bigsqcap_{Z_1 \in Res_{\max, C, \alpha}(s_1, o)} prob(SCC(z_{s_1, o}, \alpha)) &= \bigsqcap_{Z_2 \in Res_{\max, C, \alpha}(s_2, o)} prob(SCC(z_{s_2, o}, \alpha)) \end{aligned}$$
■

While $\sim_{\text{PTe}, \sqcup\sqcap}$, $\sim_{\text{PTe}, \forall\exists}$, and $\sim_{\text{PTe-tbt}, \text{dis}}$ are not conservative extensions of the testing equivalence defined in [14] for fully nondeterministic processes – which we denote by $\sim_{\text{Te}, \text{fnd}}$ – and the testing equivalence defined in [10] for fully probabilistic processes – which we denote by $\sim_{\text{Te}, \text{fpr}}$ – the other two testing equivalences are backward compatible with them.

Theorem 5.8. It holds that:

1. $\sim_{\text{PTe-tbt}} = \sim_{\text{PTe-tbt}, \sqcup\sqcap} = \sim_{\text{Te}, \text{fnd}}$ over fully nondeterministic NPLTS models.
2. $\sim_{\text{PTe-tbt}} = \sim_{\text{PTe-tbt}, \sqcup\sqcap} = \sim_{\text{Te}, \text{fpr}}$ over fully probabilistic NPLTS models.

Proof Let (S, A, \longrightarrow) be an NPLTS and $s_1, s_2 \in S$:

1. Suppose that the NPLTS is fully nondeterministic. We preliminarily recall from [14] that $s_1 \sim_{\text{Te}, \text{fnd}} s_2$ means that for every *fully nondeterministic* test $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$:
 - There exists a successful computation from (s_1, o) iff there exists a successful computation from (s_2, o) .
 - All completed computations from (s_1, o) are successful iff all completed computations from (s_2, o) are successful.

The proof is divided into two parts:

- (a) Suppose that $s_1 \sim_{\text{PTe-tbt}} s_2$. Then, in particular, for every fully nondeterministic test $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ and for all $\alpha \in A^*$:
 - For each $Z_1 \in Res_{\max, C, \alpha}(s_1, o)$ there exists $Z_2 \in Res_{\max, C, \alpha}(s_2, o)$ such that:
 $prob(SCC(z_{s_1, o}, \alpha)) = prob(SCC(z_{s_2, o}, \alpha))$
 - For each $Z_2 \in Res_{\max, C, \alpha}(s_2, o)$ there exists $Z_1 \in Res_{\max, C, \alpha}(s_1, o)$ such that:
 $prob(SCC(z_{s_2, o}, \alpha)) = prob(SCC(z_{s_1, o}, \alpha))$

Since the NPLTS under test and the considered tests are all fully nondeterministic, the resulting interaction systems are fully nondeterministic too, and hence their resolutions correspond to their 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 completed α -compatible computation from (s_1, o) there exists a completed α -compatible computation from (s_2, o) such that the two computations are both successful or both unsuccessful.
- For each completed α -compatible computation from (s_2, o) there exists a completed 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 completed computations from (s_1, o) are successful iff all completed computations from (s_2, o) are successful. In fact, assume that all completed computations from, e.g., (s_1, o) are successful. Then at least one completed computation from (s_2, o) is successful. Assume that (s_2, o) has at least two completed computations and that one of them is not successful. Then at least one completed computation from (s_1, o) would not be successful, thus contradicting the assumption that all completed computations from (s_1, o) are successful. Therefore, whenever all completed computations from (s_1, o) are successful, then all completed computations from (s_2, o) are successful. Likewise, whenever all completed computations from (s_2, o) are successful, then all completed computations from (s_1, o) are successful.

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

Suppose now that $s_1 \sim_{\text{Te, fnd}} s_2$ and consider an arbitrary NPT $\mathcal{T} = (O, A, \rightarrow_{\mathcal{T}})$ with initial state $o \in O$, an arbitrary trace $\alpha \in A^*$ such that $\text{Res}_{\max, C, \alpha}(s_1, o) \neq \emptyset$, and an arbitrary resolution $Z_1 \in \text{Res}_{\max, C, \alpha}(s_1, o)$.

Assume that $\text{Res}_{\max, C, \alpha}(s_2, o) = \emptyset$, i.e., assume that for all $Z_2 \in \text{Res}_{\max}(s_2, o)$ it holds that $\text{CCC}(z_{s_2, o}, \alpha) = \emptyset$. Let $\mathcal{T}_{\alpha} = (O, A, \rightarrow_{\mathcal{T}_{\alpha}})$ be a fully nondeterministic test obtained from \mathcal{T} in which (i) only the completed 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}_{α} yields the same α -compatible computations as \mathcal{T} in the interaction systems, the test \mathcal{T}_{α} would violate $s_1 \sim_{\text{Te, fnd}} s_2$ because at least one completed computation from (s_1, o) is successful whilst there are no completed computations from (s_2, o) that are successful. We have thus deduced that, whenever $s_1 \sim_{\text{Te, fnd}} s_2$, then the existence of $Z_1 \in \text{Res}_{\max, C, \alpha}(s_1, o)$ implies the existence of $Z_2 \in \text{Res}_{\max, C, \alpha}(s_2, o)$.

Assume now that for all $Z_2 \in \text{Res}_{\max, C, \alpha}(s_2, o)$ it holds that:

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

Observing that \mathcal{T} must have a successful α -compatible computation – otherwise it would hold that $\text{prob}(\text{SCC}(z_{s_1, o}, \alpha)) = 0 = \text{prob}(\text{SCC}(z_{s_2, o}, \alpha))$ for all $Z_2 \in \text{Res}_{\max, C, \alpha}(s_2, o)$ – from $\text{CCC}(z_{s_1, o}, \alpha) \neq \emptyset$ and $\text{CCC}(z_{s_2, o}, \alpha) = \emptyset$ we derive that $\text{prob}(\text{SCC}(z_{s_1, o}, \alpha)) > 0$ and $\text{prob}(\text{SCC}(z_{s_2, o}, \alpha)) > 0$. Denoting by Z'_1 the element of $\text{Res}_{\max}(s_1)$ that originates Z_1 , we would then have that for each $Z'_2 \in \text{Res}_{\max}(s_2)$ originating Z_2 :

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

where p is the probability of performing a successful α -compatible computation in the element Z of $\text{Res}_{\max}(o)$ that originates Z_1 . However, since the NPLTS under test is fully nondeterministic, Z'_1 and Z'_2 boil down to two α -compatible computations and it holds that:

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

which contradicts what established before.

In conclusion, whenever $s_1 \sim_{\text{Te, fnd}} s_2$, then for each $Z_1 \in \text{Res}_{\max, C, \alpha}(s_1, o)$ there exists $Z_2 \in \text{Res}_{\max, C, \alpha}(s_2, o)$ such that:

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

With a similar argument, we can prove that, whenever $s_1 \sim_{\text{Te, fnd}} s_2$, then for each $Z_2 \in \text{Res}_{\max, C, \alpha}(s_2, o)$ there exists $Z_1 \in \text{Res}_{\max, C, \alpha}(s_1, o)$ such that:

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

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

- (b) Suppose that $s_1 \sim_{\text{PTe-tbt}, \sqcap} s_2$. Then, in particular, for every fully nondeterministic test $\mathcal{T} = (O, A, \rightarrow_{\mathcal{T}})$ with initial state $o \in O$ and for all $\alpha \in A^*$ it holds that $\text{Res}_{\max, C, \alpha}(s_1, o) \neq \emptyset$ iff $\text{Res}_{\max, C, \alpha}(s_2, o) \neq \emptyset$ and:

$$\bigsqcup_{Z_1 \in Res_{\max, C, \alpha}(s_1, o)} prob(SCC(z_{s_1, o}, \alpha)) = \bigsqcup_{Z_2 \in Res_{\max, C, \alpha}(s_2, o)} prob(SCC(z_{s_2, o}, \alpha))$$

$$\prod_{Z_1 \in Res_{\max, C, \alpha}(s_1, o)} prob(SCC(z_{s_1, o}, \alpha)) = \prod_{Z_2 \in Res_{\max, C, \alpha}(s_2, o)} prob(SCC(z_{s_2, o}, \alpha))$$

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

- (\sqcup) There exists a successful α -compatible computation from (s_1, o) iff there exists a successful α -compatible computation from (s_2, o) .
- (\prod) All completed α -compatible computations from (s_1, o) are successful iff all completed α -compatible computations from (s_2, o) are successful.

From this, we immediately derive that:

- (\sqcup) There exists a successful computation from (s_1, o) iff there exists a successful computation from (s_2, o) .
- (\prod) All completed computations from (s_1, o) are successful iff all completed computations from (s_2, o) are successful.

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

Suppose now that $s_1 \sim_{Te, \text{fnd}} s_2$. Then $s_1 \sim_{PTe, \text{tbt}} s_2$ – as we have proved in the first part – and hence $s_1 \sim_{PTe, \text{tbt}, \sqcup, \prod} s_2$ – as a consequence of Thm. 5.9.

2. Suppose that \mathcal{L} is fully probabilistic. We preliminarily recall from [10] that $s_1 \sim_{Te, \text{fpr}} s_2$ means that for every *fully probabilistic* test $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$:

$$prob(SC(s_1, o)) = prob(SC(s_2, o))$$

The proof is divided into two parts:

- (a) Suppose that $s_1 \sim_{PTe, \text{tbt}} s_2$. Then, in particular, for every fully probabilistic test $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ and for all $\alpha \in A^*$:
 - For each $Z_1 \in Res_{\max, C, \alpha}(s_1, o)$ there exists $Z_2 \in Res_{\max, C, \alpha}(s_2, o)$ such that:
 $prob(SCC(z_{s_1, o}, \alpha)) = prob(SCC(z_{s_2, o}, \alpha))$
 - For each $Z_2 \in Res_{\max, C, \alpha}(s_2, o)$ there exists $Z_1 \in Res_{\max, C, \alpha}(s_1, o)$ such that:
 $prob(SCC(z_{s_2, o}, \alpha)) = prob(SCC(z_{s_1, o}, \alpha))$

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 $\alpha \in A^*$:

$$prob(SCC((s_1, o), \alpha)) = prob(SCC((s_2, o), \alpha))$$

From this, we immediately derive that:

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

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

Suppose now that $s_1 \sim_{Te, \text{fpr}} s_2$ 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 $Res_{\max, C, \alpha}(s_1, o) \neq \emptyset$, and an arbitrary resolution $Z_1 \in Res_{\max, C, \alpha}(s_1, o)$.

Assume that $Res_{\max, C, \alpha}(s_2, o) = \emptyset$, i.e., assume that for all $Z_2 \in Res_{\max}(s_2, o)$ it holds that $CCC(z_{s_2, o}, \alpha) = \emptyset$. Let $\mathcal{T}_{\alpha} = (O, A, \longrightarrow_{\mathcal{T}_{\alpha}})$ be a fully probabilistic test obtained from \mathcal{T} in which (i) only the completed 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}_α yields at least one of the α -compatible computations of \mathcal{T} in the interaction systems, the test \mathcal{T}_α would violate $s_1 \sim_{\text{Te}, \text{fpr}} s_2$ because at least one completed computation from (s_1, o) is successful whilst there are no completed computations from (s_2, o) that are successful. We have thus deduced that, whenever $s_1 \sim_{\text{Te}, \text{fpr}} s_2$, then the existence of $\mathcal{Z}_1 \in \text{Res}_{\text{max}, \mathcal{C}, \alpha}(s_1, o)$ implies the existence of $\mathcal{Z}_2 \in \text{Res}_{\text{max}, \mathcal{C}, \alpha}(s_2, o)$.

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

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

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

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

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

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

where:

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

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

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

which contradicts what established before.

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

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

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

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

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

- (b) Suppose that $s_1 \sim_{\text{PTe-tbt}, \sqcup \square} s_2$. Then, in particular, for every fully probabilistic test $\mathcal{T} = (O, A, \rightarrow_{\mathcal{T}})$ with initial state $o \in O$ and for all $\alpha \in A^*$ it holds that $\text{Res}_{\text{max}, \mathcal{C}, \alpha}(s_1, o) \neq \emptyset$ iff $\text{Res}_{\text{max}, \mathcal{C}, \alpha}(s_2, o) \neq \emptyset$ and:

$$\begin{aligned} \bigsqcup_{\mathcal{Z}_1 \in \text{Res}_{\text{max}, \mathcal{C}, \alpha}(s_1, o)} \text{prob}(\text{SCC}(z_{s_1, o}, \alpha)) &= \bigsqcup_{\mathcal{Z}_2 \in \text{Res}_{\text{max}, \mathcal{C}, \alpha}(s_2, o)} \text{prob}(\text{SCC}(z_{s_2, o}, \alpha)) \\ \prod_{\mathcal{Z}_1 \in \text{Res}_{\text{max}, \mathcal{C}, \alpha}(s_1, o)} \text{prob}(\text{SCC}(z_{s_1, o}, \alpha)) &= \prod_{\mathcal{Z}_2 \in \text{Res}_{\text{max}, \mathcal{C}, \alpha}(s_2, o)} \text{prob}(\text{SCC}(z_{s_2, o}, \alpha)) \end{aligned}$$

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 extremal probability values over maximal resolutions can be rephrased by saying that for all $\alpha \in A^*$:

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

From this, we immediately derive that:

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

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

Suppose now that $s_1 \sim_{\text{Te}, \text{fpr}} s_2$. Then $s_1 \sim_{\text{PTe-tbt}} s_2$ – as we have proved in the first part – and hence $s_1 \sim_{\text{PTe-tbt}, \sqcup \square} s_2$ – as a consequence of Thm. 5.9. \blacksquare




Figure 8: Two NPLTS models not distinguishable by $\sim_{P_{Te-tbt}}/\sim_{P_{Te-tbt}, \sqcup}$ through fully nondeterministic tests if all maximal resolutions were considered

We remark that in Defs. 5.5 to 5.7 the considered maximal resolutions are those having at least one α -computation that corresponds to a completed α -computation in the interaction system. The motivation behind this restriction is that it is not appropriate to match the 0 success probability of unsuccessful completed α -computations with the 0 success probability of α -computations that are not completed, as may happen when considering Res_{\max} instead of $Res_{\max, C, \alpha}$. Admitting all maximal resolutions would cause $\sim_{P_{Te-tbt}}$ and $\sim_{P_{Te-tbt}, \sqcup}$ not to be backward compatible with $\sim_{Te, \text{fnd}}$ when restricting attention to fully nondeterministic tests. For example, consider the two fully nondeterministic processes in Fig. 8. They are distinguished by the fully nondeterministic test in the same figure. Following the terminology of [14], the second process must pass the test, while the first one is not able to do so because the interaction system has a completed a -computation not reaching success. In the setting of $\sim_{P_{Te-tbt}}$ and $\sim_{P_{Te-tbt}, \sqcup}$, that completed a -computation in the first interaction system is not matched by any a -computation in the second interaction system because of the restriction to $Res_{\max, C, \alpha}$ – thus correctly distinguishing the two processes – but would be matched by a non-completed a -computation in the second interaction system under Res_{\max} .

We now investigate the relationships of the five testing equivalences among themselves (first two properties below) and of the three trace-by-trace testing equivalences with the three failure equivalences defined in Sect. 4 and the three trace equivalences defined in Sect. 3 (last three properties below). It turns out that $\sim_{P_{Te-\forall\exists}}$ and $\sim_{P_{Te-tbt}, \text{dis}}$ perform exactly the same identifications. Unlike the fully nondeterministic spectrum – where the testing semantics coincides with the failure semantics when all actions are visible [13] – here $\sim_{P_{Te-tbt}, \text{dis}}$ is finer than $\sim_{PF, \text{dis}}$ while $\sim_{P_{Te-tbt}}$ and $\sim_{P_{Te-tbt}, \sqcup}$ are coarser than \sim_{PF} and $\sim_{PF, \sqcup}$, respectively. It also turns out that $\sim_{P_{Te-\sqcup}}$ is incomparable with most of the other equivalences.

Theorem 5.9. It holds that:

1. $\sim_{P_{Te-\forall\exists}} \subseteq \sim_{P_{Te-\sqcup}} \subseteq \sim_{P_{Tr}, \sqcup}$.
2. $\sim_{P_{Te-\forall\exists}} = \sim_{P_{Te-tbt}, \text{dis}} \subseteq \sim_{P_{Te-tbt}} \subseteq \sim_{P_{Te-tbt}, \sqcup}$.
3. $\sim_{P_{Te-tbt}, \text{dis}} \subseteq \sim_{PF, \text{dis}}$.
4. $\sim_{PF} \subseteq \sim_{P_{Te-tbt}} \subseteq \sim_{P_{Tr}}$.
5. $\sim_{PF, \sqcup} \subseteq \sim_{P_{Te-tbt}, \sqcup} \subseteq \sim_{P_{Tr}, \sqcup}$.

Proof Let (S, A, \longrightarrow) be an NPLTS and $s_1, s_2 \in S$:

1. Suppose that $s_1 \sim_{P_{Te-\forall\exists}} s_2$. Then we immediately derive that for every NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$:

$$\{prob(\mathcal{SC}(z_{s_1, o})) \mid \mathcal{Z}_1 \in Res_{\max}(s_1, o)\} \subseteq \{prob(\mathcal{SC}(z_{s_2, o})) \mid \mathcal{Z}_2 \in Res_{\max}(s_2, o)\}$$
 and:

$$\{prob(\mathcal{SC}(z_{s_2, o})) \mid \mathcal{Z}_2 \in Res_{\max}(s_2, o)\} \subseteq \{prob(\mathcal{SC}(z_{s_1, o})) \mid \mathcal{Z}_1 \in Res_{\max}(s_1, o)\}$$
 As a consequence:

$$\{prob(\mathcal{SC}(z_{s_1, o})) \mid \mathcal{Z}_1 \in Res_{\max}(s_1, o)\} = \{prob(\mathcal{SC}(z_{s_2, o})) \mid \mathcal{Z}_2 \in Res_{\max}(s_2, o)\}$$
 and hence:

$$\begin{aligned} \bigsqcup_{Z_1 \in \text{Res}_{\max}(s_1, o)} \text{prob}(\text{SC}(z_{s_1, o})) &= \bigsqcup_{Z_2 \in \text{Res}_{\max}(s_2, o)} \text{prob}(\text{SC}(z_{s_2, o})) \\ \prod_{Z_1 \in \text{Res}_{\max}(s_1, o)} \text{prob}(\text{SC}(z_{s_1, o})) &= \prod_{Z_2 \in \text{Res}_{\max}(s_2, o)} \text{prob}(\text{SC}(z_{s_2, o})) \end{aligned}$$

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

Suppose now that $s_1 \sim_{\text{PTe-}\sqcup} s_2$ and consider an arbitrary trace $\alpha \in A^*$ and an 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 α and reaches ω . Given $s \in S$ and $Z \in \text{Res}_{\max}(s, o)$, due to the structure of \mathcal{T}_α it holds that:

$$\text{prob}(\text{SC}(z_{s, o})) = \text{prob}(\text{SCC}(z_{s, o}, \alpha)) = \text{prob}(\text{CC}(z'_s, \alpha))$$

where Z' is the element of $\text{Res}_\alpha(s)$ that generates Z . Therefore, from $s_1 \sim_{\text{PTe-}\sqcup} s_2$ it follows that:

$$\begin{aligned} \bigsqcup_{Z'_1 \in \text{Res}_\alpha(s_1)} \text{prob}(\text{CC}(z'_{s_1}, \alpha)) &= \bigsqcup_{Z'_2 \in \text{Res}_\alpha(s_2)} \text{prob}(\text{CC}(z'_{s_2}, \alpha)) \\ \prod_{Z'_1 \in \text{Res}_\alpha(s_1)} \text{prob}(\text{CC}(z'_{s_1}, \alpha)) &= \prod_{Z'_2 \in \text{Res}_\alpha(s_2)} \text{prob}(\text{CC}(z'_{s_2}, \alpha)) \end{aligned}$$

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

2. Let us prove the contrapositive of $s_1 \sim_{\text{PTe-}\forall} s_2 \implies s_1 \sim_{\text{PTe-tbt, dis}} s_2$. Suppose that $s_1 \not\sim_{\text{PTe-tbt, dis}} s_2$. Then there exist an NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ and, say, a resolution $Z_1 \in \text{Res}_{\max}(s_1, o)$ such that for each $Z_2 \in \text{Res}_{\max}(s_2, o)$ there is $\alpha_2 \in A^*$ such that $\text{CCC}(z_{s_1, o}, \alpha_2) \neq \emptyset$ and (i) $\text{CCC}(z_{s_2, o}, \alpha_2) = \emptyset$ or (ii) $\text{prob}(\text{SCC}(z_{s_1, o}, \alpha_2)) \neq \text{prob}(\text{SCC}(z_{s_2, o}, \alpha_2))$. We show that from this fact it follows that $s_1 \not\sim_{\text{PTe-}\forall} s_2$ 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 state, finitely branching, and acyclic – 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 $\text{CCC}(z_{s_1, o}, \alpha) \neq \emptyset$ and (i) $\text{CCC}(z_{s_2, o}, \alpha) = \emptyset$ in which case:

$$\text{prob}(\text{SC}(z_{s_1, o})) > 0 = \text{prob}(\text{SC}(z_{s_2, o}))$$

or (ii) it holds that:

$$\begin{aligned} \text{prob}(\text{SC}(z_{s_1, o})) &= \text{prob}(\text{SCC}(z_{s_1, o}, \alpha)) \neq \\ &\neq \text{prob}(\text{SCC}(z_{s_2, o}, \alpha)) = \text{prob}(\text{SC}(z_{s_2, o})) \end{aligned}$$

As a consequence, in both cases $s_1 \not\sim_{\text{PTe-}\forall} s_2$.

- Let $n \in \mathbb{N}_{>1}$ and suppose that the result holds for all $m = 1, \dots, n-1$. 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 completed 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 completed 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 $s_1 \not\sim_{\text{PTe-}\forall} s_2$.

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 $Z \in \text{Res}_{\max}(s, o)$ it holds that:

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

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

- For each $Z_1 \in \text{Res}_{\max}(s_1, o)$ there exists $Z_2 \in \text{Res}_{\max}(s_2, o)$ such that:
$$\begin{aligned} \text{prob}(\text{SC}(z_{s_1, o})) &= \sum_{\alpha \in A^* \text{ s.t. } \text{CCC}(z_{s_1, o}, \alpha) \neq \emptyset} \text{prob}(\text{SCC}(z_{s_1, o}, \alpha)) = \\ &= \sum_{\alpha \in A^* \text{ s.t. } \text{CCC}(z_{s_2, o}, \alpha) \neq \emptyset} \text{prob}(\text{SCC}(z_{s_2, o}, \alpha)) = \text{prob}(\text{SC}(z_{s_2, o})) \end{aligned}$$
- For each $Z_2 \in \text{Res}_{\max}(s_2, o)$ there exists $Z_1 \in \text{Res}_{\max}(s_1, o)$ such that:

$$\begin{aligned}
\text{prob}(\mathcal{SC}(z_{s_2,o})) &= \sum_{\alpha \in A^* \text{ s.t. } \mathcal{CCC}(z_{s_2,o},\alpha) \neq \emptyset} \text{prob}(\mathcal{SCC}(z_{s_2,o},\alpha)) = \\
&= \sum_{\alpha \in A^* \text{ s.t. } \mathcal{CCC}(z_{s_1,o},\alpha) \neq \emptyset} \text{prob}(\mathcal{SCC}(z_{s_1,o},\alpha)) = \text{prob}(\mathcal{SC}(z_{s_1,o}))
\end{aligned}$$

This means that $s_1 \sim_{\text{PTe-}\forall\exists} s_2$. In conclusion, $\sim_{\text{PTe-}\forall\exists} = \sim_{\text{PTe-tbt,dis}}$.

The fact that $s_1 \sim_{\text{PTe-tbt,dis}} s_2$ implies $s_1 \sim_{\text{PTe-tbt}} s_2$ is easily seen by taking the same fully matching resolutions considered in $\sim_{\text{PTe-tbt,dis}}$.

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

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

This is to say that:

- $\text{Res}_{\max,C,\alpha}(s_1,o) \neq \emptyset$ implies $\text{Res}_{\max,C,\alpha}(s_2,o) \neq \emptyset$ and:

$$\bigcup_{\mathcal{Z}_1 \in \text{Res}_{\max,C,\alpha}(s_1,o)} \{\text{prob}(\mathcal{SCC}(z_{s_1,o},\alpha))\} \subseteq \bigcup_{\mathcal{Z}_2 \in \text{Res}_{\max,C,\alpha}(s_2,o)} \{\text{prob}(\mathcal{SCC}(z_{s_2,o},\alpha))\}$$
- $\text{Res}_{\max,C,\alpha}(s_2,o) \neq \emptyset$ implies $\text{Res}_{\max,C,\alpha}(s_1,o) \neq \emptyset$ and:

$$\bigcup_{\mathcal{Z}_2 \in \text{Res}_{\max,C,\alpha}(s_2,o)} \{\text{prob}(\mathcal{SCC}(z_{s_2,o},\alpha))\} \subseteq \bigcup_{\mathcal{Z}_1 \in \text{Res}_{\max,C,\alpha}(s_1,o)} \{\text{prob}(\mathcal{SCC}(z_{s_1,o},\alpha))\}$$

Equivalently, $\text{Res}_{\max,C,\alpha}(s_1,o) \neq \emptyset$ iff $\text{Res}_{\max,C,\alpha}(s_2,o) \neq \emptyset$ and:

$$\bigcup_{\mathcal{Z}_1 \in \text{Res}_{\max,C,\alpha}(s_1,o)} \{\text{prob}(\mathcal{SCC}(z_{s_1,o},\alpha))\} = \bigcup_{\mathcal{Z}_2 \in \text{Res}_{\max,C,\alpha}(s_2,o)} \{\text{prob}(\mathcal{SCC}(z_{s_2,o},\alpha))\}$$

which implies:

$$\begin{aligned}
\bigcup_{\mathcal{Z}_1 \in \text{Res}_{\max,C,\alpha}(s_1,o)} \text{prob}(\mathcal{SCC}(z_{s_1,o},\alpha)) &= \bigcup_{\mathcal{Z}_2 \in \text{Res}_{\max,C,\alpha}(s_2,o)} \text{prob}(\mathcal{SCC}(z_{s_2,o},\alpha)) \\
\bigcap_{\mathcal{Z}_1 \in \text{Res}_{\max,C,\alpha}(s_1,o)} \text{prob}(\mathcal{SCC}(z_{s_1,o},\alpha)) &= \bigcap_{\mathcal{Z}_2 \in \text{Res}_{\max,C,\alpha}(s_2,o)} \text{prob}(\mathcal{SCC}(z_{s_2,o},\alpha))
\end{aligned}$$

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

3. The proof of $\sim_{\text{PTe-tbt,dis}} \subseteq \sim_{\text{PF,dis}}$ is divided into three steps and requires the introduction of some auxiliary definitions and notations.

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) . We say that $c \in \mathcal{C}_{\text{fin}}(z_s)$ 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 the NPLTS departing from $\text{corr}_{\mathcal{Z}}(z_i)$ is precisely R_i . We denote by $\mathcal{RTCC}(z_s, \rho)$ the set of ρ -compatible computations from z_s .

We let $s_1 \sim_{\text{PTr,dis}} s_2$ iff for each resolution $\mathcal{Z}_1 \in \text{Res}(s_1)$ there exists a resolution $\mathcal{Z}_2 \in \text{Res}(s_2)$ such that for all $\rho \in (A \times 2^A)^*$:

$$\text{prob}(\mathcal{RTCC}(z_{s_1}, \rho)) = \text{prob}(\mathcal{RTCC}(z_{s_2}, \rho))$$

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

In the first step of this proof, we show that $s_1 \sim_{\text{PTe-tbt,dis}} s_2$ implies $s_1 \sim_{\text{PTr,dis}} s_2$ 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 $(S_r, A_r, \longrightarrow_r)$ that is isomorphic to the given one up to transition labels and terminal states. A transition $s \xrightarrow{a} \mathcal{D}$ becomes $s_r \xrightarrow{a \triangleleft R}_r \mathcal{D}_r$ where $R \subseteq A$ is the set of actions labeling the outgoing transitions of s and $\mathcal{D}_r(s_r) = \mathcal{D}(s)$ for all $s \in S$. If s is a terminal state, i.e., it has no outgoing transitions, then we add a transition $s_r \xrightarrow{o \triangleleft \emptyset}_r \delta_{s_r}$ where $\delta_{s_r}(s_r) = 1$ and $\delta_{s_r}(s'_r) = 0$ for all $s' \in S \setminus \{s\}$. Transition relabeling preserves $\sim_{\text{PTe-tbt,dis}}$, i.e., $s_1 \sim_{\text{PTe-tbt,dis}} s_2$ implies $s_{1,r} \sim_{\text{PTe-tbt,dis}} s_{2,r}$, because $\sim_{\text{PTe-tbt,dis}}$ 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 $\alpha_r \in (A_r)^*$ and $R \subseteq A$, we build an NPT $\mathcal{T}_{\alpha_r, R} = (O_{\alpha_r, R}, A_r, \longrightarrow_{\alpha_r, R})$ having a single α_r -compatible computation that goes from the initial state $o_{\alpha_r, R}$ 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 $\sim_{\text{PTe-tbt,dis}}$ 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 $\mathcal{T} = (O, A_r, \longrightarrow_{\mathcal{T}})$ that has an initial τ -transition and then behaves as one of the tests $\mathcal{T}_{\alpha_r, R}$, i.e., its initial τ -transition goes from the initial state o to a state distribution whose support is the set $\{o_{\alpha_r, R} \mid \alpha_r \in (A_r)^* \wedge R \subseteq A\}$, with the probability $p_{\alpha_r, R}$ associated with $o_{\alpha_r, R}$ 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 $\alpha_{\rho, r} \in (A_r)^*$ 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 $\mathcal{Z} \in \text{Res}(s)$ it holds that:

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

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

$$\text{prob}(\mathcal{RTCC}(z_s, \rho)) = \text{prob}(\text{SCC}(z_{s_1, o}, \alpha_{\rho, r})) / p_{\alpha_{\rho, r}, R_n}$$

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

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

- For each $\mathcal{Z}_1 \in \text{Res}(s_1)$ there exists $\mathcal{Z}_2 \in \text{Res}(s_2)$ such that for all ready traces $\rho = (a_1, R_1) \dots (a_n, R_n) \in (A \times 2^A)^*$ either:

$$\text{prob}(\mathcal{RTCC}(z_{s_1}, \rho)) = 0 = \text{prob}(\mathcal{RTCC}(z_{s_2}, \rho))$$

or:

$$\begin{aligned} \text{prob}(\mathcal{RTCC}(z_{s_1}, \rho)) &= \text{prob}(\text{SCC}(z_{s_{1, r}, o}, \alpha_{\rho, r})) / p_{\alpha_{\rho, r}, R_n} = \\ &= \text{prob}(\text{SCC}(z_{s_{2, r}, o}, \alpha_{\rho, r})) / p_{\alpha_{\rho, r}, R_n} = \text{prob}(\mathcal{RTCC}(z_{s_2}, \rho)) \end{aligned}$$

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

This means that $s_1 \sim_{\text{PTr,dis}} s_2$.

We can now start the second step of this proof. 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) . We say that $c \in \mathcal{C}_{\text{fin}}(z_s)$ 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 $i = 1, \dots, n$, $\text{corr}_{\mathcal{Z}}(z_i)$ has no outgoing transitions in the NPLTS 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 let $s_1 \sim_{\text{PFTTr,dis}} s_2$ iff for each resolution $\mathcal{Z}_1 \in \text{Res}(s_1)$ there exists a resolution $\mathcal{Z}_2 \in \text{Res}(s_2)$ such that for all $\phi \in (A \times 2^A)^*$:

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

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

In this second step, we prove that $s_1 \sim_{\text{PTr,dis}} s_2$ implies $s_1 \sim_{\text{PFTTr,dis}} s_2$. Suppose that $s_1 \sim_{\text{PTr,dis}} s_2$. Then we immediately derive that:

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

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

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

This means that $s_1 \sim_{\text{PFTTr,dis}} s_2$.

The third and final step of this proof shows that $s_1 \sim_{\text{PFTTr,dis}} s_2$ implies $s_1 \sim_{\text{PF,dis}} s_2$. Suppose that $s_1 \sim_{\text{PFTTr,dis}} s_2$. Then we immediately derive that:

- For each $\mathcal{Z}_1 \in \text{Res}(s_1)$ there exists $\mathcal{Z}_2 \in \text{Res}(s_2)$ such that for all $(a_1 \dots a_n, F) \in A^* \times 2^A$:
$$\begin{aligned}
\text{prob}(\mathcal{FCC}(z_{s_1}, (a_1 \dots a_n, F))) &= \text{prob}(\mathcal{FTCC}(z_{s_1}, (a_1, \emptyset) \dots (a_{n-1}, \emptyset)(a_n, F))) \\
&= \text{prob}(\mathcal{FTCC}(z_{s_2}, (a_1, \emptyset) \dots (a_{n-1}, \emptyset)(a_n, F))) \\
&= \text{prob}(\mathcal{FCC}(z_{s_2}, (a_1 \dots a_n, F)))
\end{aligned}$$
- Symmetrically for each $\mathcal{Z}_2 \in \text{Res}(s_2)$.

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

- Let us prove the contrapositive of $s_1 \sim_{\text{PF}} s_2 \implies s_1 \sim_{\text{PTE-tbt}} s_2$. Suppose that $s_1 \not\sim_{\text{PTE-tbt}} s_2$. 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 $\mathcal{Z}_1 \in \text{Res}_{\text{max,C},\alpha}(s_1, o)$ such that $\text{Res}_{\text{max,C},\alpha}(s_2, o) = \emptyset$ or for all $\mathcal{Z}_2 \in \text{Res}_{\text{max,C},\alpha}(s_2, o)$ it holds that:

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

Observing that $\text{Res}_{\text{max,C},\alpha}(s_1, o) \neq \emptyset$, in the case that $\text{Res}_{\text{max,C},\alpha}(s_2, o) = \emptyset$ 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 $s_1 \not\sim_{\text{PF}} s_2$ in this case because, denoting by \mathcal{Z}'_1 the element of $\text{Res}(s_1)$ that originates \mathcal{Z}_1 , we have that for all $\mathcal{Z}'_2 \in \text{Res}(s_2)$:

$$\text{prob}(\mathcal{FCC}(z'_{s_1}, \varphi)) > 0 = \text{prob}(\mathcal{FCC}(z'_{s_2}, \varphi))$$

In the case that $\text{Res}_{\text{max,C},\alpha}(s_2, o) \neq \emptyset$, the failure pair $\varphi = (\alpha, \emptyset)$ shows that $s_1 \not\sim_{\text{PF}} s_2$. 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 $\text{Res}(o)$ – as the only effect of this assumption is that of possibly reducing the number of resolutions in $\text{Res}_{\text{max,C},\alpha}(s_2, o)$. 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 $\text{Res}(s_1)$ that originates \mathcal{Z}_1 , we then have that for all $\mathcal{Z}'_2 \in \text{Res}(s_2)$ originating some $\mathcal{Z}_2 \in \text{Res}_{\text{max,C},\alpha}(s_2, o)$:

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

where p is the probability of performing the α -compatible computations in the only element \mathcal{Z} of $\text{Res}(o)$ that originates \mathcal{Z}_1 and all the resolutions \mathcal{Z}_2 .

Suppose now that $s_1 \sim_{\text{PTE-tbt}} s_2$. Then, in particular, for every $\alpha \in A^*$ and NPT $\mathcal{T}_\alpha = (O, A, \longrightarrow_{\mathcal{T}_\alpha})$ with initial state $o \in O$ having a single maximal α -compatible computation that reaches success, it holds that:

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

Since for all $s \in S$, $\mathcal{Z} \in \text{Res}_{\text{max,C},\alpha}(s, o)$, and $\mathcal{Z}' \in \text{Res}(s)$ originating \mathcal{Z} in the interaction with \mathcal{T}_α it holds that:

$$\text{prob}(\text{SCC}(z_{s,o}, \alpha)) = \text{prob}(\text{CC}(z'_s, \alpha))$$

due to the structure of \mathcal{T}_α , we immediately derive that for all $\alpha \in A^*$:

- For each $Z'_1 \in \text{Res}(s_1)$ there exists $Z'_2 \in \text{Res}(s_2)$ such that:
 $\text{prob}(\text{CC}(z'_{s_1}, \alpha)) = \text{prob}(\text{CC}(z'_{s_2}, \alpha))$
- For each $Z'_2 \in \text{Res}(s_2)$ there exists $Z'_1 \in \text{Res}(s_1)$ such that:
 $\text{prob}(\text{CC}(z'_{s_2}, \alpha)) = \text{prob}(\text{CC}(z'_{s_1}, \alpha))$

This means that $s_1 \sim_{\text{PTT}} s_2$.

5. Let us prove the contrapositive of $s_1 \sim_{\text{PF}, \sqcup} s_2 \implies s_1 \sim_{\text{PTE-tbt}, \sqcup} s_2$. Suppose that $s_1 \not\sim_{\text{PTE-tbt}, \sqcup} s_2$. This means that there exist an NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ and a trace $\alpha \in A^*$ such that, for instance, $\text{Res}_{\max, C, \alpha}(s_1, o) \neq \emptyset$ and:

$$\bigsqcup_{Z_1 \in \text{Res}_{\max, C, \alpha}(s_1, o)} \text{prob}(\text{SCC}(z_{s_1, o}, \alpha)) > \bigsqcup_{Z_2 \in \text{Res}_{\max, C, \alpha}(s_2, o)} \text{prob}(\text{SCC}(z_{s_2, o}, \alpha))$$

Since $\text{Res}_{\max, C, \alpha}(s_1, o) \neq \emptyset$, in the case that $\text{Res}_{\max, C, \alpha}(s_2, o) = \emptyset$ 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 $s_1 \not\sim_{\text{PF}, \sqcup} s_2$ in this case because, denoting by Z'_1 the element of $\text{Res}_{\alpha}(s_1)$ that originates some $Z_1 \in \text{Res}_{\max, C, \alpha}(s_1, o)$, we have that:

$$\bigsqcup_{Z'_1 \in \text{Res}_{\alpha}(s_1)} \text{prob}(\text{FCC}(z'_{s_1}, \varphi)) > 0 = \bigsqcup_{Z'_2 \in \text{Res}_{\alpha}(s_2)} \text{prob}(\text{FCC}(z'_{s_2}, \varphi))$$

In the case that $\text{Res}_{\max, C, \alpha}(s_2, o) \neq \emptyset$, the failure pair $\varphi = (\alpha, \emptyset)$ shows that $s_1 \not\sim_{\text{PF}, \sqcup} s_2$. In fact, without loss of generality we can assume that the only α -compatible computations in \mathcal{T} are the ones resulting in the supremum of the success probabilities of the α -compatible computations from (s_1, o) – note that they must belong to the same element Z of $\text{Res}(o)$ – as this assumption has no effect on the relationship between the two suprema. 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 Z'_1 the element of $\text{Res}_{\alpha}(s_1)$ that originates some $Z_1 \in \text{Res}_{\max, C, \alpha}(s_1, o)$ and by Z'_2 the element of $\text{Res}_{\alpha}(s_2)$ that originates some $Z_2 \in \text{Res}_{\max, C, \alpha}(s_2, o)$, we then have that:

$$\begin{aligned} \bigsqcup_{Z'_1 \in \text{Res}_{\alpha}(s_1)} \text{prob}(\text{FCC}(z'_{s_1}, \varphi)) &= \bigsqcup_{Z_1 \in \text{Res}_{\max, C, \alpha}(s_1, o)} \text{prob}(\text{SCC}(z_{s_1, o}, \alpha)) / p \\ &> \bigsqcup_{Z_2 \in \text{Res}_{\max, C, \alpha}(s_2, o)} \text{prob}(\text{SCC}(z_{s_2, o}, \alpha)) / p \\ &= \bigsqcup_{Z'_2 \in \text{Res}_{\alpha}(s_2)} \text{prob}(\text{FCC}(z'_{s_2}, \varphi)) \end{aligned}$$

where p is the probability of performing the α -compatible computations in the only element Z of $\text{Res}(o)$ that originates all the resolutions Z_1 and Z_2 .

Suppose now that $s_1 \sim_{\text{PTE-tbt}, \sqcup} s_2$. Then, in particular, for every $\alpha \in A^*$ and NPT $\mathcal{T}_{\alpha} = (O, A, \longrightarrow_{\mathcal{T}_{\alpha}})$ with initial state $o \in O$ having a single maximal α -compatible computation that reaches success, it holds that $\text{Res}_{\max, C, \alpha}(s_1, o) \neq \emptyset$ iff $\text{Res}_{\max, C, \alpha}(s_2, o) \neq \emptyset$ and:

$$\begin{aligned} \bigsqcup_{Z_1 \in \text{Res}_{\max, C, \alpha}(s_1, o)} \text{prob}(\text{SCC}(z_{s_1, o}, \alpha)) &= \bigsqcup_{Z_2 \in \text{Res}_{\max, C, \alpha}(s_2, o)} \text{prob}(\text{SCC}(z_{s_2, o}, \alpha)) \\ \prod_{Z_1 \in \text{Res}_{\max, C, \alpha}(s_1, o)} \text{prob}(\text{SCC}(z_{s_1, o}, \alpha)) &= \prod_{Z_2 \in \text{Res}_{\max, C, \alpha}(s_2, o)} \text{prob}(\text{SCC}(z_{s_2, o}, \alpha)) \end{aligned}$$

Since for all $s \in S$, $Z \in \text{Res}_{\max, C, \alpha}(s, o)$, and $Z' \in \text{Res}_{\alpha}(s)$ originating Z in the interaction with \mathcal{T}_{α} it holds that:

$$\text{prob}(\text{SCC}(z_{s, o}, \alpha)) = \text{prob}(\text{CC}(z'_s, \alpha))$$

due to the structure of \mathcal{T}_{α} , we immediately derive that for all $\alpha \in A^*$:

$$\begin{aligned} \bigsqcup_{Z'_1 \in \text{Res}_{\alpha}(s_1)} \text{prob}(\text{CC}(z'_{s_1}, \alpha)) &= \bigsqcup_{Z'_2 \in \text{Res}_{\alpha}(s_2)} \text{prob}(\text{CC}(z'_{s_2}, \alpha)) \\ \prod_{Z'_1 \in \text{Res}_{\alpha}(s_1)} \text{prob}(\text{CC}(z'_{s_1}, \alpha)) &= \prod_{Z'_2 \in \text{Res}_{\alpha}(s_2)} \text{prob}(\text{CC}(z'_{s_2}, \alpha)) \end{aligned}$$

This means that $s_1 \sim_{\text{PTT}, \sqcup} s_2$. ■




Figure 9: Two NPLTS models distinguished by $\sim_{PF}/\sim_{PF,\sqcup\sqcap}$ and identified by $\sim_{P_{Te-tbt}}/\sim_{P_{Te-tbt},\sqcup\sqcap}$

All the inclusions in Thm. 5.9 are strict:

- Figure 2 shows that $\sim_{P_{Te-\forall\exists}}$ is strictly finer than $\sim_{P_{Te-\sqcup\sqcap}}$. It holds that $s_1 \not\sim_{P_{Te-\forall\exists}} s_2$ because the test whose initial state o has an *offer*-transition followed by a *win*₁-transition reaching success, results in a 0.5 success probability for the central maximal resolution of (s_1, o) that is not matched by any maximal resolution of (s_2, o) . In contrast, $s_1 \sim_{P_{Te-\sqcup\sqcap}} s_2$ because no test can make the central maximal resolution of s_1 come into play when the emphasis is on the maximal and minimal success probabilities.
- Figure 1 shows that $\sim_{P_{Te-\sqcup\sqcap}}$ is strictly finer than $\sim_{P_{Tr,\sqcup\sqcap}}$. It holds that $s_1 \not\sim_{P_{Te-\sqcup\sqcap}} s_2$ because the test whose initial state o has an *offer*-transition leading (i) with probability 0.9 to a state that has a *draw*-transition reaching success and (ii) with probability 0.1 to a state that has a *win*₁-transition reaching success, results in the 0.54 maximum success probability obtained in the central maximal resolution of (s_1, o) that is different from the 0.58 maximum success probability obtained in the leftmost maximal resolution of (s_2, o) . In contrast, $s_1 \sim_{P_{Tr,\sqcup\sqcap}} s_2$ because each of the various traces has the same maximum and minimum probabilities in both processes.
- Figures 1 and 2 respectively show that $\sim_{P_{Te-tbt,dis}}$ is strictly finer than $\sim_{P_{Te-tbt}}$ and $\sim_{P_{Te-tbt}}$ is strictly finer than $\sim_{P_{Te-tbt},\sqcup\sqcap}$.
- Figure 7 shows that $\sim_{P_{Te-tbt,dis}}$ is strictly finer than $\sim_{PF,dis}$. It holds that $s_1 \not\sim_{P_{Te-tbt,dis}} s_2$ because $\sim_{P_{Te-tbt,dis}}$ coincides with $\sim_{P_{Te-\forall\exists}}$ and we have already exhibited a test distinguishing the two processes with respect to $\sim_{P_{Te-\forall\exists}}$. In contrast, $s_1 \sim_{PF,dis} s_2$ because for each resolution of s_1 (resp. s_2) there exists a resolution of s_2 (resp. s_1) having precisely the same failure distribution.
- Figure 9 shows that \sim_{PF} and $\sim_{PF,\sqcup\sqcap}$ are strictly finer than $\sim_{P_{Te-tbt}}$ and $\sim_{P_{Te-tbt},\sqcup\sqcap}$, respectively. It holds that $s_1 \not\sim_{PF,\sqcup\sqcap} s_2$ (and hence $s_1 \not\sim_{PF} s_2$) because the failure pair (a, A) has maximum probability 1 in the first process and 0.5 in the second process. In contrast, $s_1 \sim_{P_{Te-tbt}} s_2$ (and hence $s_1 \sim_{P_{Te-tbt},\sqcup\sqcap} s_2$) because no test is able to distinguish the two processes due to the fact that success probabilities are computed in a trace-by-trace fashion.
- Figure 10 shows that $\sim_{P_{Te-tbt}}$ and $\sim_{P_{Te-tbt},\sqcup\sqcap}$ are strictly finer than $\sim_{P_{Tr}}$ and $\sim_{P_{Tr,\sqcup\sqcap}}$, respectively. It holds that $s_1 \not\sim_{P_{Te-tbt},\sqcup\sqcap} s_2$ (and hence $s_1 \not\sim_{P_{Te-tbt}} s_2$) because the test in the same figure results in a situation in which $Res_{\max,C,a}(s_1, o) = \emptyset$ while $Res_{\max,C,a}(s_2, o) \neq \emptyset$. In contrast, $s_1 \sim_{P_{Tr}} s_2$ (and hence $s_1 \sim_{P_{Tr,\sqcup\sqcap}} s_2$) because for each resolution of s_1 (resp. s_2) there exists a resolution of s_2 (resp. s_1) having precisely the same trace distribution.

Moreover:

- $\sim_{P_{Te-\sqcup\sqcap}}$ is incomparable with $\sim_{PF,dis}$, $\sim_{P_{Tr,dis}}$, \sim_{PF} , $\sim_{P_{Te-tbt}}$, and $\sim_{P_{Tr}}$, because in Fig. 2 it holds that $s_1 \sim_{P_{Te-\sqcup\sqcap}} s_2$ and $s_1 \not\sim_{P_{Tr}} s_2$ (and hence $s_1 \not\sim_{P_{Te-tbt}} s_2$, $s_1 \not\sim_{PF} s_2$, $s_1 \not\sim_{P_{Tr,dis}} s_2$, and $s_1 \not\sim_{PF,dis} s_2$), while in Fig. 7 it holds that $s_1 \not\sim_{P_{Te-\sqcup\sqcap}} s_2$ and $s_1 \sim_{PF,dis} s_2$ (and hence $s_1 \sim_{P_{Tr,dis}} s_2$, $s_1 \sim_{PF} s_2$, $s_1 \sim_{P_{Te-tbt}} s_2$, and $s_1 \sim_{P_{Tr}} s_2$).




Figure 10: Two NPLTS models distinguished by $\sim_{P_{Te-tbt}}/\sim_{P_{Te-tbt, \sqcup \cap}}$ and identified by $\sim_{P_{Tr}}/\sim_{P_{Tr, \sqcup \cap}}$

- $\sim_{P_{Te-\sqcup \cap}}$ is incomparable with $\sim_{P_{F, \sqcup \cap}}$ and $\sim_{P_{Te-tbt, \sqcup \cap}}$ too. Indeed, in Fig. 11 it holds that $s_1 \sim_{P_{Te-\sqcup \cap}} s_2$ while $s_1 \not\sim_{P_{Te-tbt, \sqcup \cap}} s_2$ (and hence $s_1 \not\sim_{P_{F, \sqcup \cap}} s_2$) due to the test shown in the figure, which yields different minimum success probabilities for trace ab over the maximal resolutions of the two interaction systems having completed computations compatible with that trace. In contrast, in Fig. 1 it holds that $s_1 \not\sim_{P_{Te-\sqcup \cap}} s_2$ and $s_1 \sim_{P_{F, \sqcup \cap}} s_2$ (and hence $s_1 \sim_{P_{Te-tbt, \sqcup \cap}} s_2$).
- $\sim_{P_{Te-tbt}}$ and $\sim_{P_{Te-tbt, \sqcup \cap}}$ are incomparable with $\sim_{P_{Tr, dis}}$, because in Fig. 9 it holds that $s_1 \sim_{P_{Te-tbt}} s_2$ (and hence $s_1 \sim_{P_{Te-tbt, \sqcup \cap}} s_2$) and $s_1 \not\sim_{P_{Tr, dis}} s_2$, while in Fig. 10 it holds that $s_1 \not\sim_{P_{Te-tbt, \sqcup \cap}} s_2$ (and hence $s_1 \not\sim_{P_{Te-tbt}} s_2$) and $s_1 \sim_{P_{Tr, dis}} s_2$.
- $\sim_{P_{Te-tbt}}$ is incomparable with $\sim_{P_{F, \sqcup \cap}}$, because in Fig. 9 it holds that $s_1 \sim_{P_{Te-tbt}} s_2$ and $s_1 \not\sim_{P_{F, \sqcup \cap}} s_2$, while in Fig. 2 it holds that $s_1 \not\sim_{P_{Te-tbt}} s_2$ and $s_1 \sim_{P_{F, \sqcup \cap}} s_2$.

6. Bisimulation Equivalences for NPLTS Models

Bisimulation equivalences capture the ability of two processes of mimicking each other's behavior stepwise. Similar to the trace and failure cases, given two states there are three different approaches to the definition of these bisimilarities, each following the style of [28] based on equivalence relations. The first approach is to match transitions on the basis of *class distributions*, which means that for each transition of one of the two states there must exist an equally labeled transition of the other state such that, *for every equivalence class*, the two transitions have the same probability of reaching a state in that class. In other words, matching transitions of the two states are related by the fully probabilistic version of bisimilarity (fully matching transitions). The second approach is to consider *a single equivalence class at a time*, i.e., to anticipate the quantification over classes with respect to the quantification over transitions. In this way, a transition departing from one of the two states is allowed to be matched, with respect to the probabilities of reaching different classes, by several different transitions departing from the other state (partially matching




Figure 11: Two NPLTS models distinguished by $\sim_{P_{F, \sqcup \cap}}/\sim_{P_{Te-tbt, \sqcup \cap}}$ and identified by $\sim_{P_{Te-\sqcup \cap}}$

transitions). The third approach is to compare only the *extremal probabilities* of reaching each class over all possible transitions labeled with a certain action (max-min-matching transition sets).

Unlike [28], we shall consider *groups of equivalence classes* rather than individual equivalence classes. This does not change the discriminating power in the case of the first approach, while it increases the discriminating power thereby resulting in desirable logical characterizations in the case of the other two approaches [12, 38, 37, 8]. Given an NPLTS (S, A, \longrightarrow) and a distribution $\mathcal{D} \in \text{Distr}(S)$, in the following we let $\mathcal{D}(S') = \sum_{s \in S'} \mathcal{D}(s)$ for $S' \subseteq S$. Moreover, given an equivalence relation \mathcal{B} over S and a group of equivalence classes $\mathcal{G} \in 2^{S/\mathcal{B}}$, we also let $\bigcup \mathcal{G} = \bigcup_{C \in \mathcal{G}} C$.

Definition 6.1. (*Probabilistic group-distribution bisimilarity* – $\sim_{\text{PB,dis}}$ – [34])

$s_1 \sim_{\text{PB,dis}} s_2$ iff (s_1, s_2) belongs to the largest probabilistic group-distribution bisimulation. An equivalence relation \mathcal{B} over S is a *probabilistic group-distribution bisimulation* iff, whenever $(s_1, s_2) \in \mathcal{B}$, then for each $s_1 \xrightarrow{a} \mathcal{D}_1$ there exists $s_2 \xrightarrow{a} \mathcal{D}_2$ such that for all $\mathcal{G} \in 2^{S/\mathcal{B}}$ it holds that $\mathcal{D}_1(\bigcup \mathcal{G}) = \mathcal{D}_2(\bigcup \mathcal{G})$. ■

Definition 6.2. (*Probabilistic bisimilarity* – \sim_{PB} – [8])

$s_1 \sim_{\text{PB}} s_2$ iff (s_1, s_2) belongs to the largest probabilistic bisimulation. An equivalence relation \mathcal{B} over S is a *probabilistic bisimulation* iff, whenever $(s_1, s_2) \in \mathcal{B}$, then for all $\mathcal{G} \in 2^{S/\mathcal{B}}$ it holds that for each $s_1 \xrightarrow{a} \mathcal{D}_1$ there exists $s_2 \xrightarrow{a} \mathcal{D}_2$ such that $\mathcal{D}_1(\bigcup \mathcal{G}) = \mathcal{D}_2(\bigcup \mathcal{G})$. ■

Definition 6.3. (*Probabilistic $\sqcup\sqcap$ -bisimilarity* – $\sim_{\text{PB},\sqcup\sqcap}$ – [8])

$s_1 \sim_{\text{PB},\sqcup\sqcap} s_2$ iff (s_1, s_2) belongs to the largest probabilistic $\sqcup\sqcap$ -bisimulation. An equivalence relation \mathcal{B} over S is a *probabilistic $\sqcup\sqcap$ -bisimulation* iff, whenever $(s_1, s_2) \in \mathcal{B}$, then for all $\mathcal{G} \in 2^{S/\mathcal{B}}$ and $a \in A$ it holds that $s_1 \xrightarrow{a}$ iff $s_2 \xrightarrow{a}$ and:

$$\begin{aligned} \bigcup_{s_1 \xrightarrow{a} \mathcal{D}_1} \mathcal{D}_1(\bigcup \mathcal{G}) &= \bigcup_{s_2 \xrightarrow{a} \mathcal{D}_2} \mathcal{D}_2(\bigcup \mathcal{G}) \\ \bigcap_{s_1 \xrightarrow{a} \mathcal{D}_1} \mathcal{D}_1(\bigcup \mathcal{G}) &= \bigcap_{s_2 \xrightarrow{a} \mathcal{D}_2} \mathcal{D}_2(\bigcup \mathcal{G}) \end{aligned}$$

■

The three bisimulation equivalences defined above are all backward compatible with the bisimulation equivalences respectively defined in [21] for fully nondeterministic processes – which we denote by $\sim_{\text{B,fnd}}$ – and in [19] for fully probabilistic processes – which we denote by $\sim_{\text{B,fpr}}$.

Theorem 6.4. It holds that:

1. $\sim_{\text{PB,dis}} = \sim_{\text{PB}} = \sim_{\text{PB},\sqcup\sqcap} = \sim_{\text{B,fnd}}$ over fully nondeterministic NPLTS models.
2. $\sim_{\text{PB,dis}} = \sim_{\text{PB}} = \sim_{\text{PB},\sqcup\sqcap} = \sim_{\text{B,fpr}}$ over fully probabilistic NPLTS models.

Proof Let (S, A, \longrightarrow) be an NPLTS and $s_1, s_2 \in S$:

1. Suppose that the NPLTS is fully nondeterministic. We preliminarily recall from [21] (and adapt to the NPLTS setting) that $s_1 \sim_{\text{B,fnd}} s_2$ means that there exists an fnd-bisimulation \mathcal{B} over S such that $(s_1, s_2) \in \mathcal{B}$. Let us denote by δ_s the Dirac distribution for $s \in S$, i.e., let $\delta_s(s) = 1$ and $\delta_s(s') = 0$ for all $s' \in S \setminus \{s\}$. A relation \mathcal{B} over S is an fnd-bisimulation iff, whenever $(s'_1, s'_2) \in \mathcal{B}$, then:
 - For each $s'_1 \xrightarrow{a} \delta_{s'_1}$ there exists $s'_2 \xrightarrow{a} \delta_{s'_2}$ such that $(s''_1, s''_2) \in \mathcal{B}$.
 - For each $s'_2 \xrightarrow{a} \delta_{s'_2}$ there exists $s'_1 \xrightarrow{a} \delta_{s'_1}$ such that $(s''_1, s''_2) \in \mathcal{B}$.

The property $\sim_{\text{PB,dis}} = \sim_{\text{PB}} = \sim_{\text{PB},\sqcup\sqcap} = \sim_{\text{B,fnd}}$ over the considered fully nondeterministic NPLTS is a straightforward consequence of the fact that, since in this model every transition can reach with probability greater than 0 only one state and hence only one class of any equivalence relation – which are thus reached with probability 1 – the reflexive, symmetric, and transitive closure of an fnd-bisimulation is trivially a probabilistic group-distribution bisimulation, a probabilistic bisimulation, and a probabilistic $\sqcup\sqcap$ -bisimulation.

- Suppose that the NPLTS is fully probabilistic. We preliminarily recall from [19] (and adapt to the NPLTS setting) that $s_1 \sim_{\text{B}, \text{fpr}} s_2$ means that there exists an fpr-bisimulation \mathcal{B} over S such that $(s_1, s_2) \in \mathcal{B}$. An equivalence relation \mathcal{B} over S is an fpr-bisimulation iff, whenever $(s'_1, s'_2) \in \mathcal{B}$, then for all equivalence classes $C \in S/\mathcal{B}$ it holds for each $s'_1 \xrightarrow{a} \mathcal{D}_1$ there exists $s'_2 \xrightarrow{a} \mathcal{D}_2$ such that $\mathcal{D}_1(C) = \mathcal{D}_2(C)$.

The property $\sim_{\text{PB}, \text{dis}} = \sim_{\text{PB}} = \sim_{\text{PB}, \sqcup \sqcap} = \sim_{\text{B}, \text{fpr}}$ over the considered fully probabilistic NPLTS is a straightforward consequence of the fact that, since in this model every state has at most one outgoing transition, an fpr-bisimulation is trivially a probabilistic group-distribution bisimulation, a probabilistic bisimulation, and a probabilistic $\sqcup \sqcap$ -bisimulation. \blacksquare

We now investigate the relationships of the three bisimulation equivalences among themselves (first property below) and with the five testing equivalences defined in Sect. 5, the three failure equivalences defined in Sect. 4, and the three trace equivalences defined in Sect. 3 (second property below). It turns out that only $\sim_{\text{PB}, \text{dis}}$ is related to the other equivalences, while \sim_{PB} and $\sim_{\text{PB}, \sqcup \sqcap}$ are incomparable with them.

Theorem 6.5. It holds that:

- $\sim_{\text{PB}, \text{dis}} \subseteq \sim_{\text{PB}} \subseteq \sim_{\text{PB}, \sqcup \sqcap}$.
- $\sim_{\text{PB}, \text{dis}} \subseteq \sim_{\text{PTe-tbt}, \text{dis}}$.

Proof Let (S, A, \longrightarrow) be an NPLTS and $s_1, s_2 \in S$:

- The fact that $s_1 \sim_{\text{PB}, \text{dis}} s_2$ implies $s_1 \sim_{\text{PB}} s_2$ is a straightforward consequence of the fact that a probabilistic group-distribution bisimulation is a probabilistic bisimulation, as can be easily seen by taking the same fully matching transitions considered in the group-distribution bisimulation. Suppose now that $s_1 \sim_{\text{PB}} s_2$. This means that there exists a probabilistic bisimulation \mathcal{B} over S such that $(s_1, s_2) \in \mathcal{B}$. In other words, whenever $(s'_1, s'_2) \in \mathcal{B}$, then for all $\mathcal{G} \in 2^{S/\mathcal{B}}$ it holds that:

- For each $s'_1 \xrightarrow{a} \mathcal{D}_1$ there exists $s'_2 \xrightarrow{a} \mathcal{D}_2$ such that $\mathcal{D}_1(\bigcup \mathcal{G}) = \mathcal{D}_2(\bigcup \mathcal{G})$.
- For each $s'_2 \xrightarrow{a} \mathcal{D}_2$ there exists $s'_1 \xrightarrow{a} \mathcal{D}_1$ such that $\mathcal{D}_2(\bigcup \mathcal{G}) = \mathcal{D}_1(\bigcup \mathcal{G})$.

This is to say that, whenever $(s'_1, s'_2) \in \mathcal{B}$, then for all $\mathcal{G} \in 2^{S/\mathcal{B}}$ and $a \in A$ it holds that:

- If $s'_1 \xrightarrow{a}$, then $s'_2 \xrightarrow{a}$ and $\bigcup_{s'_1 \xrightarrow{a} \mathcal{D}_1} \{\mathcal{D}_1(\bigcup \mathcal{G})\} \subseteq \bigcup_{s'_2 \xrightarrow{a} \mathcal{D}_2} \{\mathcal{D}_2(\bigcup \mathcal{G})\}$.
- If $s'_2 \xrightarrow{a}$, then $s'_1 \xrightarrow{a}$ and $\bigcup_{s'_2 \xrightarrow{a} \mathcal{D}_2} \{\mathcal{D}_2(\bigcup \mathcal{G})\} \subseteq \bigcup_{s'_1 \xrightarrow{a} \mathcal{D}_1} \{\mathcal{D}_1(\bigcup \mathcal{G})\}$.

Equivalently, $s'_1 \xrightarrow{a}$ iff $s'_2 \xrightarrow{a}$ and:

$$\bigcup_{s'_1 \xrightarrow{a} \mathcal{D}_1} \{\mathcal{D}_1(\bigcup \mathcal{G})\} = \bigcup_{s'_2 \xrightarrow{a} \mathcal{D}_2} \{\mathcal{D}_2(\bigcup \mathcal{G})\}$$

which implies:

$$\begin{aligned} \bigcup_{s'_1 \xrightarrow{a} \mathcal{D}_1} \mathcal{D}_1(\bigcup \mathcal{G}) &= \bigcup_{s'_2 \xrightarrow{a} \mathcal{D}_2} \mathcal{D}_2(\bigcup \mathcal{G}) \\ \bigcap_{s'_1 \xrightarrow{a} \mathcal{D}_1} \mathcal{D}_1(\bigcup \mathcal{G}) &= \bigcap_{s'_2 \xrightarrow{a} \mathcal{D}_2} \mathcal{D}_2(\bigcup \mathcal{G}) \end{aligned}$$

Therefore, \mathcal{B} is also a probabilistic $\sqcup \sqcap$ -bisimulation, i.e., $s_1 \sim_{\text{PB}, \sqcup \sqcap} s_2$.

- Suppose that $s_1 \sim_{\text{PB}, \text{dis}} s_2$. Notice that states related by $\sim_{\text{PB}, \text{dis}}$ have the same set of actions labeling their outgoing transitions, and that states not enjoying this property are trivially distinguished by $\sim_{\text{PTe-tbt}, \text{dis}}$. Consider an arbitrary NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$. Since $\sim_{\text{PB}, \text{dis}}$ is a congruence with respect to parallel composition [34], for all $s'_1, s'_2 \in S$ such that $s'_1 \sim_{\text{PB}, \text{dis}} s'_2$ and for all $o \in O$ it holds that $(s'_1, o) \sim_{\text{PB}, \text{dis}} (s'_2, o)$ due to some probabilistic group-distribution bisimulation \mathcal{B} over $S \times O$. Since configurations related by $\sim_{\text{PB}, \text{dis}}$ have the same set of actions labeling their outgoing transitions, this induces projections of \mathcal{B} that are fpr-bisimulations [19] over pairs of matching resolutions of the interaction system that are both maximal. As a consequence, whenever $((s'_1, o), (s'_2, o)) \in \mathcal{B}$, then:

- For each $Z_1 \in Res_{\max}(s'_1, o)$ there exists $Z_2 \in Res_{\max}(s'_2, o)$ such that the equivalence relation $\mathcal{B}_{1,2}$ over $Z = Z_1 \cup Z_2$ corresponding to \mathcal{B} projected onto $Z \times Z$ is an fpr-bisimulation, i.e., whenever $(z_{s'_1, o'}, z_{s'_2, o'}) \in \mathcal{B}_{1,2}$, then for each $z_{s'_1, o'} \xrightarrow{a} \mathcal{D}_1$ there exists $z_{s'_2, o'} \xrightarrow{a} \mathcal{D}_2$ such that for all equivalence classes $C \in Z/\mathcal{B}_{1,2}$ it holds that $\mathcal{D}_1(C) = \mathcal{D}_2(C)$.
- Symmetrically for each $Z_2 \in Res_{\max}(s'_2, o)$.

In particular, it holds that:

- For each $Z_1 \in Res_{\max}(s'_1, o)$ there exists $Z_2 \in Res_{\max}(s'_2, o)$ such that for all $\alpha \in A^*$ it holds that $CCC(z_{s'_1, o}, \alpha) \neq \emptyset$ implies $CCC(z_{s'_2, o}, \alpha) \neq \emptyset$.
- Symmetrically for each $Z_2 \in Res_{\max}(s'_2, o)$.

Given $s'_1, s'_2 \in S$ and $o \in O$ such that $((s'_1, o), (s'_2, o)) \in \mathcal{B}$, and given $Z_1 \in Res_{\max}(s'_1, o)$ and $Z_2 \in Res_{\max}(s'_2, o)$ such that $z_{s'_1, o}$ and $z_{s'_2, o}$ are related by one of the projections of \mathcal{B} , we prove that for all $\alpha \in A^*$ such that $CCC(z_{s'_1, o}, \alpha) \neq \emptyset \neq CCC(z_{s'_2, o}, \alpha)$ it holds that:

$$prob(SCC(z_{s'_1, o}, \alpha)) = prob(SCC(z_{s'_2, o}, \alpha))$$

by proceeding by induction on the length n of α :

- If $n = 0$, i.e., $\alpha = \varepsilon$, then:

$$prob(SCC(z_{s'_1, o}, \alpha)) = prob(SCC(z_{s'_2, o}, \alpha)) = \begin{cases} 1 & \text{if } o = \omega \\ 0 & \text{if } o \neq \omega \end{cases}$$

- Let $n \in \mathbb{N}_{>0}$ and suppose that the result holds for all traces of length $m = 0, \dots, n-1$ that label completed computations starting from pairs of states of Z related by one of the projections of \mathcal{B} . Assume that $\alpha = a\alpha'$. Given $s \in S$ and $Z \in Res_{\max}(s, o)$ such that $CCC(z_{s, o}, \alpha) \neq \emptyset$, it holds that, whenever $z_{s, o} \xrightarrow{a} \mathcal{D}$, then:

$$\begin{aligned} prob(SCC(z_{s, o}, \alpha)) &= \sum_{z_{s', o'} \in Z} \mathcal{D}(z_{s', o'}) \cdot prob(SCC(z_{s', o'}, \alpha')) \\ &= \sum_{[z_{s', o'}] \in Z/\mathcal{B}'} \mathcal{D}([z_{s', o'}]) \cdot prob(SCC(z_{s', o'}, \alpha')) \end{aligned}$$

where \mathcal{B}' is a projection of \mathcal{B} and the factorization of $prob(SCC(z_{s', o'}, \alpha'))$ with respect to the specific representative $z_{s', o'}$ of the equivalence class $[z_{s', o'}]$ stems from the application of the induction hypothesis on α' to all states of that equivalence class. Since $z_{s'_1, o}$ and $z_{s'_2, o}$ are related by a projection $\mathcal{B}_{1,2}$ of \mathcal{B} , it follows that, whenever $z_{s'_1, o} \xrightarrow{a} \mathcal{D}_1$, then $z_{s'_2, o} \xrightarrow{a} \mathcal{D}_2$ and:

$$\begin{aligned} prob(SCC(z_{s'_1, o}, \alpha)) &= \sum_{[z_{s', o'}] \in Z/\mathcal{B}_{1,2}} \mathcal{D}_1([z_{s', o'}]) \cdot prob(SCC(z_{s', o'}, \alpha')) = \\ &= \sum_{[z_{s', o'}] \in Z/\mathcal{B}_{1,2}} \mathcal{D}_2([z_{s', o'}]) \cdot prob(SCC(z_{s', o'}, \alpha')) = prob(SCC(z_{s'_2, o}, \alpha)) \end{aligned}$$

Therefore $s_1 \sim_{\text{PTE-tbt,dis}} s_2$. ■

All the inclusions in Thm. 6.5 are strict:

- Figure 1 shows that $\sim_{\text{PB,dis}}$ is strictly finer than \sim_{PB} . It holds that $s_1 \not\sim_{\text{PB,dis}} s_2$ because for instance the group distribution of the leftmost *offer*-transition of s_1 – which assigns probability 1 to each group containing both the state with the outgoing *draw*-transition and the state with the outgoing *win*₁-transition, probability 0.4 to each group containing the state with the outgoing *draw*-transition but not the state with the outgoing *win*₁-transition, probability 0.6 to each group containing the state with the outgoing *win*₁-transition but not the state with the outgoing *draw*-transition, and probability 0 to any other group – is not matched by the group distribution of any of the three *offer*-transitions of s_2 . In contrast, $s_1 \sim_{\text{PB}} s_2$ because, given an arbitrary group \mathcal{G} , for each *offer*-transition of s_1 (resp. s_2) there exists an *offer*-transition of s_2 (resp. s_1) such that \mathcal{G} has the same probability of being reached by both transitions. For example, the leftmost *offer*-transition of s_1 is matched by (i) the leftmost *offer*-transition of s_2 with respect to every group containing both the state with the outgoing *draw*-transition and the state with the outgoing *win*₁-transition,




Figure 12: Two NPLTS models distinguished by $\sim_{PB,dis}/\sim_{PB}/\sim_{PB,\sqcup\cap}$ and identified by testing/failure/trace equivalences

(ii) the central *offer*-transition of s_2 with respect to every group containing the state with the outgoing *draw*-transition but not the state with the outgoing *win*₁-transition, and (iii) the rightmost *offer*-transition on the side of s_2 with respect to every group containing the state with the outgoing *win*₁-transition but not the state with the outgoing *draw*-transition.

- Figure 2 shows that \sim_{PB} is strictly finer than $\sim_{PB,\sqcup\cap}$. It holds that $s_1 \not\sim_{PB} s_2$ because for instance the probability 0.5 of reaching the group containing the state with the outgoing *win*₁-transition but not the state with the outgoing *win*₂-transition after the central *offer*-transition of s_1 is not matched by the probability of reaching the same group after any of the two *offer*-transitions of s_2 . In contrast, $s_1 \sim_{PB,\sqcup\cap} s_2$ because, given an arbitrary group \mathcal{G} , the maximum probability and the minimum probability of reaching \mathcal{G} over all *offer*-transitions of the two processes are respectively the same in both processes. For example, the group containing the state with the outgoing *win*₁-transition but not the state with the outgoing *win*₂-transition has maximum probability 0.7 and minimum probability 0.3 in both processes.
- Figure 12 shows that $\sim_{PB,dis}$ is strictly finer than $\sim_{PTe-tbt,dis}$. It holds that $s_1 \not\sim_{PB,dis} s_2$ because the leftmost state with outgoing *b*-transitions reachable from s_2 is not group-distribution bisimilar to the two states with outgoing *b*-transitions reachable from s_1 . In contrast, $s_1 \sim_{PTe-tbt,dis} s_2$ because success probabilities are computed in a trace-by-trace fashion without adding up over different traces.

Moreover:

- \sim_{PB} and $\sim_{PB,\sqcup\cap}$ are incomparable with the five testing equivalences, the three failure equivalences, and the three trace equivalences. Indeed, in Fig. 13 it holds that $s_1 \sim_{PB} s_2$ (and hence $s_1 \sim_{PB,\sqcup\cap} s_2$) – as can be seen by taking the equivalence relation that pairs states having equally labeled transitions leading to the same distribution – and $s_1 \not\sim_{PT,\sqcup\cap} s_2$ (and hence s_1 and s_2 are also distinguished by the other two trace equivalences, the three failure equivalences, and the five testing equivalences) – due to trace abc having maximum probability 0.68 in the first process and 0.61 in the second process. In contrast, in Fig. 12 it holds that $s_1 \not\sim_{PB,\sqcup\cap} s_2$ (and hence $s_1 \not\sim_{PB} s_2$) – as the leftmost state with outgoing *b*-transitions reachable from s_2 is not $\sqcup\cap$ -bisimilar to the two states with outgoing *b*-transitions reachable from s_1 – and $s_1 \sim_{PTe-tbt,dis} s_2$ (and hence s_1 and s_2 are also identified by the other four testing equivalences, the three failure equivalences, and the three trace equivalences).

7. A Spectrum of Strong Behavioral Equivalences

The relationships among the different equivalences that we have considered in the previous sections are summarized in Fig. 14. In the so-called spectrum, following the terminology of [39], the absence of (chains of) arrows represents incomparability, bidirectional arrows connecting boxes indicate coincidence, and ordinary arrows stand for the strictly-more-discriminating-than relation. Continuous hexagonal boxes contain

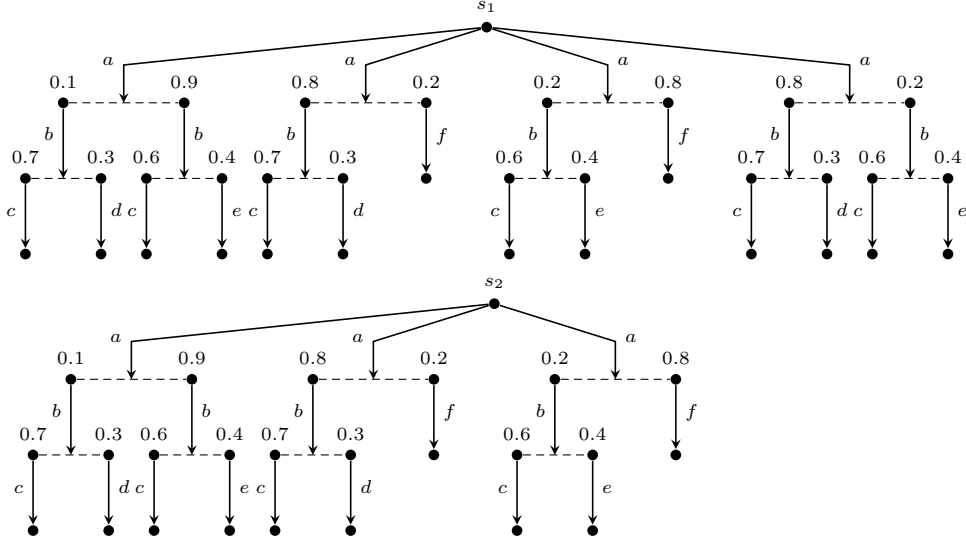


Figure 13: Two NPLTS models distinguished by testing/failure/trace equivalences and identified by $\sim_{PB}/\sim_{PB,\sqcup\cap}$

well known equivalences that compare probability distributions of all equivalence-specific events. In contrast, continuous rounded boxes contain more recent equivalences assigning a weaker role to schedulers that compare separately the probabilities of individual equivalence-specific events. Continuous rectangular boxes instead contain old and new equivalences based on extremal probabilities. The only hybrid box is the one containing $\sim_{PTe-\forall\exists}$, as this equivalence is half way between the first two definitional approaches. Dashed boxes contain equivalences that we have introduced in this paper to better assess the different impact of the approaches themselves.

The informed reader would have certainly noticed that the equivalences we considered are only a small subset of those considered in [39] (4 vs. 12). It is obviously tempting to take the challenge and study also the probabilistic variants of all the other equivalences in [39]. Indeed, while this paper was in print, we enlarged the spectrum by considering variants of trace equivalences (completed-trace equivalences), additional decorated-trace equivalences (failure-trace, readiness, and ready-trace equivalences), and variants of bisimulation equivalences (kernels of simulation, completed-simulation, failure-simulation, and ready-simulation preorders). Moreover, we studied how the spectrum changes when randomized schedulers are used instead of deterministic ones. The outcome of our studies are reported in a companion paper [7]. There definitions, theorems, proofs, and counterexamples are spelt in full detail. For the sake of completeness, in this section we summarize the new results for the additional equivalences, while in the final section we briefly discuss the impact on the spectra of randomized schedulers.

Completed-Trace Equivalences

Completed-trace equivalence is a variant of trace equivalence that additionally considers completed computations. It was introduced in the literature of fully nondeterministic models to guarantee that trace equivalence be deadlock sensitive. We have studied the following probabilistic variants of completed-trace equivalence:

- *Probabilistic completed-trace-distribution equivalence* – $\sim_{PCTr,dis}$
- *Probabilistic completed-trace equivalence* – \sim_{PCTr}
- *Probabilistic $\sqcup\cap$ -completed-trace equivalence* – $\sim_{PCTr,\sqcup\cap}$

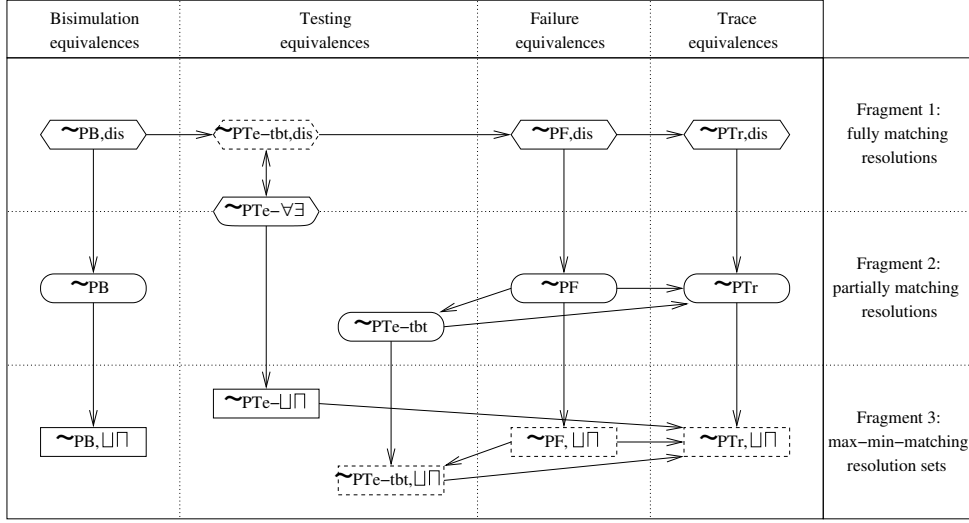


Figure 14: A spectrum of strong behavioral equivalences for NPLTS models (deterministic schedulers)

It turns out that, like in the fully nondeterministic spectrum [39], completed-trace semantics is comprised between failure semantics and trace semantics. This holds in particular for the completed-trace equivalence based on fully matching resolutions, although completed-trace semantics coincides with trace semantics in the fully probabilistic spectrum [26, 22].

Additional Decorated-Trace Equivalences

Failure semantics generalizes completed-trace equivalence towards arbitrary safety properties. An extension of failure semantics is failure-trace semantics, which takes into account failure traces. A *failure trace* is an element $\phi \in (A \times 2^A)^*$. A computation c is compatible with a given failure trace $\phi = (a_1, F_1) \dots (a_n, F_n)$ if and only if c is an $(a_1 \dots a_n)$ -compatible computation and the state reached by c after the i -th step has no outgoing transitions labeled with an action in F_i . For this semantics, we have studied the following probabilistic variants:

- *Probabilistic failure-trace-distribution equivalence* – $\sim_{PFTr,dis}$
- *Probabilistic failure-trace equivalence* – \sim_{PFTr}
- *Probabilistic $\sqcup\sqcap$ -failure-trace equivalence* – $\sim_{PFTr,\sqcup\sqcap}$

A different generalization of completed-trace semantics to capture liveness properties is readiness semantics, which considers the set of actions that can be accepted after performing a trace. We call *ready pair* an element $\varrho \in A^* \times 2^A$ formed by a trace α and a decoration R called *ready set*. A computation c is compatible with $\varrho = (\alpha, A)$ iff c is an α -compatible computation and the set of actions labeling the transitions departing from the last state in c is precisely R . For readiness semantics, we have studied the following probabilistic variants:

- *Probabilistic readiness-distribution equivalence* – $\sim_{PR,dis}$
- *Probabilistic readiness equivalence* – \sim_{PR}
- *Probabilistic $\sqcup\sqcap$ -readiness equivalence* – $\sim_{PR,\sqcup\sqcap}$

A generalization of readiness equivalences is the one based on ready traces that considers sequences of ready pairs rather than just ready pairs. A *ready trace* is an element $\rho \in (A \times 2^A)^*$. A computation c is compatible with a given ready trace $\rho = (a_i, R_1) \dots (a_n, R_n)$ if and only if c is an $(a_1 \dots a_n)$ -compatible

computation and the transitions departing from the i -th state in c is precisely R_i . For ready trace semantics, we have studied the following probabilistic variants:

- *Probabilistic ready-trace-distribution equivalence* – $\sim_{\text{PRT},\text{dis}}$
- *Probabilistic ready-trace equivalence* – \sim_{PRT}
- *Probabilistic $\sqcup\sqcap$ -ready-trace equivalence* – $\sim_{\text{PRT},\sqcup\sqcap}$

As in the fully probabilistic spectrum [26, 22], for the decorated-trace equivalences based on fully matching resolutions it holds that readiness semantics coincides with failure semantics, and this extends to ready-trace semantics and failure-trace semantics. In contrast, for the other decorated-trace equivalences based on partially matching resolutions or extremal probabilities, unlike the fully nondeterministic spectrum [39] it turns out that ready-trace equivalence and readiness equivalence are incomparable with most of the other equivalences.

Simulation-Based Equivalences

The variant of bisimulation equivalence in which only one direction is considered is called simulation preorder, which is a refinement of trace inclusion. Simulation equivalence can be defined as the kernel of the simulation preorder. In the probabilistic setting, simulation equivalence was defined by means of weight functions [24]. Instead, we have followed an alternative characterization introduced in [17], which relies on preorders as well as on closed sets, and studied the following probabilistic variants:

- *Probabilistic set-distribution similarity* – $\sim_{\text{PS},\text{dis}}$ – [34]
- *Probabilistic similarity* – \sim_{PS} – [38]
- *Probabilistic \sqcup -similarity* – $\sim_{\text{PS},\sqcup}$

Similar to trace semantics, a number of variants of simulation semantics can be defined in which the sets of actions that can be refused or accepted by states are also considered. Given $s \in S$, in the following we let $\text{init}(s) = \{a \in A \mid s \xrightarrow{a}\}$. Observing that $\text{init}(s_1) \subseteq \text{init}(s_2)$ whenever s_1 and s_2 are related by a simulation semantics, the additional constraints are the following, where the names of the obtained variants are reported in parentheses:

- $\text{init}(s_1) = \emptyset \implies \text{init}(s_2) = \emptyset$, for *completed simulation* ($\sim_{\text{PCS},\text{dis}}$, \sim_{PCS} , $\sim_{\text{PCS},\sqcup}$).
- $\text{init}(s_1) \cap F = \emptyset \implies \text{init}(s_2) \cap F = \emptyset$ for all $F \in 2^A$, for *failure simulation* ($\sim_{\text{PFS},\text{dis}}$, \sim_{PFS} , $\sim_{\text{PFS},\sqcup}$).
- $\text{init}(s_1) = \text{init}(s_2)$, for *ready simulation* ($\sim_{\text{PRS},\text{dis}}$, \sim_{PRS} , $\sim_{\text{PRS},\sqcup}$).

Of those variants, only $\sim_{\text{PFS},\text{dis}}$ has been already considered in the literature concerned with nondeterministic and probabilistic processes [15, 11].

Every simulation-based equivalence relying on partially matching transitions coincides with the corresponding simulation-based equivalence relying on extremal probabilities. Moreover, ready-simulation semantics coincides with failure-simulation semantics, but the various simulation-based semantics do not collapse to bisimulation semantics as in the case of fully probabilistic processes [24]. Each of the simulation-based equivalences relying on fully matching transitions is comprised between bisimilarity and the corresponding trace equivalence, as in the fully nondeterministic spectrum [39]. In contrast, the simulation-based equivalences relying on partially matching transitions or extremal probabilities are incomparable with most of the other equivalences.

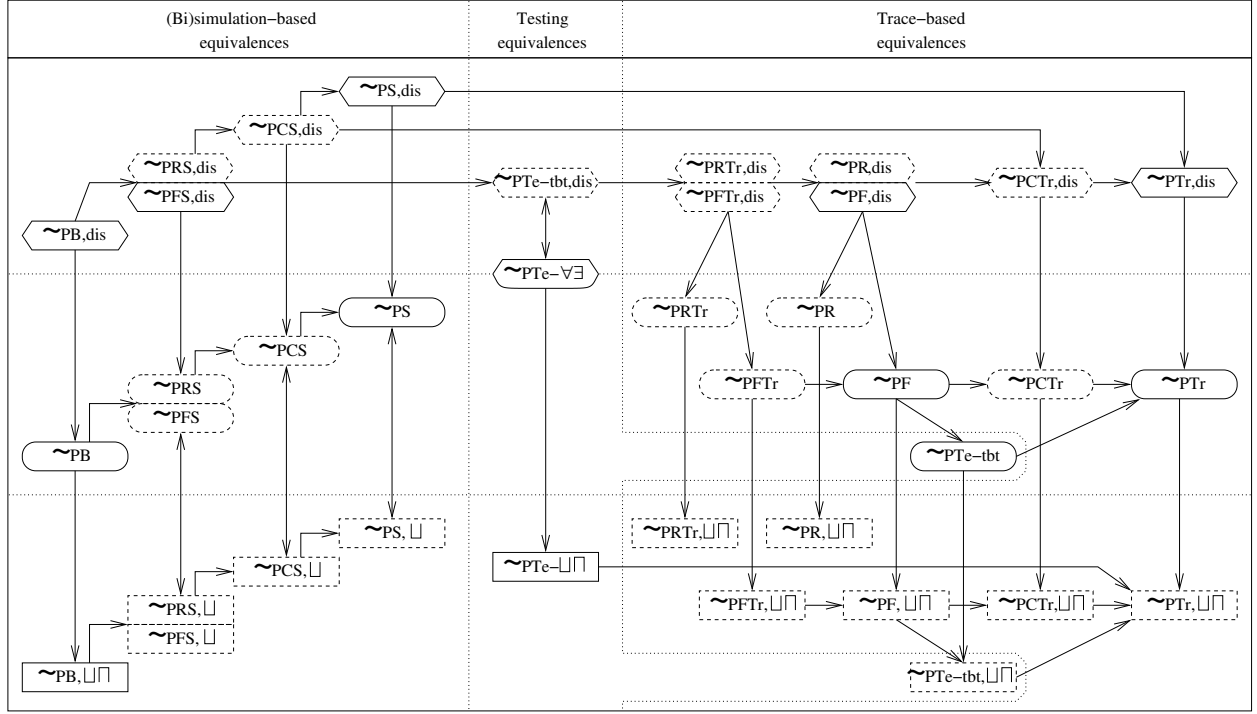


Figure 15: Full spectrum of strong behavioral equivalences for finitely-branching NPLTS models (deterministic schedulers)

A Full Spectrum

The full spectrum that considers the variants listed in this section together with those already studied in the previous ones is presented in Fig. 15. There, we follow the same graphical conventions mentioned at the beginning of this section when commenting the reduced spectrum of Fig. 14. As additional notation, we use adjacency of boxes within the same fragment with the same meaning as bidirectional arrows connecting boxes of different fragments, i.e., coincidence. It is, perhaps, worth noticing that there are many more dashed boxes than in Fig. 14; indeed, many relations have been introduced for the first time in this paper and in its companion [7].

The top fragment of the spectrum in Fig. 15 refers to equivalences that are based on fully matching resolutions. Similar to the spectrum for fully probabilistic processes in [26, 22], many equivalences collapse into a single one; in particular, ready-simulation semantics coincides with failure-simulation semantics, ready-trace semantics coincides with failure-trace semantics, and readiness semantics coincides with failure semantics. Different from the fully probabilistic spectrum, in the top fragment we have that the various simulation-based semantics do not coincide with bisimulation semantics [24] and that completed-trace semantics does not coincide with trace semantics [26, 22]. Moreover, testing semantics turns out to be finer than failure semantics.

The central and the bottom fragments of the spectrum in Fig. 15 instead refer to equivalences that are based on partially matching resolutions and extremal probabilities, respectively. These equivalences are coarser than those in the top fragment and do not flatten the specificity of the intuition behind the original definition of the behavioral equivalences for LTS models. Therefore, these two fragments preserve much of the original spectrum for fully nondeterministic processes in [39], with testing semantics being coarser than failure semantics. It is worth noting the coincidence of corresponding simulation-based equivalences in the two fragments (due to the fact that the comparison operator \leq is used in their definitions), whereas this is not the case for the two bisimulation equivalences (as the comparison operator $=$ is used instead in their definitions). We finally stress the isolation of bisimulation semantics, simulation semantics, ready-trace semantics, and readiness semantics in the two fragments, as well as the partial isolation of $\sim_{\text{PTE-}\sqcup\sqcap}$.

8. Concluding Remarks

We have studied the relationships among the strong equivalences that stem from three significantly different approaches to the definition of behavioral relations for NPLTS models. The specificity of the three approaches is determined by the way they deal with the probabilities associated with the resolutions of nondeterminism. For each approach, we have first considered trace, failure, testing, and bisimulation semantics, then following the work in [39] we have also addressed completed-trace, failure-trace, readiness, ready-trace, simulation, completed-simulation, failure-simulation, and ready-simulation semantics. The resulting spectrum in Fig. 15 obtained by assuming deterministic schedulers shows a much wider variety of options, together with certain resemblances in some of its three fragments, with respect to the fully non-deterministic spectra in [13, 39] and the fully probabilistic spectra in [26, 22]. The spectrum also presents some peculiarities, due to a number of equivalences that are incomparable with most of the others.

Randomized Schedulers. In the paper, we have considered strong equivalences that, for resolving nondeterminism, rely on deterministic schedulers. More recently, we have however examined also the impact on the equivalences of randomized schedulers and we would like to sum up here our findings.

Randomized schedulers are defined in such a way that each of them selects at each state a convex combination of equally labeled transitions, called a *combined transition* [31]. The reader is referred to [7] for their formal definition in the NPLTS setting; we would only like to remark that a deterministic scheduler is a special case of randomized one in which every selected combination involves a single ordinary transition.

For each strong behavioral equivalence, say \sim , introduced so far, we denote by \sim^{ct} the corresponding equivalence based on combined transitions (ct-equivalence for short), i.e., in which nondeterminism is resolved by means of randomized schedulers. For the eighteen trace-based equivalences and the five testing equivalences, the only modification in their definitions is the use of Res^{ct} in place of Res , where Res^{ct} is the set of resolutions of a state obtained via a randomized scheduler. For the fifteen (bi)simulation-based equivalences, the only modification in their definitions is the direct use of combined transitions (denoted by \rightarrow_{c}) instead of ordinary transitions.

All the results connecting the various equivalences and the counterexamples showing strict inclusion or incomparability are still valid for the ct-equivalences. A notable exception is given by the counterexamples based on Fig. 2, as the central *offer*-transition of s_1 can now be obtained as a convex combination of the two *offer*-transitions of s_2 with both coefficients equal to 0.5. Indeed, no ct-equivalence can be finer than the corresponding equivalence arising from deterministic schedulers, as matching ordinary transitions implies matching combined transitions.

While every ct-equivalence based on fully matching resolutions is still strictly finer than the corresponding ct-equivalences based on partially matching resolutions or extremal probabilities (the counterexample provided by Fig. 1 is still valid), it turns out that every ct-equivalence based on partially matching resolutions coincides with the corresponding ct-equivalence based on extremal probabilities. As far as $\sim_{\text{PTE-}\sqcup}$ and $\sim_{\text{PTE-}\forall}$ are concerned, their ct-variants coincide as well. In other words, when moving to randomized schedulers, the central fragment and the bottom fragment of the spectrum in Fig. 15 collapse. Pictorially, all the ordinary arrows in Fig. 15 going from the central fragment to the bottom one become bidirectional in the presence of randomized schedulers. Moreover, it holds that every ct-equivalence based on extremal probabilities coincides with the corresponding equivalence in the bottom fragment of the spectrum in Fig. 15.

Future work. We plan to investigate also the spectrum of weak behavioral equivalences, for which the choice of randomized schedulers has been shown to be more appropriate.

In addition, we would like to study the impact of the different equivalences on a simple process description language with operators for probabilistic and nondeterministic choice. Indeed, as motivated in [41] and [40], mutual distributivity of these operators has a key role in the definition of denotational models of nondeterministic and probabilistic systems. Thus, it would be interesting to see which of the many equivalences that we have considered guarantees distributivity. This should also enable us to better understand the correspondence between abstract denotational models and more concrete operational ones.

Finally, it would be interesting to compare the discriminating power of the various equivalences on a model richer than NPLTS, in which – similar to the general model of [31] – transitions are elements of

$S \times \text{Distr}(A \times S)$ instead of $S \times A \times \text{Distr}(S)$. The formalization of the equivalences on this more expressive model may be made less intricate by resorting to coalgebraic reasoning like in [36]. An even more general framework in which to perform comparisons is that of ULTRAS [5], as it has been shown to encompass many behavioral equivalences for models such as labeled transition systems, discrete-/continuous-time Markov chains, and discrete-/continuous-time Markov decision processes without/with internal nondeterminism.

Acknowledgment: We are very grateful to the anonymous referees for their helpful and constructive comments. We also would like to thank Valeria Vignudelli for pointing out an error in the previous version of the counterexample in Fig. 12. This work has been partially funded by EU Collaborative FET project no. 257414 *Autonomic Service-Component Ensemble* (ASCENS) and by MIUR PRIN project *Compositionality, Interaction, Negotiation, Autonomicity for the Future ICT Society* (CINA).

References

- [1] S. Abramsky. Observational equivalence as a testing equivalence. *Theoretical Computer Science*, 53:225–241, 1987.
- [2] C. Baier, J.-P. Katoen, H. Hermanns, and V. Wolf. Comparative branching-time semantics for Markov chains. *Information and Computation*, 200:149–214, 2005.
- [3] J.A. Bergstra, A. Ponse, and S.A. Smolka (editors). *Handbook of Process Algebra*. Elsevier, 2001.
- [4] 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. Full version: Technical Report 04/2013, IMT Institute for Advanced Studies Lucca, available on-line at <http://eprints.imtlucca.it/1523/>.
- [5] 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.
- [6] M. Bernardo, R. De Nicola, and M. Loreti. The spectrum of strong behavioral equivalences for nondeterministic and probabilistic processes. In *Proc. of the 11th Int. Workshop on Quantitative Aspects of Programming Languages and Systems (QAPL 2013)*, volume 117 of *EPTCS*, pages 81–96, 2013.
- [7] M. Bernardo, R. De Nicola, and M. Loreti. A companion of “Relating strong behavioral equivalences for processes with nondeterminism and probabilities”. *CoRR*, abs/1305.0538, 2013. Available on-line at <http://arxiv.org/abs/1305.0538>.
- [8] M. Bernardo, R. De Nicola, and M. Loreti. Group-by-group probabilistic bisimilarities and their logical characterizations. In *Proc. of the 8th Int. Symp. on Trustworthy Global Computing (TGC 2013)*, volume 8358 of *LNCS*. Springer, 2013. Full version: Technical Report 06/2013, IMT Institute for Advanced Studies Lucca, available on-line at <http://eprints.imtlucca.it/1553/>.
- [9] 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.
- [10] R. Cleaveland, Z. Dayar, S.A. Smolka, and S. Yuen. Testing preorders for probabilistic processes. *Information and Computation*, 154:93–148, 1999.
- [11] S. Crafa and F. Ranzato. A spectrum of behavioral relations over LTSs on probability distributions. In *Proc. of the 22nd Int. Conf. on Concurrency Theory (CONCUR 2011)*, volume 6901 of *LNCS*, pages 124–139. Springer, 2011.
- [12] L. de Alfaro, R. Majumdar, V. Raman, and M. Stoelinga. Game refinement relations and metrics. *Logical Methods in Computer Science*, 4(3:7):1–28, 2008.
- [13] R. De Nicola. Extensional equivalences for transition systems. *Acta Informatica*, 24:211–237, 1987.
- [14] R. De Nicola and M. Hennessy. Testing equivalences for processes. *Theoretical Computer Science*, 34:83–133, 1984.
- [15] 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.
- [16] C. Derman. *Finite State Markovian Decision Processes*. Academic Press, 1970.
- [17] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Approximating labelled Markov processes. *Information and Computation*, 184:160–200, 2003.
- [18] S. Georgievska and S. Andova. Retaining the probabilities in probabilistic testing theory. In *Proc. of the 13th Int. Conf. on Foundations of Software Science and Computation Structures (FOSSACS 2010)*, volume 6014 of *LNCS*, pages 79–93. Springer, 2010.
- [19] 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.
- [20] 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.
- [21] M. Hennessy and R. Milner. Algebraic laws for nondeterminism and concurrency. *Journal of the ACM*, 32:137–162, 1985.
- [22] D.T. Huynh and L. Tian. On some equivalence relations for probabilistic processes. *Fundamenta Informaticae*, 17:211–234, 1992.
- [23] H. Jifeng, K. Seidel, and A. McIver. Probabilistic models for the guarded command language. *Science of Computer Programming*, 28:171–192, 1997.
- [24] B. Jonsson and K.G. Larsen. Specification and refinement of probabilistic processes. In *Proc. of the 6th IEEE Symp. on Logic in Computer Science (LICS 1991)*, pages 266–277. IEEE-CS Press, 1991.

- [25] 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.
- [26] 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.
- [27] R.M. Keller. Formal verification of parallel programs. *Communications of the ACM*, 19:371–384, 1976.
- [28] K.G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94:1–28, 1991.
- [29] N. Lopez and M. Nuñez. An overview of probabilistic process algebras and their equivalences. In *Validation of Stochastic Systems*, volume 2925 of *LNCS*, pages 89–123. Springer, 2004.
- [30] 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.
- [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] R. Segala and A. Turrini. Comparative analysis of bisimulation relations on alternating and non-alternating probabilistic models. In *Proc. of the 2nd Int. Conf. on the Quantitative Evaluation of Systems (QEST 2005)*, pages 44–53. IEEE-CS Press, 2005.
- [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] L. Song, L. Zhang, and J.C. Godskesen. Bisimulations meet PCTL equivalences for probabilistic automata. In *Proc. of the 22nd Int. Conf. on Concurrency Theory (CONCUR 2011)*, volume 6901 of *LNCS*, pages 108–123. Springer, 2011.
- [38] M. Tracol, J. Desharnais, and A. Zhioua. Computing distances between probabilistic automata. In *Proc. of the 9th Int. Workshop on Quantitative Aspects of Programming Languages (QAPL 2011)*, volume 57 of *EPTCS*, pages 148–162, 2011.
- [39] R.J. van Glabbeek. The linear time – branching time spectrum I. In *Handbook of Process Algebra*, pages 3–99. Elsevier, 2001.
- [40] D. Varacca. *Probability, Nondeterminism and Concurrency. Two Denotational Models for Probabilistic Computation*. PhD Thesis, 2003.
- [41] D. Varacca and G. Winskel. Distributing probability over non-determinism. *Mathematical Structures in Computer Science*, 16:87–113, 2006.
- [42] 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.
- [43] V. Wolf. Testing theory for probabilistic systems. In *Model-Based Testing of Reactive Systems*, volume 3472 of *LNCS*, pages 233–275. Springer, 2005.
- [44] 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.