

Extending Backward Compatibility of Probabilistic Testing via Coherent Resolutions^{*}

Marco Bernardo

Dipartimento di Scienze Pure e Applicate, Università di Urbino, Italy

Abstract. Testing equivalence for processes featuring both nondeterminism and probabilities is not insensitive to the moment of occurrence of nondeterministic or probabilistic choices among identical actions. Therefore, it is only partially backward compatible with testing equivalences for fully nondeterministic processes and for fully probabilistic processes. We illustrate how its backward compatibility can be extended through the joint use of coherent resolutions of nondeterminism and additional decorations for transitions, to ensure the insensitivity to the aforementioned internal choices. We also show that full backward compatibility cannot be achieved by exhibiting a counterexample with external choices too, inspired by failure semantics for fully nondeterministic processes.

1 Introduction

Behavioral relations play a fundamental role in concurrency theory [3]. They formalize observational mechanisms that permit relating models that, despite their different representations in the same mathematical domain, cannot be distinguished by external entities when abstracting from certain internal details. Moreover, they support system modeling and verification by providing a means to relate system descriptions expressed at different levels of abstraction, as well as to reduce the size of a system representation while preserving specific properties to be assessed later.

Several approaches to the definition of behavioral relations have appeared in the literature, together with the investigation of their compositional, equational, and logical characteristics. Comparative concurrency theory is devoted to the study of the discriminating power and of the mutual relationships of behavioral relations. In the case of fully nondeterministic processes, from the first work on this subject [15] to the elaboration of the full spectrum [22], a number of equivalences have emerged that range from the branching-time – i.e., (bi)simulation-based – endpoint [33,32] to the linear-time – i.e., trace-based – endpoint [11] passing through testing relations [16].

The spectrum becomes simpler when considering fully probabilistic processes [29,24,2], whereas as shown in [7] it is much more variegated in the case of processes with nondeterminism and probabilities. The reason is that only

^{*} Copyright © 2020 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

after resolving nondeterminism via a scheduler it is possible to compute the probability of equivalence-specific events. Examples of such events are reaching via given actions certain sets of equivalent states (bisimulation semantics), executing specific action sequences (trace semantics), and passing tests formalized themselves as processes (testing semantics), with states/traces possibly being decorated with additional information.

Regardless of the specific approach, there are at least three alternative ways of applying a behavioral equivalence to nondeterministic and probabilistic processes, based on how the resolutions of nondeterminism of those processes are compared. The three alternative ways have been addressed in the spectrum of [7]:

- The first option, coming from [38,36,37], examines the probability distributions of *all* equivalence-specific events calculated over resolutions. Two processes are considered equivalent if, for each resolution of either process, 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*). The resulting portion of the spectrum closely resembles the spectrum for fully probabilistic processes.
- The second option, deriving from [41,40,6,8], compares resolutions on the basis of the probabilities of *individual* equivalence-specific events. A resolution of either process can be matched, with respect to *different* equivalence-specific events, by *different* resolutions of the other process (*partially matching resolutions*). The resulting equivalences are less discriminating than those arising from fully matching resolutions and retrieve nice logical characterizations for bisimilarity and compositionality properties for trace semantics.
- The third option, stemming from testing theories in [43,27,37] and adapted to other semantics in [7], instead of comparing individual resolutions, takes into account only the *extremal probabilities* of equivalence-specific events computed over all resolutions of the two processes (*max-min matching resolutions*). The resulting equivalences are less discriminating than the ones originated from partially matching resolutions, but both portions of the spectrum corresponding to these two families of equivalences feature many analogies with the spectrum for fully nondeterministic processes.

In this paper, we focus on testing semantics for nondeterministic and probabilistic processes, for which the third of the aforementioned options is the most applied one [43,27,37]. Each test is formalized as a finite nondeterministic and probabilistic process extended with success states, which is run in parallel with the process under test thus resulting in an interaction (or testing) system in which the process and the test have to synchronize on every action. The probability of reaching success is not unique, but depends on the specific resolution of nondeterminism considered within the interaction system. In the third option above, only the two maximal resolutions respectively yielding the maximum and the minimum success probabilities are taken into account.

It is well known that behavioral equivalences for nondeterministic and probabilistic processes tend to be overdiscriminating, thereby hampering the achievement of desirable properties. For example, in [28,17] it has been shown that the

testing equivalences of [43,27,37] can be characterized in terms of branching-time, simulation-like relations, which is consistent with the fact that those equivalences are not insensitive to the moment of occurrence of a nondeterministic or probabilistic choice that is *internal*, i.e., among identical actions. This is partly due to centralized schedulers coming into play after assembling the testing system – where a process is composed in parallel with a test, thus giving the possibility to make decisions in either component on the basis of those made in the other – which may be avoided via distributed schedulers [14,12,21]. Most importantly, it is a consequence of a special instance of the copying capability [1], which shows up in the presence of a nondeterministic choice in either component that synchronizes with a probabilistic choice in the other, thus creating copies of a state possessing several outgoing transitions where different decisions can be made.

This has a negative impact on the property, expected of any equivalence for nondeterministic and probabilistic processes, of being *backward compatible* with the corresponding equivalences for fully nondeterministic processes and for fully probabilistic processes. As recalled in [6], the testing equivalences of [43,27,37] turn out to coincide with the testing equivalence of [16] on fully nondeterministic processes only if tests are restricted to be fully nondeterministic, and with the testing equivalence of [13] on fully probabilistic processes only if tests are restricted to be fully probabilistic. The two testing equivalences of [16,13] are insensitive to internal nondeterministic or probabilistic choices, respectively.

We show that backward compatibility can be extended by making the testing equivalences of [43,27,37] insensitive to the moment in which a nondeterministic or probabilistic choice among identical actions occurs. Instead of resorting to a different definition of probabilistic testing equivalence like in [6], where backward compatibility stems from comparing success probabilities in a trace-by-trace fashion rather than cumulatively on all traces, we reuse the notion of coherent resolution of nondeterminism for probabilistic trace semantics developed in [4,5]. In the case of testing semantics, coherency must be accompanied by the introduction of additional transition decorations, so that the same decisions are made by schedulers in distinct copies of the same state of a process or a test occurring in a choice within the testing system. This is similar to the technique employed in [20] for processes in which action, nondeterministic, and probabilistic branchings alternate, with the remarkable difference that our decoration procedure is much simpler. Consistent with the ready-trace semantics characterization of [20], a counterexample inspired by failure semantics for fully nondeterministic processes shows that full backward compatibility cannot be achieved in the presence of certain synchronizations among external choices.

This paper is organized as follows. In Sect. 2, we recall background definitions for nondeterministic and probabilistic processes as well as resolutions of nondeterminism. In Sect. 3, we present testing equivalence for those processes together with its limitations about backward compatibility. In Sect. 4, we illustrate an adaptation of coherent resolutions to testing systems in which transitions are suitably decorated, so as to gain insensitivity to the moment of occurrence of

internal nondeterministic or probabilistic choices. This results in a higher level of backward compatibility with respect to [43,27,37].

2 Nondeterministic and Probabilistic Processes

Processes featuring nondeterminism and probability are typically described by extending the labeled transition system (LTS) model [30] in such a way that every action-labeled transition goes from a source state to a probability distribution over target states [31,35] rather than to a single target state. These models are essentially Markov decision processes [19], or probabilistic automata in the sense of [34], that additionally allow for internal nondeterminism, i.e., equally labeled transitions departing from the same state.

In the literature, they have been represented through a number of slightly different probabilistic computational entities such as, e.g., concurrent Markov chains [42], strictly alternating models [23], probabilistic automata in the sense of [35], and the denotational probabilistic models of [25]; see [39] for an overview. We formalize them through a variant of simple probabilistic automata [35], in which we do not distinguish between external and internal actions.

Definition 1. A nondeterministic and probabilistic labeled transition system, *NPLTS for short*, is a triple (S, A, \longrightarrow) where $S \neq \emptyset$ is an at most countable set of states, $A \neq \emptyset$ is a countable set of transition-labeling actions, and $\longrightarrow \subseteq S \times A \times \text{Distr}(S)$ is a transition relation, with $\text{Distr}(S)$ being the set of discrete probability distributions over S . ■

A transition (s, a, Δ) is written $s \xrightarrow{a} \Delta$. We say that $s' \in S$ is not reachable from s via that a -transition if $\Delta(s') = 0$, otherwise we say that it is reachable with probability $p = \Delta(s')$. The reachable states form the support of the target distribution Δ , i.e., $\text{supp}(\Delta) = \{s' \in S \mid \Delta(s') > 0\}$. An NPLTS can be depicted as a directed graph in which vertices represent states and action-labeled edges represent transitions, with states in the support of the same target distribution being linked by a dashed line and decorated with the respective probabilities when these are different from 1 (see the forthcoming Figs. 1 to 4).

The nondeterministic choice among all the transitions departing from state s can be influenced by the external environment, while the probabilistic choice of the target state for a specific outgoing transition of s takes place internally. An NPLTS represents a *fully nondeterministic* process when every transition has a target distribution with a singleton support, while it represents a *fully probabilistic* process when every state has at most one outgoing transition.

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 $s \xrightarrow{a} \Delta$.

Definition 2. Let $\mathcal{L} = (S, A, \longrightarrow)$ be an NPLTS and $s, s' \in S$. We say that the finite sequence of steps:

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

is a computation of \mathcal{L} of length $n \in \mathbb{N}$ from $s = s_0$ to $s' = s_n$ compatible with

trace $\alpha = a_1 a_2 \dots a_n \in A^*$, written $c \in \mathcal{CC}(s, \alpha)$, iff for each step $s_{i-1} \xrightarrow{a_i} s_i$ in c there is a transition $s_{i-1} \xrightarrow{a_i} \Delta_i$ in \mathcal{L} such that $s_i \in \text{supp}(\Delta_i)$, $1 \leq i \leq n$, where:

- $\Delta_i(s_i)$ is the execution probability of step $s_{i-1} \xrightarrow{a_i} s_i$ conditioned on the selection of transition $s_{i-1} \xrightarrow{a_i} \Delta_i$ at state s_{i-1} , or simply the execution probability of that step if \mathcal{L} is fully probabilistic.
- $\text{prob}(c) = \prod_{1 \leq i \leq n} \Delta_i(s_i)$ is the execution probability of c if \mathcal{L} is fully probabilistic, assuming $\text{prob}(c) = 1$ when $n = 0$.
- For $C \subseteq \mathcal{CC}(s, \alpha)$, we let $\text{prob}(C) = \sum_{c \in C} \text{prob}(c)$ if \mathcal{L} is fully probabilistic, provided that no computation in C is a proper prefix of one of the others. ■

When several transitions depart from the same state s of an NPLTS \mathcal{L} , they describe a nondeterministic choice among different behaviors. A *resolution* of s is the result of a possible way of resolving nondeterministic choices starting from s , as if a *scheduler* were applied that decides which activity has to be performed next. A resolution of nondeterminism can thus be formalized as a fully probabilistic NPLTS \mathcal{Z} with a tree-like structure, whose branching points correspond to target distributions of transitions deriving from those of \mathcal{L} .

There are two ways of resolving nondeterminism. The structure-preserving approach constructs a resolution by importing states and transitions from the original NPLTS via a *deterministic scheduler*. In a resolution of the structure-modifying approach, (i) a transition can be produced by probabilistically combining transitions of the original model via a *randomized scheduler* [35], or (ii) a state can be obtained by probabilistically splitting states of the original model via an *interpolating scheduler* [18], or (iii) a combination thereof [10].

In this paper, we focus on structure-preserving resolutions arising from centralized, memoryless, deterministic schedulers. At each step, a scheduler of this kind selects one of the transitions departing from the current state, or no transitions at all thus stopping the execution. As a consequence, the resulting resolution is isomorphic to a submodel of the original model (or of its unfolding, should cycles be present), thereby preserving the structure of the original model (or of its unfolding). If the model is fully nondeterministic, each of its resolutions coincides with a computation of the model; if the model is fully probabilistic, its maximal resolution coincides with (the unfolding of) the entire model.

Following [26,9] we introduce a *correspondence function* $\text{corr}_{\mathcal{Z}} : Z \rightarrow S$ from the acyclic state space of the resolution $\mathcal{Z} = (Z, A, \longrightarrow_{\mathcal{Z}})$ being built, to the possibly cyclic state space of the considered model $\mathcal{L} = (S, A, \longrightarrow_{\mathcal{L}})$. For each transition $z \xrightarrow{a}_{\mathcal{Z}} \Delta$, the function $\text{corr}_{\mathcal{Z}}$ must preserve the probabilities of all the states corresponding to those in $\text{supp}(\Delta)$ and must be injective over $\text{supp}(\Delta)$. In the absence of injectivity, the original structure may not be preserved in the case that the target distribution of a transition assigns the same probability to several inequivalent states. This is exemplified in Fig. 1. The correspondence function that maps z to s , z'_1 and z'_2 to s'_1 , and z''_1 and z''_2 to s''_1 would cause the rightmost NPLTS to be considered a legal resolution of the leftmost NPLTS, which is not correct as the former is not isomorphic to any submodel of the latter.




Fig. 1. Lack of injectivity breaks structure preservation

Definition 3. Let $\mathcal{L} = (S, A, \xrightarrow{\mathcal{L}})$ be an NPLTS and $s \in S$. An acyclic NPLTS $\mathcal{Z} = (Z, A, \xrightarrow{\mathcal{Z}})$ is a structure-preserving resolution of s , written $\mathcal{Z} \in Res_{sp}(s)$, iff there exists a correspondence function $corr_{\mathcal{Z}} : Z \rightarrow S$ such that $s = corr_{\mathcal{Z}}(z_s)$, for some $z_s \in Z$ acting as the initial state of \mathcal{Z} , and for all $z \in Z$ it holds that:

- If $z \xrightarrow{\mathcal{Z}} \Delta$ then $corr_{\mathcal{Z}}(z) \xrightarrow{\mathcal{L}} \Gamma$, with $corr_{\mathcal{Z}}$ being injective over $supp(\Delta)$ and satisfying $\Delta(z') = \Gamma(corr_{\mathcal{Z}}(z'))$ for all $z' \in supp(\Delta)$.
- At most one transition departs from z .

\mathcal{Z} is maximal, written $\mathcal{Z} \in Res_{sp,max}(s)$, iff for all $z \in Z$, whenever z has no outgoing transitions, then $corr_{\mathcal{Z}}(z)$ has no outgoing transitions either. ■

3 Partial Compatibility of NPLTS Testing Equivalence

The testing theories developed in [43,27,37] for nondeterministic and probabilistic processes are based on comparing the extremal probabilities of passing a test. We formalize both processes and tests as NPLTS models, with the difference that a test has finitely many states and transitions, features an acyclic graph structure, and may contain occurrences of a success state. A test is passed by a process with a certain probability if there exists a resolution of nondeterminism of the parallel composition of the process and the test, with synchronization being enforced on any action, in which the probability of reaching a state having success in its test component is equal to the given probability.

Definition 4. A nondeterministic and probabilistic test, *NPT for short*, is an acyclic NPLTS $\mathcal{T} = (O, A, \xrightarrow{\mathcal{T}})$ where both O and $\xrightarrow{\mathcal{T}}$ are finite, with O containing a distinguished success state denoted by ω having no outgoing transitions. We say that a computation of \mathcal{T} is successful iff its last state is ω . ■

Definition 5. Let $\mathcal{L} = (S, A, \xrightarrow{\mathcal{L}})$ be an NPLTS and $\mathcal{T} = (O, A, \xrightarrow{\mathcal{T}})$ be an NPT. The interaction system of \mathcal{L} and \mathcal{T} is the NPLTS $\mathcal{I}(\mathcal{L}, \mathcal{T}) = (S \times O, A, \xrightarrow{\mathcal{I}})$ where:

- Every $(s, o) \in S \times O$ is called a configuration, which is successful iff $o = \omega$.
- $(s, o) \xrightarrow{\mathcal{I}} \Delta$ iff $s \xrightarrow{\mathcal{L}} \Delta_1$ and $o \xrightarrow{\mathcal{T}} \Delta_2$ with $\Delta(s', o') = \Delta_1(s') \cdot \Delta_2(o')$ for all $(s', o') \in S \times O$.

- A computation of $\mathcal{I}(\mathcal{L}, \mathcal{T})$ is successful iff so is its last configuration. ■

We observe that $\mathcal{I}(\mathcal{L}, \mathcal{T})$ and any $\mathcal{Z} \in Res_{sp}(s, o)$ have finitely many computations due to the test structure; we denote by $\mathcal{SC}(z_{s,o})$ the set of successful computations from the initial state $z_{s,o}$ of \mathcal{Z} . Only maximal resolutions of non-determinism are taken into account within interaction systems, because the ones that are not maximal do not expose all successful computations and hence may erroneously lead to conclude that the minimal success probability is zero. We respectively denote by \sqcup and \sqcap the supremum and infimum of a set of numbers.

Definition 6. Let $\mathcal{L} = (S, A, \longrightarrow_{\mathcal{L}})$ be an NPLTS. States $s_1, s_2 \in S$ are probabilistic testing equivalent, written $s_1 \sim_{PTe-\sqcup\sqcap} s_2$, iff for every NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ it holds that:

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

As shown in Thm. 4.4 of [6], the discriminating power of $\sim_{PTe-\sqcup\sqcap}$ does not change if randomized schedulers are used in place of deterministic ones. As further shown in Thm. 4.8 of [6], $\sim_{PTe-\sqcup\sqcap}$ is backward compatible with the testing equivalence of [16] on fully nondeterministic processes (in the sense that the two equivalences coincide on those processes) only if tests are restricted to be fully nondeterministic in $\sim_{PTe-\sqcup\sqcap}$. Likewise, $\sim_{PTe-\sqcup\sqcap}$ is backward compatible with the testing equivalence of [13] on fully probabilistic processes only if tests are restricted to be fully probabilistic in $\sim_{PTe-\sqcup\sqcap}$. We recall below the definition of testing equivalence for the two considered classes of processes.

Definition 7. Let $\mathcal{L} = (S, A, \longrightarrow_{\mathcal{L}})$ be a fully nondeterministic NPLTS. States $s_1, s_2 \in S$ are fully nondeterministic testing equivalent, written $s_1 \sim_{FNDTe} s_2$, iff for every fully nondeterministic NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ it holds that:

- There exists a successful computation from (s_1, o) iff there exists a successful computation from (s_2, o) – known as may testing.
- All maximal computations from (s_1, o) are successful iff all maximal computations from (s_2, o) are successful – known as must testing. ■

Definition 8. Let $\mathcal{L} = (S, A, \longrightarrow_{\mathcal{L}})$ be a fully probabilistic NPLTS. States $s_1, s_2 \in S$ are fully probabilistic testing equivalent, written $s_1 \sim_{FPTe} s_2$, iff for every fully probabilistic NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ it holds that:

$$\text{prob}(\mathcal{SC}(s_1, o)) = \text{prob}(\mathcal{SC}(s_2, o)) \quad \blacksquare$$

The reason for the aforementioned incompleteness of backward compatibility can be illustrated through some counterexamples. For the two fully nondeterministic NPLTS models in Fig. 2(A) whose initial states are s_1 and s_2 , it is well known that $s_1 \sim_{FNDTe} s_2$, but $s_1 \not\sim_{PTe-\sqcup\sqcap} s_2$ because the fully probabilistic

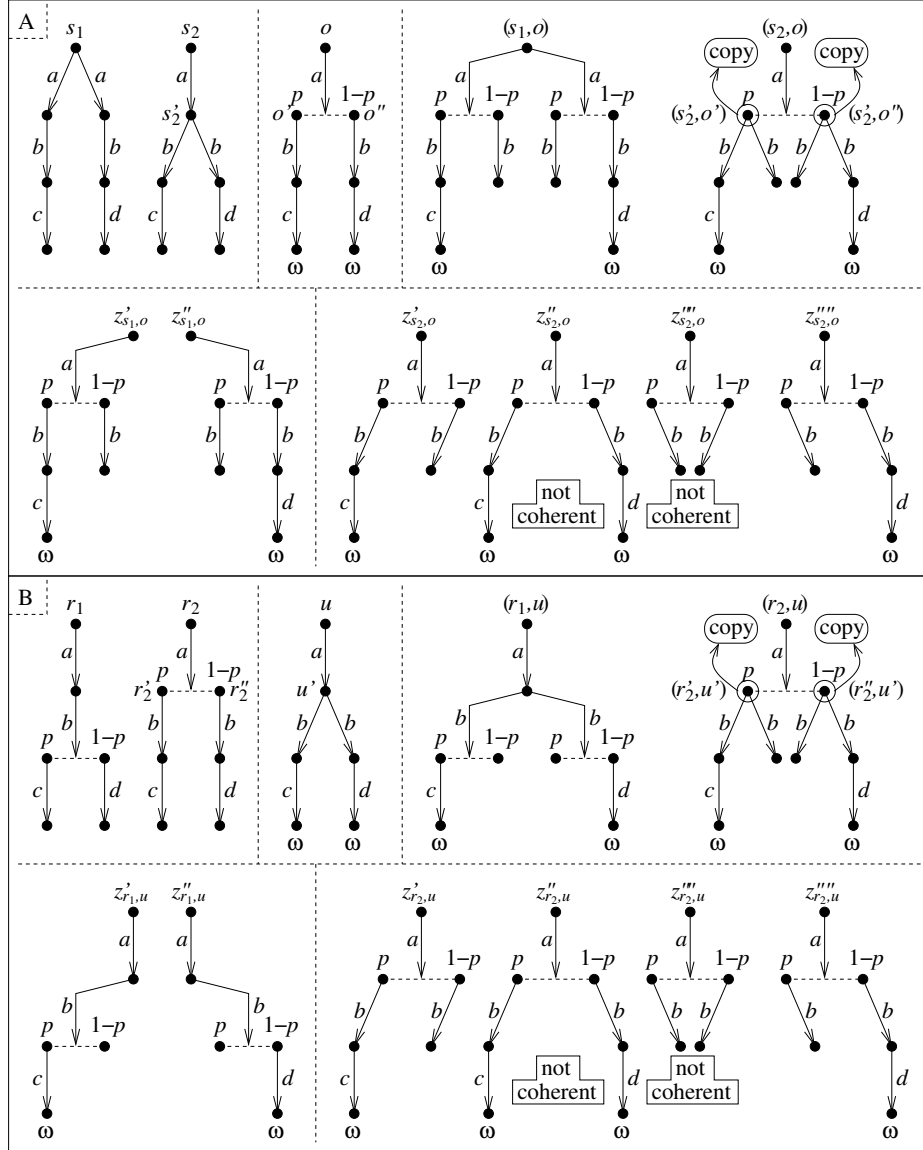


Fig. 2. (A) Two \sim_{FNDTE} -equivalent fully nondeterministic NPLTS models that are $\sim_{\text{PTE-}\sqcup}$ -distinguished by a fully probabilistic test. (B) Two \sim_{FPTe} -equivalent fully probabilistic NPLTS models $\sim_{\text{PTE-}\sqcup}$ -distinguished by a fully nondeterministic test. In both cases, an internal nondeterministic choice on b synchronizes with an internal probabilistic choice on b . This originates copies of the same state in the corresponding interaction systems, as well as sensitivity to the moment of occurrence of the internal choice in the original systems under test.

NPT with initial state o tells them apart. Assuming $p \geq 1 - p$, the interaction system with initial state (s_1, o) has two maximal resolutions yielding $\sqcup = p$ and $\sqcap = 1 - p$, while the interaction system with initial state (s_2, o) has four maximal resolutions yielding $\sqcup = 1$ and $\sqcap = 0$ instead. The synchronization of the nondeterministic choice between the two b -transitions reachable from s_2 with the probabilistic choice between the two b -transitions reachable from o creates two copies of state s'_2 in the second interaction system. The same internal nondeterministic choice is enabled in either copy, thereby giving the scheduler the opportunity of performing the incoherent selections that lead to the two maximal resolutions respectively yielding $\sqcup = 1$ and $\sqcap = 0$.

The situation is similar in Fig. 2(B) with the two \sim_{FPTe} -equivalent fully probabilistic NPLTS models whose initial states are r_1 and r_2 . They are distinguished with respect to $\sim_{\text{PTe-}\sqcup\sqcap}$ by the fully nondeterministic NPT whose initial state is u . This is due to the two copies of u' in the second interaction system, in each of which the same internal nondeterministic choice is enabled.

4 Extending Compatibility via Coherent Resolutions

The anomalies shown in Fig. 2 are due to the freedom of schedulers of making different decisions in states enabling the same actions. For these situations, in [4] we proposed to limit the excessive power of schedulers by restricting them to yield *coherent resolutions*. Intuitively, this means that, if several states in the support of the target distribution of a transition are equivalent, then the decisions made by the scheduler in those states have to be coherent with each other, so that the states to which they correspond in any resolution are equivalent too. Although developed for trace semantics, we now show that the notion of coherent resolution applies to testing semantics as well.

The restriction to coherent maximal resolutions of interaction systems is the basis for developing a variant $\sim_{\text{PTe-}\sqcup\sqcap}^c$ of probabilistic testing equivalence whose backward compatibility with \sim_{FNDTe} and \sim_{FPTe} is higher than $\sim_{\text{PTe-}\sqcup\sqcap}$. Similar to [20], in addition to coherency, within resolutions of interaction systems we need suitable decorations to differentiate among identically labeled transitions departing from states deriving from copies of a state of the process or the test.

In Fig. 2, for instance, both states (s'_2, o') and (s'_2, o'') embody a copy of s'_2 . Therefore, with respect to a scheduler, in those two states only the choice of their two left b -transitions or right b -transitions should be considered coherent, which can be achieved by decorating in the same way corresponding transitions departing from the two considered states. The situation is similar for (r'_2, u') and (r''_2, u') , with the difference that the state being copied comes from the test.

Unlike [20], our decoration procedure is very simple. The decoration of each transition of the process and of the test is just a serial number, then each transition of the interaction system inherits the serial numbers of the two transitions from which it is originated. This is illustrated in Fig. 3, where in the maximal resolution whose initial state is $z'_{s_2, o}$ the two b -transitions are coherent with each other because they both derive from the b -transition of s'_2 decorated with 1, while




Fig. 3. Application of the simplified decoration procedure based on serial numbers

this is not the case in the maximal resolution whose initial state is $z''_{s_2, o}$ because the two b -transitions respectively stem from the two b -transitions of s'_2 .

The coherency constraints behind the formalization of coherent resolutions have been introduced in [4]. They rely on coherent trace distributions, which are suitable families of sets of traces weighted with their execution probabilities in a given resolution, built through the following operations. To take decorations into account, which we assume to be unique within any NPLTS, we replace A with $B = A \times \mathbb{N} \times \mathbb{N}$ and adapt definitions accordingly.

Definition 9. For $b \in B$, $p \in \mathbb{R}$, $TD \subseteq 2^{B^* \times \mathbb{R}}$, and $T \subseteq B^* \times \mathbb{R}$ we define:

$$\begin{aligned} b \cdot TD &= \{b \cdot T \mid T \in TD\} & b \cdot T &= \{(b \beta, p') \mid (\beta, p') \in T\} \\ p \cdot TD &= \{p \cdot T \mid T \in TD\} & p \cdot T &= \{(\beta, p \cdot p') \mid (\beta, p') \in T\} \\ \text{dtr}(TD) &= \{\text{dtr}(T) \mid T \in TD\} & \text{dtr}(T) &= \{\beta \in B^* \mid \exists p' \in \mathbb{R}. (\beta, p') \in T\} \end{aligned}$$

while for $TD_1, TD_2 \subseteq 2^{B^* \times \mathbb{R}}$ we define:

$$TD_1 + TD_2 = \begin{cases} \{T_1 + T_2 \mid T_1 \in TD_1 \wedge T_2 \in TD_2 \wedge \text{dtr}(T_1) \equiv \text{dtr}(T_2)\} & \text{if } \text{dtr}(TD_1) \equiv \text{dtr}(TD_2) \\ \{T_1 + T_2 \mid T_1 \in TD_1 \wedge T_2 \in TD_2\} & \text{otherwise} \end{cases}$$

where for $T_1, T_2 \subseteq B^* \times \mathbb{R}$ we define:

$$\begin{aligned} T_1 + T_2 &= \{(\beta_1, p_1 + p_2) \mid (\beta_1, p_1) \in T_1 \wedge (\beta_2, p_2) \in T_2 \wedge \beta_1 \equiv \beta_2\} \cup \\ &\quad \{(\beta, p) \in T_1 \cup T_2 \mid \text{there is no } (\xi, q) \text{ in the other trace set s.t. } \beta \equiv \xi\} \end{aligned}$$

with:

- $\beta_1 \equiv \beta_2$ iff either $\beta_1 = \beta_2 = \varepsilon$, or $\beta_1 = \langle a, h_1, k_1 \rangle \beta'_1$, $\beta_2 = \langle a, h_2, k_2 \rangle \beta'_2$, $h_1 = h_2 \vee k_1 = k_2$, and $\beta'_1 \equiv \beta'_2$.
- $\text{dtr}(T_1) \equiv \text{dtr}(T_2)$ iff for each $\beta_1 \in \text{dtr}(T_1)$ there exists $\beta_2 \in \text{dtr}(T_2)$ such that $\beta_1 \equiv \beta_2$, and vice versa.
- $\text{dtr}(TD_1) \equiv \text{dtr}(TD_2)$ iff for each $T_1 \in TD_1$ there exists $T_2 \in TD_2$ such that $\text{dtr}(T_1) \equiv \text{dtr}(T_2)$, and vice versa. ■

Weighted trace set addition $T_1 + T_2$ is commutative and associative, with probabilities of equivalent traces in the two summands being always added up for coherency purposes. In contrast, trace distribution addition is only commutative. Essentially, the two summands in $TD_1 + TD_2$ represent two families of

sets of weighted traces executable in the resolutions of two states in the support of a target distribution. Every weighted trace set $T_1 \in TD_1$ is summed with every weighted trace set $T_2 \in TD_2$ – so as to characterize an overall resolution – unless TD_1 and TD_2 have equivalent families of trace sets, in which case summation is restricted to weighted trace sets featuring equivalent traces for the sake of coherency. Due to the lack of associativity, in the definition below all trace distributions $\Delta(s') \cdot TD_{n-1}^c(s')$ exhibiting a trace set family equivalent to Θ have to be summed up first, which is ensured by the presence of a double summation.

Definition 10. *Let (S, B, \longrightarrow) be an NPLTS and $s \in S$. The coherent decorated trace distribution of s is the subset of $2^{B^* \times \mathbb{R}_{[0,1]}}$ defined as follows:*

$$TD^c(s) = \bigcup_{n \in \mathbb{N}} TD_n^c(s)$$

with the coherent decorated trace distribution of s whose traces have length at most n , i.e., $TD_n^c(s)$, being defined as:

$$\left\{ \begin{array}{l} (\varepsilon, 1) \dagger \bigcup_{s \xrightarrow{a,h,k} \Delta} \langle a, h, k \rangle \cdot \left(\sum_{\Theta \in \text{dtr}(\Delta, n-1)} \sum_{s' \in \text{supp}(\Delta)}^{\text{dtr}(TD_{n-1}^c(s')) \equiv \Theta} \Delta(s') \cdot TD_{n-1}^c(s') \right) \\ \quad \text{if } n > 0 \text{ and } s \text{ has outgoing transitions} \\ \{ \{ (\varepsilon, 1) \} \} \\ \quad \text{otherwise} \end{array} \right.$$

where $\text{dtr}(\Delta, n-1) = \{ \text{dtr}(TD_{n-1}^c(s')) \mid s' \in \text{supp}(\Delta) \}$ and the operator $(\varepsilon, 1) \dagger$ is such that $(\varepsilon, 1) \dagger TD = \{ \{ (\varepsilon, 1) \} \} \cup T \mid T \in TD$. ■

As shown by several examples in [4], the coherency constraints should involve all $TD_n^c(-)$ distributions separately – rather than $TD^c(-)$ – and should not consider the probabilities contained in those trace distributions – which are instead necessary for alternative characterizations of probabilistic trace semantics. In [5] it was further shown that the coherency constraints should be based on a monotonic construction in which any $TD_n^c(-)$ incrementally builds on $TD_{n-1}^c(-)$, in the sense that every weighted trace set in the former should include as a subset a weighted trace set in the latter. This is achieved through a fully coherent variant of trace distribution, which we adapt below.

Definition 11. *Let (S, B, \longrightarrow) be an NPLTS and $s \in S$. The fully coherent decorated trace distribution of s is the subset of $2^{B^* \times \mathbb{R}_{[0,1]}}$ defined as follows:*

$$TD^{\text{fc}}(s) = \bigcup_{n \in \mathbb{N}} TD_n^{\text{fc}}(s)$$

with the fully coherent decorated trace distribution of s whose traces have length at most n being the subset of $TD_n^c(s)$ defined as:

$$TD_n^{\text{fc}}(s) = \left\{ \begin{array}{l} \{ T \in TD_n^c(s) \mid \exists T' \in TD_{n-1}^{\text{fc}}(s). T' \subseteq T \} \\ \quad \text{if } n > 0 \text{ and } s \text{ has outgoing transitions} \\ \{ \{ (\varepsilon, 1) \} \} \\ \quad \text{otherwise} \end{array} \right. \quad \blacksquare$$

Since testing semantics makes use of maximal resolutions only, with respect to [4,5] the first coherency constraint suffices and is adapted as follows.

Definition 12. *Let $\mathcal{L} = (S, B, \longrightarrow_{\mathcal{L}})$ be an NPLTS, $s \in S$, $\mathcal{Z} = (Z, B, \longrightarrow_{\mathcal{Z}}) \in \text{Res}_{\text{sp}, \max}(s)$ with correspondence function $\text{corr}_{\mathcal{Z}} : Z \rightarrow S$. We say that \mathcal{Z} is a*

coherent maximal resolution of s , written $\mathcal{Z} \in Res_{sp, \max}^c(s)$, iff for all $z \in \mathcal{Z}$, whenever $z \xrightarrow{a, h, k} \Delta$, then for all $n \in \mathbb{N}$ and $z', z'' \in \text{supp}(\Delta)$ it holds that:
 $dtr(TD_n^{fc}(\text{corr}_{\mathcal{Z}}(z'))) \equiv dtr(TD_n^{fc}(\text{corr}_{\mathcal{Z}}(z''))) \implies dtr(TD_n^{fc}(z')) \equiv dtr(TD_n^{fc}(z''))$ ■

Definition 13. Let $\mathcal{L} = (S, A \times \mathbb{N}, \longrightarrow_{\mathcal{L}})$ be an NPLTS and $\mathcal{T} = (O, A \times \mathbb{N}, \longrightarrow_{\mathcal{T}})$ be an NPT. The decorated interaction system of \mathcal{L} and \mathcal{T} is the NPLTS $\mathcal{I}(\mathcal{L}, \mathcal{T}) = (S \times O, B, \longrightarrow)$ where $(s, o) \xrightarrow{a, h, k} \Delta$ iff $s \xrightarrow{a, h} \Delta_1$ and $o \xrightarrow{a, k} \Delta_2$ with $\Delta(s', o') = \Delta_1(s') \cdot \Delta_2(o')$ for all $(s', o') \in S \times O$. ■

Definition 14. Let $\mathcal{L} = (S, A \times \mathbb{N}, \longrightarrow_{\mathcal{L}})$ be an NPLTS. States $s_1, s_2 \in S$ are coherent probabilistic testing equivalent, written $s_1 \sim_{\text{PTe-}\sqcup\sqcap}^c s_2$, iff for every NPT $\mathcal{T} = (O, A \times \mathbb{N}, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ it holds that:

$$\begin{aligned} \bigsqcup_{\mathcal{Z}_1 \in Res_{sp, \max}^c(s_1, o)} \text{prob}(\mathcal{SC}(z_{s_1, o})) &= \bigsqcup_{\mathcal{Z}_2 \in Res_{sp, \max}^c(s_2, o)} \text{prob}(\mathcal{SC}(z_{s_2, o})) \\ \prod_{\mathcal{Z}_1 \in Res_{sp, \max}^c(s_1, o)} \text{prob}(\mathcal{SC}(z_{s_1, o})) &= \prod_{\mathcal{Z}_2 \in Res_{sp, \max}^c(s_2, o)} \text{prob}(\mathcal{SC}(z_{s_2, o})) \end{aligned}$$
 ■

As an example, $dtr(TD_1^{fc}(s'_2, o')) = \{\{\varepsilon, \langle b, 1, 1 \rangle\}, \{\varepsilon, \langle b, 2, 1 \rangle\}\}$ is identified via \equiv with $dtr(TD_1^{fc}(s'_2, o'')) = \{\{\varepsilon, \langle b, 1, 2 \rangle\}, \{\varepsilon, \langle b, 2, 2 \rangle\}\}$, hence the states to which they correspond in any coherent maximal resolution of (s_2, o) must result in an analogous identification. In contrast, $\langle b, 1, 1 \rangle$ cannot be identified with $\langle b, 2, 2 \rangle$ because $1 = h_1 \neq h_2 = 2$ and $1 = k_1 \neq k_2 = 2$. Likewise, $\langle b, 2, 1 \rangle$ cannot be identified with $\langle b, 1, 2 \rangle$ because $2 = h_1 \neq h_2 = 1$ and $1 = k_1 \neq k_2 = 2$. As a consequence, the two maximal resolutions of (s_2, o) in Fig. 2 respectively having initial states $z''_{s_2, o}$ and $z'''_{s_2, o}$ and success probabilities 1 and 0, with the former appearing also in Fig. 3 together with its decorations, are not coherent. It thus turns out that $s_1 \sim_{\text{PTe-}\sqcup\sqcap}^c s_2$; for similar reasons, $r_1 \sim_{\text{PTe-}\sqcup\sqcap}^c r_2$.

We finally prove that the joint use of coherency and decorations makes $\sim_{\text{PTe-}\sqcup\sqcap}^c$ insensitive to the moment of occurrence of internal nondeterministic or probabilistic choices. Before that, we show that full backward compatibility with \sim_{FNDe} and \sim_{FPDe} cannot be achieved, though. Consider the two fully nondeterministic NPLTS models with initial states t_1 and t_2 in Fig. 4. These two models are known to be failure equivalent, i.e., identified by the must-part of \sim_{FNDe} [15], hence we may expect them to be identified by $\sim_{\text{PTe-}\sqcup\sqcap}^c$ too. However, this is not the case, as witnessed by the fully probabilistic NPT with initial state w because of the maximal resolution of (t_2, w) with success probability 1, in which the (external) nondeterministic choice between the b -transition and the c -transition after the central a -transition of t_2 synchronizes with the (external) probabilistic choice between the b -transition and the c -transition departing from the two states in the support of the a -transition of w .

The backward compatibility of $\sim_{\text{PTe-}\sqcup\sqcap}^c$ extends till the point in which the following property \mathcal{S}_{ext} holds: whenever an external nondeterministic choice of the process (resp. test) synchronizes with an external probabilistic choice of the test (resp. process), then all the states in the support of the target distribution of the resulting (interaction system) transition enable the same set of actions.




Fig. 4. Limit to the extension of the backward compatibility of $\sim_{P_{Te-U\cap}}^c$

Theorem 1. Let $\mathcal{L} = (S, A \times \mathbb{N}, \longrightarrow_{\mathcal{L}})$ be an NPLTS and admit only NPTs $\mathcal{T} = (O, A \times \mathbb{N}, \longrightarrow_{\mathcal{T}})$ such that $\mathcal{I}(\mathcal{L}, \mathcal{T}) = (S \times O, B, \longrightarrow)$ meets \mathcal{S}_{ext} . For $s_1, s_2 \in S$:

1. If \mathcal{L} is fully nondeterministic, then $s_1 \sim_{P_{Te-U\cap}}^c s_2 \iff s_1 \sim_{\text{FNDe}}$.
2. If \mathcal{L} is fully probabilistic, then $s_1 \sim_{P_{Te-U\cap}}^c s_2 \iff s_1 \sim_{\text{FPTe}} s_2$. ■

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. Coherent resolutions of nondeterminism. In *Proc. of the 16th European Performance Engineering Workshop (EPEW 2019)*, volume 12039 of LNCS, pages 16–32. Springer, 2019.
5. M. Bernardo. Alternative characterizations of probabilistic trace equivalences on coherent resolutions of nondeterminism. In *Proc. of the 17th Int. Conf. on the Quantitative Evaluation of Systems (QEST 2020)*, LNCS. Springer, 2020. To appear.
6. M. Bernardo, R. De Nicola, and M. Loreti. Revisiting trace and testing equivalences for nondeterministic and probabilistic processes. *Logical Methods in Computer Science*, 10(1:16):1–42, 2014.
7. M. Bernardo, R. De Nicola, and M. Loreti. Relating strong behavioral equivalences for processes with nondeterminism and probabilities. *Theoretical Computer Science*, 546:63–92, 2014.
8. M. Bernardo, R. De Nicola, and M. Loreti. Revisiting bisimilarity and its modal logic for nondeterministic and probabilistic processes. *Acta Informatica*, 52:61–106, 2015.
9. M. Bernardo, D. Sangiorgi, and V. Vignudelli. On the discriminating power of testing equivalences for reactive probabilistic systems: Results and open problems. In *Proc. of the 11th Int. Conf. on the Quantitative Evaluation of Systems (QEST 2014)*, volume 8657 of LNCS, pages 281–296. Springer, 2014.
10. F. Bonchi, A. Sokolova, and V. Vignudelli. The theory of traces for systems with nondeterminism and probability. In *Proc. of the 34th ACM/IEEE Symp. on Logic in Computer Science (LICS 2019)*, pages (19:62)1–14. IEEE-CS Press, 2019.

11. 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.
12. L. Cheung, N.A. Lynch, R. Segala, and F. Vaandrager. Switched PIOA: Parallel composition via distributed scheduling. *Theoretical Computer Science*, 365:83–108, 2006.
13. R. Cleaveland, Z. Dayar, S.A. Smolka, and S. Yuen. Testing preorders for probabilistic processes. *Information and Computation*, 154:93–148, 1999.
14. L. de Alfaro, T.A. Henzinger, and R. Jhala. Compositional methods for probabilistic systems. In *Proc. of the 12th Int. Conf. on Concurrency Theory (CONCUR 2001)*, volume 2154 of *LNCS*, pages 351–365. Springer, 2001.
15. R. De Nicola. Extensional equivalences for transition systems. *Acta Informatica*, 24:211–237, 1987.
16. R. De Nicola and M. Hennessy. Testing equivalences for processes. *Theoretical Computer Science*, 34:83–133, 1984.
17. 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.
18. Y. Deng, R.J. van Glabbeek, C. Morgan, and C. Zhang. Scalar outcomes suffice for finitary probabilistic testing. In *Proc. of the 16th European Symp. on Programming (ESOP 2007)*, volume 4421 of *LNCS*, pages 363–378. Springer, 2007.
19. C. Derman. *Finite State Markovian Decision Processes*. Academic Press, 1970.
20. S. Georgievska and S. Andova. Probabilistic may/must testing: Retaining probabilities by restricted schedulers. *Formal Aspects of Computing*, 24:727–748, 2012.
21. S. Giro and P.R. D’Argenio. On the expressive power of schedulers in distributed probabilistic systems. In *Proc. of the 7th Int. Workshop on Quantitative Aspects of Programming Languages (QAPL 2009)*, volume 253(3) of *ENTCS*, pages 45–71. Elsevier, 2009.
22. R.J. van Glabbeek. The linear time – branching time spectrum I. In *Handbook of Process Algebra*, pages 3–99. Elsevier, 2001.
23. 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.
24. D.T. Huynh and L. Tian. On some equivalence relations for probabilistic processes. *Fundamenta Informaticae*, 17:211–234, 1992.
25. H. Jifeng, K. Seidel, and A. McIver. Probabilistic models for the guarded command language. *Science of Computer Programming*, 28:171–192, 1997.
26. B. Jonsson, C. Ho-Stuart, and Wang Yi. Testing and refinement for nondeterministic and probabilistic processes. In *Proc. of the 3rd Int. Symp. on Formal Techniques in Real Time and Fault Tolerant Systems (FTRTFT 1994)*, volume 863 of *LNCS*, pages 418–430. Springer, 1994.
27. B. Jonsson and Wang 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.
28. B. Jonsson and Wang Yi. Testing preorders for probabilistic processes can be characterized by simulations. *Theoretical Computer Science*, 282:33–51, 2002.
29. 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.
30. R.M. Keller. Formal verification of parallel programs. *Communications of the ACM*, 19:371–384, 1976.

31. K.G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94:1–28, 1991.
32. R. Milner. *Communication and Concurrency*. Prentice Hall, 1989.
33. D. Park. Concurrency and automata on infinite sequences. In *Proc. of the 5th GI Conf. on Theoretical Computer Science*, volume 104 of *LNCS*, pages 167–183. Springer, 1981.
34. M.O. Rabin. Probabilistic automata. *Information and Control*, 6:230–245, 1963.
35. R. Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems*. PhD Thesis, 1995.
36. 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.
37. 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.
38. 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.
39. 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.
40. L. Song, L. Zhang, J.C. Godskesen, and F. Nielson. Bisimulations meet PCTL equivalences for probabilistic automata. *Logical Methods in Computer Science*, 9(2:7):1–34, 2013.
41. 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.
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. Wang 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.