

OBSERVING BRANCHING STRUCTURE THROUGH PROBABILISTIC CONTEXTS*

NANCY LYNCH[†], ROBERTO SEGALA[‡], AND FRITS VAANDRAGER[§]

Abstract. Probabilistic automata (PAs) constitute a general framework for modeling and analyzing discrete event systems that exhibit both nondeterministic and probabilistic behavior, such as distributed algorithms and network protocols. The behavior of PAs is commonly defined using schedulers (also called adversaries or strategies), which resolve all nondeterministic choices based on past history. From the resulting purely probabilistic structures, trace distributions can be extracted, whose intent is to capture the observable behavior of a PA. However, when PAs are composed via an (asynchronous) parallel composition operator, a global scheduler may establish strong correlations between the behavior of system components and, for example, resolve nondeterministic choices in one PA based on the outcome of probabilistic choices in the other. It is well known that, as a result of this, the (linear-time) trace distribution precongruence is not compositional for PAs. In his 1995 Ph.D. thesis, Segala has shown that the (branching-time) probabilistic simulation preorder is compositional for PAs. In this paper, we establish that the simulation preorder is, in fact, the coarsest refinement of the trace distribution preorder that is compositional. We prove our characterization result by providing (1) a context of a given PA \mathcal{A} , called the *tester*, which may announce the state of \mathcal{A} to the outside world, and (2) a specific global scheduler, called the *observer*, which ensures that the state information that is announced is actually correct. Now when another PA \mathcal{B} is composed with the tester, it may generate the same external behavior as the observer only when it is able to simulate \mathcal{A} in the sense that whenever \mathcal{A} goes to some state s , \mathcal{B} can go to a corresponding state u , from which it may generate the same external behavior. Our result shows that probabilistic contexts together with global schedulers are able to exhibit the branching structure of PAs.

Key words. probabilistic automata, trace distributions, compositionality, simulation preorder, concurrency theory

AMS subject classifications. 68Q05, 68Q10, 68Q85

DOI. 10.1147/S0097539704446487

1. Introduction. Labeled transition systems (automata) are studied extensively within concurrency theory as underlying operational models of concurrent systems [27]: a system is described as a state machine whose transitions are labeled by *actions*, where each action describes potential communication with the external environment. An important aspect of concurrency theory is the study of relationships between systems, namely equivalence and preorder relations, with the objective of understanding whether a system can be used in place of another one or as an implementation of some

*Received by the editors November 8, 2004; accepted for publication (in revised form) December 2, 2006, published electronically September 14, 2007. A preliminary version of this paper appeared as [24].

<http://www.siam.org/journals/sicomp/37-4/44648.html>

[†]Computer Science and Artificial Intelligence Laboratory, MIT, Cambridge, MA 02139 (lynch@theory.csail.mit.edu). The work of this author was supported by AFOSR contract F49620-00-1-0097, AFRL award FA9550-04-1-0121, NSF grants CCR-0121277 and CCR-0326277, NSF award CNS-0614414, NSF/ITR grants CCR-0121277 and CCR-0326277, USAF/AFRL award FA9550-04-1-0121, and DARPA/AFOSR MURI F49620-02-1-0325.

[‡]Dipartimento di Informatica, Università di Verona, Italy (roberto.segala@univr.it). This author was supported by MURST projects MEFISTO and CoVer, MIUR project AIDA, and by Inria project ProNoBis.

[§]Institute for Computing and Information Sciences, Radboud University Nijmegen, P.O. Box 9010, 6500 GL Nijmegen, The Netherlands (F.Vaandrager@cs.ru.nl). This author was supported by PROGRESS project TES4999: Verification of Hard and Softly Timed Systems (HaaST) and by DFG/NWO bilateral cooperation project Validation of Stochastic Systems (VOSS and VOSS2).

more abstract description. Several relations are studied in the literature, but the most important classes of relations are represented by simulations and bisimulations [27] and by language (trace) inclusion and equivalence [17]. An extensive classification of existing relations appears in [12], where, in particular, relations are classified as either *branching*, which observe the places where nondeterminism is resolved, or *linear*, which are insensitive to the actual places where nondeterminism is resolved. For instance, language inclusion and language equivalence are linear relations, while simulations and bisimulations are branching relations.

During the last fifteen years there has been a growing interest in the extension of concurrent models with probabilities, mainly motivated by the fact that several applications included randomized behaviors. Some of the most relevant proposals of operational models with probability and nondeterminism are reactive, generative, and stratified systems [13], concurrent labeled Markov chains [15], alternating automata [40, 29], probabilistic automata (PAs) [32], and probabilistic reactive modules [10]. Extensive comparative studies that include these models appear in [36, 4, 35].

Simulation, bisimulation, and language inclusion relations have been extended to the probabilistic case as well. In particular, [22] defines strong bisimulation on reactive systems, [34] defines strong and weak simulation and bisimulation relations on PAs, including a notion of branching bisimulation, [15] defines strong bisimulation on labeled concurrent Markov chains, [29] defines strong and weak bisimulation on alternating automata, and [2] defines branching bisimulation on alternating automata. Although the above definitions are quite different, it turns out that they can all be seen in a uniform way by viewing reactive systems, labeled concurrent Markov chains, and alternating automata as special cases of PAs [35]. For extensive comparative studies we refer the reader again to [36, 4, 35].

In this paper we are interested in extensions of language inclusion to the probabilistic case. On ordinary nondeterministic automata the resolution of nondeterminism produces sequences of alternating states and actions called *executions*; then, by restricting those sequences to visible actions, we obtain the so-called *traces*. Implementation and equivalence of nondeterministic automata can be defined in terms of inclusion and equality of sets of traces. This approach was first proposed in the context of process algebras [17] and is used extensively in the area of I/O automata [25].

An attempt to extend language inclusion to PAs appears in [31], where it is proposed that the probabilistic extension of a trace should be a probability measure over traces. Indeed, the resolution of nondeterminism on PAs produces a stochastic process that induces a probability measure over executions (a *probabilistic execution*), and the restriction of a probabilistic execution to the externally visible actions leads to a probability measure over traces (a *trace distribution*). Then, the proposal of [31] is to compare PAs based on inclusion and equality of sets of trace distributions. This is consistent with ordinary nondeterministic automata since an execution can be seen as a probabilistic execution that assigns probability 1 to a single element, and similarly a trace can be seen as a trace distribution that assigns probability 1 to a single element.

There are several arguments in favor of the point of view that probabilistic executions in PAs should play the role that executions play in ordinary nondeterministic automata, and thus in favor of the notion of trace distribution as well. One element is that this point of view leads to the definition of *weak transitions*, used to extend weak simulations and bisimulation to the probabilistic case. Another element of evidence comes from the area of distributed algorithms, where the probability of termination

of an algorithm is studied under any scheduling policy: in this context a scheduler is the entity that resolves nondeterminism, for example, by choosing the order in which processes take steps, and a probabilistic execution is the natural object where the probability of termination can be computed, as demonstrated by several case studies [23, 30, 1, 39, 20, 28, 6] and by the ongoing research on automatic verification tools for probabilistic systems [16]. Finally, again in the area of distributed algorithms, the approach of [31] to language inclusion turns out to be useful for the modular analysis of complex algorithms [30].

An important requirement for an implementation relation on systems is *compositionality*, that is, the relation is preserved by parallel composition. For labeled transition systems, the trace, simulation, and bisimulation preorders are all compositional [17, 27]. For PAs, various simulation and bisimulation preorders are known to be compositional [34]. A problem with the trace-based relations proposed in [31] is that they are not compositional; that is, they are not preserved by parallel composition. A typical solution to the problem, followed by [31], is to define a notion of *trace distribution precongrence* as the coarsest precongrence included in the trace distribution inclusion. Unfortunately, such implicit definition does not provide much insight about the structure of the relation. For this reason, there have been several attempts to characterize it in more concrete terms. In [32] trace distribution precongrence is characterized in terms of the set of trace distributions observable in a certain *principal context*—a rudimentary PA that makes very limited nondeterministic and probabilistic choices; in [33] a testing scenario is proposed. However, these indirect characterizations still do not provide much insight into the structure of trace distribution precongrence; for example, they do not explain its branching structure. Indeed, trace distribution precongrence is not a linear relation since it distinguishes ordinary nondeterministic automata that are trace equivalent.

In this paper, we provide an explicit characterization of the trace distribution precongrence, \leq_{DC} , for PAs, which completely explains its branching structure. Namely, we show that $\mathcal{P}_1 \leq_{DC} \mathcal{P}_2$ iff there exists a *weak probabilistic (forward) simulation relation* from \mathcal{P}_1 to \mathcal{P}_2 . Moreover, we provide a similar characterization of \leq_{DC} for nondeterministic automata in terms of the existence of a weak (nonprobabilistic) simulation relation. It was previously known that simulation relations are sound for \leq_{DC} [32], for both nondeterministic and probabilistic automata; we show the surprising fact that they are also *complete*. That is, we show that, for both nondeterministic and probabilistic automata, probabilistic contexts can observe all the distinctions that can be expressed using simulation relations.

Our proofs of completeness rely on special contexts for PAs, called *testers*. The tester of a PA \mathcal{P} , under the action of an appropriate scheduler, can reveal the branching structure of \mathcal{P} via a trace distribution. Such a scheduler is called an *observer scheduler*. Informally, the tester \mathcal{C} of a PA \mathcal{P} announces the outcome of each probabilistic choice of \mathcal{P} by performing an action with the name of the state reached, and flips coins to propose and announce how \mathcal{P} should resolve its nondeterministic choices. The ability of another PA \mathcal{P}' to comply with the requirements of \mathcal{C} , that is, that the trace distribution induced by the observer scheduler is also a trace distribution of $\mathcal{P}' \parallel \mathcal{C}$ (the parallel composition of \mathcal{P}' and \mathcal{C}), reveals whether \mathcal{P}' has at least the same possibilities for solving nondeterministic choices as \mathcal{P} . If $\mathcal{P} \leq_{DC} \mathcal{P}'$, then we extract a probabilistic forward simulation from \mathcal{P} to \mathcal{P}' by observing how $\mathcal{P}' \parallel \mathcal{C}$ produces the trace distribution induced by the observer scheduler.

An interesting observation about tester automata is that probabilistic choices of a PA are observed via nondeterministic choices of the tester automaton, while

nondeterministic choices of a PA are observed via probabilistic choices of the tester automaton. Thus, the branching structure of a PA is observed via a probabilistic context.

The rest of the paper is structured as follows. Sections 2 and 3 contain basic definitions and results for nondeterministic and probabilistic automata, respectively, and for the preorders we consider. These sections contain no new material, but recall definitions and theorems from the literature. For a more leisurely introduction see [25, 26, 38, 36]. Section 4 introduces the concept of tester automaton and the scheduler for a PA \mathcal{P} and its tester that reveals the structure of \mathcal{P} . Sections 5 and 6 contain, respectively, our characterization results for nondeterministic and probabilistic automata. Since the proof of the characterization result for the general case of PAs with internal actions is highly complex, we first present a proof for the special case of nondeterministic automata without internal actions (section 5.1). Then we successively show how we can also handle internal actions (section 5.2) and probabilistic choice (section 6.1) before dealing with the general case of PAs with internal actions (section 6.2). Section 7 contains our conclusions.

2. Definitions and basic results for nondeterministic automata. In this section we recall definitions and basic results for nondeterministic automata. We impose a few restrictions to avoid confusion and unnecessary complications in the rest of the paper. For more information the reader is referred to [25, 26].

2.1. Nondeterministic automata, executions, and traces. A *nondeterministic automaton* is a tuple $\mathcal{A} = (Q, \bar{q}, E, H, D)$, where

- Q is a countable set of *states*,
- $\bar{q} \in Q$ is a *start state*,
- E is a countable set of *external actions*,
- H is a countable set of *internal (hidden) actions* with $E \cap H = \emptyset$, and
- $D \subseteq Q \times (E \cup H) \times Q$ is a *transition relation*.

We denote $E \cup H$ by A , and we refer to it as the set of *actions*. We denote a transition (q, a, q') of D by $q \xrightarrow{a} q'$. We write $q \rightarrow q'$ if $q \xrightarrow{a} q'$ for some a , and we write $q \rightarrow q'$ for some q' .

We assume finite branching: for each state q the number of pairs (a, q') such that $q \xrightarrow{a} q'$ is finite. We denote the elements of a nondeterministic automaton \mathcal{A} by $Q_{\mathcal{A}}, \bar{q}_{\mathcal{A}}, E_{\mathcal{A}}, H_{\mathcal{A}}, D_{\mathcal{A}}, A_{\mathcal{A}}, \xrightarrow{a}_{\mathcal{A}}$. Often we use the name \mathcal{A} for a generic nondeterministic automaton; in this case, we usually omit the subscripts, writing simply Q, \bar{q}, E, H, D, A , and \xrightarrow{a} . We extend this convention to allow indices and primes as well; thus, the set of states of a nondeterministic automaton \mathcal{A}'_i is denoted by Q'_i .

Remark 2.1. In the definition of nondeterministic automaton above, we have imposed some restrictions that are not strictly necessary for this paper but rather avoid unnecessary complications. The restriction on the cardinality of the sets of states and actions is imposed to ensure that a nondeterministic automaton has at most countably many finite execution fragments (see definition later), which simplifies the use of measure theory later. The finite branching restriction is imposed to simplify the construction of the tester automaton in section 4; however, the results of this paper generalize to countable branching at the cost of adding complexity to the proofs (cf. Remark 4.5). We have also chosen to define nondeterministic automata with a single initial state rather than a set of initial states. Sets of initial states do not add any technical insight, but they complicate notation slightly.

An *execution fragment* of a nondeterministic automaton \mathcal{A} is a finite or infinite sequence $\alpha = q_0 a_1 q_1 a_2 q_2 \cdots$ of alternating states and actions, starting with a state and, if the sequence is finite, ending in a state, where each $(q_i, a_{i+1}, q_{i+1}) \in D$. State q_0 , the first state of α , is denoted by $fstate(\alpha)$. If α is a finite sequence, then the last state of α is denoted by $lstate(\alpha)$. An *execution* of \mathcal{A} is an execution fragment whose first state is the start state \bar{q} . We let $frags(\mathcal{A})$ denote the set of execution fragments of \mathcal{A} and $frags^*(\mathcal{A})$ the set of finite execution fragments. Similarly, we let $execs(\mathcal{A})$ denote the set of executions of \mathcal{A} and $execs^*(\mathcal{A})$ the set of finite executions.

Execution fragment α is a *prefix* of execution fragment α' , denoted by $\alpha \leq \alpha'$, if sequence α is a prefix of sequence α' . Finite execution fragment $\alpha_1 = q_0 a_1 q_1 \cdots a_k q_k$ and execution fragment α_2 can be concatenated if $fstate(\alpha_2) = q_k$. In this case the *concatenation* of α_1 and α_2 , $\alpha_1 \frown \alpha_2$, is the execution fragment $q_0 a_1 q_1 \cdots a_k \alpha_2$. Given an execution fragment α and a finite prefix α' , $\alpha \triangleright \alpha'$ (read as “ α after α' ”) is defined to be the unique execution fragment α'' such that $\alpha = \alpha' \frown \alpha''$.

The *trace* of an execution fragment α of a nondeterministic automaton \mathcal{A} , written $trace_{\mathcal{A}}(\alpha)$, or just $trace(\alpha)$ when \mathcal{A} is clear from context, is the sequence obtained by restricting α to the set of external actions of \mathcal{A} . For a set S of executions of a nondeterministic automaton \mathcal{A} , $traces_{\mathcal{A}}(S)$, or just $traces(S)$ when \mathcal{A} is clear from context, is the set of traces of the executions in S . We say that β is a trace of a nondeterministic automaton \mathcal{A} if there is an execution α of \mathcal{A} with $trace(\alpha) = \beta$. Let $traces(\mathcal{A})$ denote the set of traces of \mathcal{A} . We define the *trace preorder* relation on nondeterministic automata as follows: $\mathcal{A}_1 \leq_T \mathcal{A}_2$ iff $E_1 = E_2$ and $traces(\mathcal{A}_1) \subseteq traces(\mathcal{A}_2)$. We use \equiv_T to denote the kernel of \leq_T . That is, $\mathcal{A}_1 \equiv_T \mathcal{A}_2$ iff $\mathcal{A}_1 \leq_T \mathcal{A}_2$ and $\mathcal{A}_2 \leq_T \mathcal{A}_1$. A similar convention will be adopted to denote the kernels of other preorder relations used in the paper.

If $\beta \in A^*$, then $q \xrightarrow{\beta} q'$ iff there exists an execution fragment α such that $fstate(\alpha) = q$, $lstate(\alpha) = q'$, and $trace(\alpha) = trace(\beta)$. (Here and elsewhere, we abuse notation slightly by extending the *trace* function to arbitrary sequences.) We call $q \xrightarrow{\beta} q'$ a *weak transition*. If β is the empty sequence, then we write alternatively $q \Rightarrow q'$. Observe that the definition of $q \xrightarrow{\beta} q'$ depends only on the external actions that occur in β . We have chosen to define weak transitions for any sequence β , including internal actions as well, for notational convenience in later definitions.

We let tr range over either transitions or weak transitions. For a transition $tr = (q, a, q')$, we denote q by $source(tr)$ and q' by $target(tr)$.

2.2. Composition. We define composition of nondeterministic automata by synchronizing them on common external actions. There are several ways to do this, but the simplest approach that is followed in several papers is to synchronize nondeterministic automata on common actions and impose the restriction that no internal action of a component is an action of the other component as well. This restriction can easily be eliminated, for example, by renaming internal actions if necessary.

Nondeterministic automata \mathcal{A}_1 and \mathcal{A}_2 are *compatible* if $H_1 \cap A_2 = A_1 \cap H_2 = \emptyset$. The (*parallel*) *composition* of compatible nondeterministic automata \mathcal{A}_1 and \mathcal{A}_2 , denoted by $\mathcal{A}_1 \parallel \mathcal{A}_2$, is the nondeterministic automaton $\mathcal{A} \triangleq (Q_1 \times Q_2, (\bar{q}_1, \bar{q}_2), E_1 \cup E_2, H_1 \cup H_2, D)$, where D is the set of triples (q, a, q') such that, for $i \in \{1, 2\}$,

$$a \in A_i \Rightarrow (\pi_i(q), a, \pi_i(q')) \in D_i \quad \text{and} \quad a \notin A_i \Rightarrow \pi_i(q) = \pi_i(q'),$$

where π_i is the projection function on states of \mathcal{A} defined by $\pi_i(q_1, q_2) = q_i$.

Let α be an execution fragment of $\mathcal{A}_1 \parallel \mathcal{A}_2$, $i \in \{1, 2\}$. Then $\pi_i(\alpha)$, the i th projection of α , is the sequence obtained from α by projecting each state onto its

i th component, and removing each action not in A_i together with its following state. Sometimes we denote this projection by $\alpha[\mathcal{A}_i]$.

PROPOSITION 2.2. *Let \mathcal{A}_1 and \mathcal{A}_2 be nondeterministic automata, with $\mathcal{A}_1 \leq_T \mathcal{A}_2$. Then, for each nondeterministic automaton \mathcal{C} compatible with both \mathcal{A}_1 and \mathcal{A}_2 , $\mathcal{A}_1 \parallel \mathcal{C} \leq_T \mathcal{A}_2 \parallel \mathcal{C}$.*

2.3. Simulation relations. We define two kinds of simulation relations: forward simulations, which provide a step-by-step correspondence, and weak forward simulations, which are insensitive to the occurrence of internal steps. Namely, relation $R \subseteq Q_1 \times Q_2$ is a *forward simulation* (resp., *weak forward simulation*) from \mathcal{A}_1 to \mathcal{A}_2 iff $E_1 = E_2$ and both of the following hold:

1. $\bar{q}_1 R \bar{q}_2$.
2. If $q_1 R q_2$ and $q_1 \xrightarrow{a} q'_1$, then there exists q'_2 such that $q_2 \xrightarrow{a} q'_2$ (resp., $q_2 \xRightarrow{a} q'_2$) and $q'_1 R q'_2$.

We write $\mathcal{A}_1 \leq_F \mathcal{A}_2$ (resp., $\mathcal{A}_1 \leq_{wF} \mathcal{A}_2$) when there is a forward simulation (resp., a weak forward simulation) from \mathcal{A}_1 to \mathcal{A}_2 . It is easy to prove that both \leq_F and \leq_{wF} are preorders, that is, reflexive and transitive. Since all simulation relations in this paper are forward simulations, we often omit the word “forward.”

PROPOSITION 2.3. *Let \mathcal{A}_1 and \mathcal{A}_2 be nondeterministic automata. Then*

1. *if $\mathcal{A}_1 \leq_F \mathcal{A}_2$, then $\mathcal{A}_1 \leq_{wF} \mathcal{A}_2$;*
2. *if $H_1 = H_2 = \emptyset$, then $\mathcal{A}_1 \leq_F \mathcal{A}_2$ iff $\mathcal{A}_1 \leq_{wF} \mathcal{A}_2$;*
3. *if $\mathcal{A}_1 \leq_{wF} \mathcal{A}_2$, then $\mathcal{A}_1 \leq_T \mathcal{A}_2$.*

Proof. The proof is standard; for instance, see [26]. \square

2.4. Tree-structured automata. We say that a nondeterministic automaton is *tree-structured* if each state is reached via (i.e., occurs as a final state of) a unique execution.

The *unfolding* of nondeterministic automaton \mathcal{A} , denoted by $Unfold(\mathcal{A})$, is the tree-structured nondeterministic automaton \mathcal{B} obtained from \mathcal{A} by unfolding its transition graph into a tree. Formally,

- $Q_{\mathcal{B}} = \text{execs}^*(\mathcal{A})$,
- $\bar{q}_{\mathcal{B}} = \bar{q}_{\mathcal{A}}$,
- $E_{\mathcal{B}} = E_{\mathcal{A}}$,
- $H_{\mathcal{B}} = H_{\mathcal{A}}$, and
- $D_{\mathcal{B}} = \{(\alpha, a, \alpha a q) \mid (\text{lstate}(\alpha), a, q) \in D_{\mathcal{A}}\}$.

PROPOSITION 2.4. $\mathcal{A} \equiv_F Unfold(\mathcal{A})$.

Proof. See [26]. It is easy to check that the relation R , where $\alpha R q$ iff $\text{lstate}(\alpha) = q$, is a forward simulation from $Unfold(\mathcal{A})$ to \mathcal{A} and that the inverse relation of R is a forward simulation from \mathcal{A} to $Unfold(\mathcal{A})$. \square

PROPOSITION 2.5. $\mathcal{A} \equiv_T Unfold(\mathcal{A})$.

Proof. The proof follows by Proposition 2.4 and Proposition 2.3, parts 1 and 3. \square

3. Definitions and basic results for probabilistic automata.

3.1. Preliminaries and notation on measure theory. We recall a few basic definitions and notation for measure theory that can be retrieved from any standard book on the subject (e.g., [11]).

A σ -field over a set X is a set $\mathcal{F} \subseteq 2^X$ that contains the empty set and is closed under complement and countable union. A pair (X, \mathcal{F}) , where \mathcal{F} is a σ -field over X , is called a *measurable space*. A *measure* on a measurable space (X, \mathcal{F}) is a function $\mu : \mathcal{F} \rightarrow [0, \infty]$ such that $\mu(\emptyset) = 0$ and μ is countably additive: for each

countable family $\{C_i\}_i$ of pairwise disjoint elements of \mathcal{F} , $\mu(\cup_i C_i) = \sum_i \mu(C_i)$. A *probability measure* on (X, \mathcal{F}) is a measure μ on (X, \mathcal{F}) such that $\mu(X) = 1$. A *subprobability measure* on (X, \mathcal{F}) is a measure μ on (X, \mathcal{F}) such that $\mu(X) \leq 1$. A *discrete probability measure* on a set X is a probability measure μ on $(X, 2^X)$. A *discrete subprobability measure* on X is a subprobability measure μ on $(X, 2^X)$. We denote the set of discrete probability measures and discrete subprobability measures on X by $Disc(X)$ and $SubDisc(X)$, respectively. We denote the *support* of a discrete measure μ , that is, the set of elements of X that have nonzero measure, by $supp(\mu)$. We let $\delta(q)$ denote the *Dirac measure* for q , the discrete probability measure that assigns probability 1 to $\{q\}$. Finally, if X is nonempty and finite, then $\mathcal{U}(X)$ denotes the *uniform distribution* over X , the discrete measure that assigns probability $1/|X|$ to each element of X . Given two discrete probability measures μ_1, μ_2 on $(X, 2^X)$ and $(Y, 2^Y)$, respectively, we denote by $\mu_1 \times \mu_2$ the *product measure*, that is, the measure on $(X \times Y, 2^{(X \times Y)})$ such that $\mu_1 \times \mu_2((x, y)) = \mu_1(x)\mu_2(y)$ for each $x \in X, y \in Y$.

Sometimes it is useful to know the probability μ of some event C , knowing that some other event C' takes place. We call this the measure of C *conditional* on C' and denote it by $\mu(C | C')$. Such probability is defined to be 0 if $\mu(C') = 0$, and $\mu(C \cap C')/\mu(C')$ otherwise.

A function $f : X \rightarrow Y$ is said to be *measurable* from (X, \mathcal{F}_X) to (Y, \mathcal{F}_Y) if the inverse image of each element of \mathcal{F}_Y is an element of \mathcal{F}_X , that is, for each $C \in \mathcal{F}_Y$, $f^{-1}(C) \in \mathcal{F}_X$. In such a case, given a measure μ on (X, \mathcal{F}_X) , the function $f(\mu)$ defined on \mathcal{F}_Y by $f(\mu)(C) = \mu(f^{-1}(C))$ for each $C \in \mathcal{F}_Y$ is a measure on (Y, \mathcal{F}_Y) and is called the *image measure* of μ under f .

Given a countable collection of measures $\{\mu_i\}_i$ on (X, \mathcal{F}_X) and a countable collection $\{p_i\}_i$ of real numbers in $[0, \infty)$, denote by $\sum_i p_i \mu_i$ a new function μ such that, for each element $C \in \mathcal{F}_X$, $\mu(C) = \sum_i p_i \mu_i(C)$. We state a few elementary properties.

PROPOSITION 3.1. *The following hold:*

1. $\sum_i p_i \mu_i$ is a measure on (X, \mathcal{F}_X) .
2. If each μ_i is a (sub)probability measure and $\sum_i p_i = 1$, then $\sum_i p_i \mu_i$ is a (sub)probability measure.
3. If f is a measurable function from (X, \mathcal{F}_X) to (Y, \mathcal{F}_Y) , then $f(\sum_i p_i \mu_i) = \sum_i p_i f(\mu_i)$.

3.2. PAs, executions, and traces. A *probabilistic automaton (PA)* is a tuple $\mathcal{P} = (Q, \bar{q}, E, H, D)$, where all components are exactly as for nondeterministic automata, except that the following holds:

- D , the *transition relation*, is a subset of $Q \times (E \cup H) \times Disc(Q)$.

We define A as before. Also, we use the name \mathcal{P} for a generic PA, and we refer to its components by writing simply Q, \bar{q}, E, H, D, A , and \xrightarrow{a} . We extend this convention to allow indices and primes as well; thus, the set of states of a PA \mathcal{P}'_i is denoted by Q'_i . We denote a transition (q, a, μ) by $q \xrightarrow{a} \mu$. We assume finite branching: for each state q the number of pairs (a, μ) such that $q \xrightarrow{a} \mu$ is finite. Given a transition $tr = (q, a, \mu)$, we denote q by *source*(tr) and μ either by *target*(tr) or by μ_{tr} .

Thus, a PA differs from a nondeterministic automaton in that a transition leads to a probability measure over states rather than to a single state. A nondeterministic automaton can be viewed as a special case of a PA, where the last component of each transition is a Dirac measure. Conversely, we can associate a nondeterministic automaton with each PA by replacing transition relation D by the relation D' given

by

$$(q, a, q') \in D' \Leftrightarrow (\exists \mu)[(q, a, \mu) \in D \wedge \mu(q') > 0].$$

Using this correspondence, notions such as execution fragments and traces carry over from nondeterministic automata to PAs.¹ For instance, an execution fragment of a PA is simply an execution fragment of its associated nondeterministic automaton. Along the same lines we write $q \xrightarrow{a} q'$ whenever there exists a measure μ such that $q \xrightarrow{a} \mu$ and $q' \in \text{supp}(\mu)$.

An execution fragment of a PA is the result of resolving nondeterministic as well as probabilistic choices; however we are interested also in the outcome of the resolution of nondeterministic choices only. We can think of resolving nondeterminism by unfolding the transition relation of a PA and then choosing only one transition at each point. From the formal point of view it is more convenient to define a function, called a *scheduler*, that chooses transitions based on the past history (i.e., the current position in the unfolding of the transition relation).

A *scheduler* for a PA \mathcal{P} is a function $\sigma : \text{frags}^*(\mathcal{P}) \rightarrow \text{SubDisc}(D)$ such that $tr \in \text{supp}(\sigma(\alpha))$ implies $\text{source}(tr) = \text{lstate}(\alpha)$. A scheduler σ is said to be *deterministic* if for each finite execution fragment α either $\sigma(\alpha)(D) = 0$ or else $\sigma(\alpha) = \delta(tr)$ (the Dirac measure for tr) for some $tr \in D$. A scheduler σ is *memoryless* if it depends only on the last state of its argument, that is, for each pair α_1, α_2 of finite execution fragments, if $\text{lstate}(\alpha_1) = \text{lstate}(\alpha_2)$, then $\sigma(\alpha_1) = \sigma(\alpha_2)$.

Informally, $\sigma(\alpha)$ describes the rule for choosing a transition after α has occurred. The rule itself may be randomized. Since $\sigma(\alpha)$ is a subprobability measure, it is possible that with some nonzero probability no transition is chosen, which corresponds to terminating the computation (with what in nondeterministic automata is called a finite execution fragment). Deterministic schedulers are not allowed to use randomization in their choices, while memoryless schedulers are not allowed to look at the past history in their choices. Deterministic and memoryless schedulers are easier to analyze compared to general schedulers, and several properties (e.g., reachability) can be studied by referring to deterministic memoryless schedulers only. Note that a deterministic memoryless scheduler can be represented alternatively as a partial function from Q to D .

A scheduler σ and a discrete probability measure over states μ induce a measure ϵ on the σ -field generated by cones of execution fragments as follows. If α is a finite execution fragment, then the *cone* of α is defined by $C_\alpha = \{\alpha' \in \text{frags}(\mathcal{P}) \mid \alpha \leq \alpha'\}$. The measure ϵ of a cone C_α is defined to be $\mu(q)$ if $\alpha = q$ for some state $q \in Q$, and if α is of the form $\alpha'a'q'$, it is defined by the recursive equation

$$(1) \quad \epsilon(C_\alpha) = \epsilon(C_{\alpha'}) \sum_{tr \in D(a')} \sigma(\alpha')(tr) \mu_{tr}(q'),$$

where $D(a')$ denotes the set of transitions of D that are labeled by a' . Roughly speaking, the measure of a cone C_α equals the probability of doing α when using σ to resolve nondeterminism. Standard measure theoretical arguments ensure that ϵ is well defined. We call the measure ϵ a *probabilistic execution fragment* of \mathcal{P} , and we say that ϵ is *generated* by σ and μ . We also denote by $\epsilon_{\sigma, \mu}$ the probabilistic execution fragment generated by σ and μ .

¹The correspondence between nondeterministic automata and PAs is worked out in great detail in [4].

PROPOSITION 3.2. *Let σ be a scheduler and μ be a discrete probability measure over states. Then $fstate(\epsilon_{\sigma,\mu}) = \mu$.*

Proof. The proof follows immediately by definition of $\epsilon_{\sigma,\mu}$ after observing that the inverse image under $fstate$ of a state q is the set C_q . \square

We call the measure $fstate(\epsilon)$ the *first state* of ϵ . If $fstate(\epsilon)$ is the Dirac measure over the start state \bar{q} , then ϵ is called a *probabilistic execution*. We often write $\epsilon_{\sigma,q}$ for $\epsilon_{\sigma,\delta(q)}$, and we say that $\epsilon_{\sigma,q}$ is generated by σ and q .

Example 3.1 (the cone construction is rich). The cone construction produces a very rich set of measurable events. The event “action a occurs at least once,” that is, the set of execution fragments where an action a occurs at least once, is measurable since it can be expressed as a union of cones and there are at most countably many cones in a PA. Similarly the event “action a occurs at least n times” is measurable for any natural number n . The event “action a occurs exactly n times” is measurable since it is the intersection of “action a occurs at least n times” with the complement of “action a occurs at least $n + 1$ times.” Also the event “action a occurs finitely (infinitely) many times” is measurable since it is the countable union (intersection) of “action a occurs exactly (at least) n times.” Similar arguments hold for occurrences of states rather than actions.

Any singleton set is measurable since for an infinite execution fragment α the set $\{\alpha\}$ is the intersection of the cones of all its finite prefixes, while for a finite execution fragment α the set $\{\alpha\}$ is the intersection of C_α with the complement of the union of the cones of the extensions of α . Thus, also the set of finite execution fragments is measurable, and the set of infinite execution fragments is measurable as well.

Observe that the probability of a finite execution fragment α is the probability that α occurs and then the computation terminates. Thus, the probability of the set of finite execution fragments represents the probability of termination in a probabilistic execution fragment. This leads to the idea that a probabilistic execution fragment should be called *finite* if the probability of the set of finite execution fragments is 1.

We now show how to obtain a probability measure over traces from a probabilistic execution fragment. The measurable space is the pair $(E^* \cup E^\omega, \mathcal{F})$, where \mathcal{F} is the σ -field generated by cones of traces. More precisely, the cone of a finite trace β is defined by $C_\beta = \{\beta' \in E^* \cup E^\omega \mid \beta \leq \beta'\}$, where \leq denotes the prefix ordering on sequences. It is easy to check that the *trace* function is measurable since the inverse image of a cone C_β is a union of cones, specifically those cones C_α such that $\beta \leq trace(\alpha)$, and since there are countably many finite execution fragments in a PA.

Given a probabilistic execution fragment ϵ , we define the *trace distribution* of ϵ , $tdist(\epsilon)$, to be the image measure of ϵ under $trace$. We denote the set of trace distributions of probabilistic executions of a PA \mathcal{P} by $tdists(\mathcal{P})$. We define the *trace distribution preorder* relation on PAs by $\mathcal{P}_1 \leq_D \mathcal{P}_2$ iff $E_1 = E_2$ and $tdists(\mathcal{P}_1) \subseteq tdists(\mathcal{P}_2)$.

An example of a measurable set of traces that is used extensively throughout the paper is the set $E^*a(E^* \cup E^\omega)$ of traces in which a specific action a occurs. We denote this set by $\diamond a$. The inverse image under $trace$ of $\diamond a$ can be expressed as a disjoint union of cones of executions, namely the cones of the minimal executions with trace in $\diamond a$. Thus, we have the following proposition, whose elementary proof is omitted.

PROPOSITION 3.3. *Let η be the trace distribution of a probabilistic execution ϵ of a PA \mathcal{P} , and let Θ_a be the set of finite executions of \mathcal{P} with a single occurrence of*

action a whose last transition is labeled by a . Then

$$(2) \quad \eta(\diamond a) = \sum_{\alpha \in \Theta_a} \epsilon(C_\alpha).$$

3.3. Combined, weak, and hyper transitions. We define three new kinds of transitions that play crucial roles in the paper. Informally, a combined transition is a convex combination of transitions that are labeled by the same action, a weak combined transition abstracts from internal computation and is obtained by performing several, possibly zero, combined transitions, while a hyper-transition is a generalization of combined transitions and weak combined transitions where the starting point is a measure over states rather than a single state.

3.3.1. Combined transitions. Let $\{q \xrightarrow{a} \mu_i\}_{i \in I}$ be a collection of transitions of a PA \mathcal{P} , and let $\{p_i\}_{i \in I}$ be a collection of probabilities such that $\sum_{i \in I} p_i = 1$. Then the triple $(q, a, \sum_{i \in I} p_i \mu_i)$ is called a *combined transition* of \mathcal{P} .

3.3.2. Weak transitions. Consider a probabilistic execution fragment ϵ of a PA \mathcal{P} , with first state $\delta(q)$, that assigns probability 1 to the set of all finite execution fragments with trace $trace(\beta)$ for some $\beta \in A^*$. Let μ be the discrete measure on Q defined by $\mu(q') = \epsilon(\{\alpha \mid lstate(\alpha) = q'\})$. Then $q \xrightarrow{\beta} \mu$ is a *weak combined transition* of \mathcal{P} . We refer to ϵ as a *representation* of $q \xrightarrow{\beta} \mu$. Observe that the measure μ can be seen alternatively as the image measure of ϵ under $lstate$. This is an abuse of notation because $lstate$ is not defined for infinite execution fragments; however, since ϵ assigns measure 1 to the set of finite execution fragments, we can extend the definition of $lstate$ to infinite execution fragments for this purpose: for instance, we define the last state of any infinite execution fragment to be \bar{q} .

The notion of weak combined transition that we have just defined for PAs is a conservative extension of the corresponding notion defined for nondeterministic automata. Indeed, it is routine to check that whenever $q \xrightarrow{\beta} q'$ is a weak transition of a nondeterministic automaton \mathcal{A} , then $q \xrightarrow{\beta} \delta(q')$ is a weak combined transition of \mathcal{A} viewed as a PA.

PROPOSITION 3.4. *Let $\{tr_i = (q, a, \mu_i)\}_{i \in I}$ be a collection of weak combined transitions of a PA \mathcal{P} , and let $\{p_i\}_{i \in I}$ be probabilities such that $\sum_{i \in I} p_i = 1$. Then $(q, a, \sum_{i \in I} p_i \mu_i)$, written $\sum_{i \in I} p_i tr_i$, is a weak combined transition of \mathcal{P} .*

Proof. For each $i \in I$, let ϵ_i be a representation of tr_i , and σ_i be a scheduler that, together with state q , induces ϵ_i . We omit the index set I in the rest of the proof. For each finite execution fragment α , let $N(\alpha) \triangleq \sum_i p_i \epsilon_i(C_\alpha)$. Define a new scheduler σ as follows:

$$\sigma(\alpha) = \begin{cases} \sum_i \frac{p_i \epsilon_i(C_\alpha)}{N(\alpha)} \sigma_i(\alpha) & \text{if } N(\alpha) > 0, \\ \text{arbitrary} & \text{otherwise.} \end{cases}$$

Informally, the weight that $\sigma(\alpha)$ gives to the choice $\sigma_i(\alpha)$ is the normalized probability with which σ_i contributes to the generation of α . Let ϵ be the probabilistic execution fragment induced by σ and q . Let α be a finite execution fragment of \mathcal{P} . We first prove by induction on the length of α that $\epsilon(C_\alpha) = N(\alpha)$. The base case is trivial since $\epsilon(C_q) = 1$ and for each i , $\epsilon_i(C_q) = 1$, which implies $N(q) = \sum_i p_i \epsilon_i(C_q) = 1$; similarly, for each state $q' \neq q$, $\epsilon(C_{q'}) = 0$ and for each i , $\epsilon_i(C_{q'}) = 0$. For the inductive step, let $\alpha = \alpha' a' q'$. If $\epsilon(C_{\alpha'}) = 0$, then, by induction, $N(\alpha') = \sum_i p_i \epsilon_i(C_{\alpha'}) = 0$, which implies that for each i , $p_i \epsilon_i(C_{\alpha'}) = 0$. By the definition of the measure of a cone,

in (1), $\epsilon(C_\alpha) = 0$. Furthermore, for each i , if $p_i = 0$, then $p_i \epsilon_i(C_\alpha) = 0$ trivially, and if $p_i > 0$, then $\epsilon_i(C_{\alpha'}) = 0$, and by definition of measure of a cone, in (1), $\epsilon_i(C_\alpha) = 0$, which implies $p_i \epsilon_i(C_\alpha) = 0$. Thus, $N(\alpha) = 0$ as needed. If $\epsilon(C_{\alpha'}) > 0$, then, by the definition of the measure of a cone, in (1),

$$\epsilon(C_\alpha) = \epsilon(C_{\alpha'}) \sum_{tr \in D(\alpha')} \sigma(\alpha')(tr) \mu_{tr}(q').$$

By the induction hypothesis, $N(\alpha') > 0$. Thus, by expanding $\sigma(\alpha')(tr)$ with the definition of σ , we obtain

$$\epsilon(C_\alpha) = \epsilon(C_{\alpha'}) \sum_{tr \in D(\alpha')} \left(\sum_i \frac{p_i \epsilon_i(C_{\alpha'})}{N(\alpha')} \sigma_i(\alpha')(tr) \right) \mu_{tr}(q').$$

By standard algebraic manipulations (exchanges of sums and rearrangements of constants) we obtain

$$\epsilon(C_\alpha) = \frac{\epsilon(C_{\alpha'})}{N(\alpha')} \sum_i \sum_{tr \in D(\alpha')} p_i \epsilon_i(C_{\alpha'}) \sigma_i(\alpha')(tr) \mu_{tr}(q').$$

By induction, $\epsilon(C_{\alpha'}) = N(\alpha')$. Thus, by simplifying (removing) the leftmost term and rearranging constants, we obtain

$$\epsilon(C_\alpha) = \sum_i p_i \left(\epsilon_i(C_{\alpha'}) \sum_{tr \in D(\alpha')} \sigma_i(\alpha')(tr) \mu_{tr}(q') \right).$$

Finally, by the definition of the measure of a cone, in (1), we get the desired equation

$$\epsilon(C_\alpha) = N(\alpha) = \sum_i p_i \epsilon_i(C_\alpha).$$

Thus, $\epsilon = \sum_i p_i \epsilon_i$. Since each ϵ_i assigns probability 1 to the set of finite execution fragments of \mathcal{P} with trace $trace(a)$, then so does ϵ . Furthermore, by Proposition 3.1.3, $lstate(\epsilon) = \sum_i p_i lstate(\epsilon_i)$. That is, ϵ is a representation of a weak combined transition $(q, a, \sum_i p_i lstate(\epsilon_i))$, which, since $lstate(\epsilon_i) = \mu_i$, is the triplet $(q, a, \sum_i p_i \mu_i)$. Hence, $\sum_i p_i tr_i$ is a weak combined transition of \mathcal{P} . \square

3.3.3. Hyper-transitions. Let \mathcal{P} be a PA with $a \in A$, and let $\mu \in Disc(Q)$. For each $q \in supp(\mu)$, suppose that $q \xrightarrow{\alpha} \mu_q$ is a combined transition of \mathcal{P} . Let μ' be $\sum_{q \in supp(\mu)} \mu(q) \mu_q$. Then $\mu \xrightarrow{\alpha} \mu'$ is called a *hyper-transition* of \mathcal{P} . Also, let $\beta \in A^*$, and for each $q \in supp(\mu)$, suppose that $q \xrightarrow{\beta} \mu_q$ is a weak combined transition of \mathcal{P} . Let μ' be $\sum_{q \in supp(\mu)} \mu(q) \mu_q$. Then $\mu \xrightarrow{\beta} \mu'$ is called a *weak hyper-transition* of \mathcal{P} .

We now prove two technical properties of weak hyper-transitions. The first property gives an alternative definition of weak hyper-transition and is used to prove the second property; the second property states that weak hyper-transitions can be concatenated. It will be used in section 6.2.

PROPOSITION 3.5. *There is a weak hyper-transition $\mu \xrightarrow{\beta} \mu'$ iff there is a scheduler σ such that $\epsilon_{\sigma, \mu}$ assigns probability 1 to the set of finite execution fragments with trace β , and $lstate(\epsilon_{\sigma, \mu}) = \mu'$. We say that $\epsilon_{\sigma, \mu}$ represents $\mu \xrightarrow{\beta} \mu'$.*

Proof. Let $\{q_i\}_I$ be an enumeration of the states in $supp(\mu)$. We prove the two implications separately.

⇒ For each i let σ_i be a scheduler such that ϵ_{σ_i, q_i} represents $q_i \xrightarrow{\beta} \mu_i$ and $\mu' = \sum_I \mu(q_i)\mu_i$. Let σ be a new scheduler defined as follows:

$$\sigma(\alpha) = \begin{cases} \sigma_i(\alpha) & \text{if } fstate(\alpha) = q_i \text{ for some } i \in I, \\ 0 & \text{otherwise.} \end{cases}$$

We first prove that $\epsilon_{\sigma, \mu} = \sum_i \mu(q_i)\epsilon_{\sigma_i, q_i}$ by demonstrating that $\epsilon_{\sigma, \mu}(C_\alpha) = \sum_i \mu(q_i)\epsilon_{\sigma_i, q_i}(C_\alpha)$ for each finite execution fragment α . The proof is by induction on the length of α . For the base case, let $\alpha = q$ for some state q . By the definition of the measure of a cone, $\epsilon_{\sigma, \mu}(C_q) = \mu(q)$, and, for each i , $\epsilon_{\sigma_i, q_i}(C_q)$ is 1 if $q = q_i$ and 0 otherwise. Thus, $\epsilon_{\sigma, \mu}(C_q) = \sum_i \mu(q_i)\epsilon_{\sigma_i, q_i}(C_q)$ trivially. For the inductive step, let α be $\alpha'aq$. If $fstate(\alpha') \notin supp(\mu)$, then trivially $\epsilon_{\sigma, \mu}(C_\alpha) = 0$ and, for each i , $\epsilon_{\sigma_i, q_i}(C_\alpha) = 0$. Thus, $\epsilon_{\sigma, \mu}(C_\alpha) = \sum_i \mu(q_i)\epsilon_{\sigma_i, q_i}(C_\alpha)$. If $fstate(\alpha') \in supp(\mu)$, then let j be the index of $fstate(\alpha')$. By the definition of the measure of a cone, $\epsilon_{\sigma, \mu}(C_\alpha) = \epsilon_{\sigma, \mu}(C_{\alpha'}) \sum_{tr \in D(a)} \sigma(\alpha')(tr)\mu_{tr}(q)$. By induction and the definition of σ ,

$$\epsilon_{\sigma, \mu}(C_\alpha) = \sum_i \mu(i)\epsilon_{\sigma_i, q_i}(C_{\alpha'}) \sum_{tr \in D(a)} \sigma_j(\alpha')(tr)\mu_{tr}(q).$$

Since only $\epsilon_{\sigma_j, q_j}(C_{\alpha'})$ may be different from 0, we get

$$\epsilon_{\sigma, \mu}(C_\alpha) = \mu(q_j)\epsilon_{\sigma_j, q_j}(C_{\alpha'}) \sum_{tr \in D(a)} \sigma_j(\alpha')(tr)\mu_{tr}(q) = \mu(q_j)\epsilon_{\sigma_j, q_j}(C_\alpha).$$

For the same reason, $\sum_i \mu(q_i)\epsilon_{\sigma_i, q_i}(C_\alpha) = \mu(q_j)\epsilon_{\sigma_j, q_j}(C_\alpha)$. Thus, $\epsilon_{\sigma, \mu}(C_\alpha) = \sum_i \mu(i)\epsilon_{\sigma_i, q_i}(C_\alpha)$, as needed.

Since each ϵ_{σ_i, q_i} assigns probability 1 to the set of finite execution fragments with trace β , then also $\epsilon_{\sigma, \mu}$ assigns probability 1 to the set of finite execution fragments with trace β . Furthermore, by Proposition 3.1.3, $lstate(\epsilon_{\sigma, \mu}) = \sum_i \mu(q_i)lstate(\epsilon_{\sigma_i, q_i}) = \mu'$. Thus, $lstate(\epsilon_{\sigma, \mu})$ represents $\mu \xrightarrow{\beta} \mu'$.

⇐ Let σ be a scheduler that represents $\mu \xrightarrow{\beta} \mu'$. For each i , let $\sigma_i = \sigma$. Observe that for each finite execution fragment α ,

$$\epsilon_{\sigma, \mu}(C_\alpha) = \mu(fstate(\alpha))\epsilon_{\sigma, fstate(\alpha)}(C_\alpha).$$

Thus, as in the previous case, $\epsilon_{\sigma, \mu} = \sum_i \mu(i)\epsilon_{\sigma_i, q_i}$. Let $q_i \xrightarrow{\beta} \mu_i$ be the weak transition represented by ϵ_{σ_i, q_i} . Since $\mu' = lstate(\epsilon_{\sigma, \mu})$ and for each i , $\mu_i = lstate(\epsilon_{\sigma_i, q_i})$, by Proposition 3.1.3, $\mu' = \sum_i \mu(q_i)\mu_i$. This suffices. □

PROPOSITION 3.6. Suppose that $\mu_1 \xrightarrow{\beta_1} \mu_2$ and $\mu_2 \xrightarrow{\beta_2} \mu_3$ are weak hyper-transitions of a PA \mathcal{P} . Then $\mu_1 \xrightarrow{\beta_1\beta_2} \mu_3$ is a weak hyper-transition of \mathcal{P} .

Proof. Let ϵ_1 and ϵ_2 be the probabilistic execution fragments that represent $\mu_1 \xrightarrow{\beta_1} \mu_2$ and $\mu_2 \xrightarrow{\beta_2} \mu_3$, respectively, and let σ_1 and σ_2 be the schedulers that generate ϵ_1 and ϵ_2 , respectively. Let $N(\alpha) \triangleq \epsilon_1(C_\alpha) + \sum_{\alpha' < \alpha} \epsilon_1(\alpha')\epsilon_2(C_{\alpha \triangleright \alpha'} | C_{\alpha' \triangleright \alpha'})$, where we recall that $\epsilon_2(C_{\alpha \triangleright \alpha'} | C_{\alpha' \triangleright \alpha'})$ denotes the probability of $C_{\alpha \triangleright \alpha'}$ conditional on $C_{\alpha' \triangleright \alpha'}$. Define a new scheduler σ as follows:

$$(3) \quad \sigma(\alpha) = \begin{cases} \frac{\epsilon_1(C_\alpha)\sigma_1(\alpha) + \sum_{\alpha' < \alpha} \epsilon_1(\alpha')\epsilon_2(C_{\alpha \triangleright \alpha'} | C_{\alpha' \triangleright \alpha'})\sigma_2(\alpha \triangleright \alpha')}{N(\alpha)} & \text{if } N(\alpha) > 0, \\ 0 & \text{otherwise.} \end{cases}$$

Informally, σ is a scheduler that represents a concatenation of ϵ_1 and ϵ_2 . A finite execution fragment α can be reached in the concatenation of ϵ_1 and ϵ_2 either because ϵ_1 reaches it or because ϵ_1 terminates at a prefix of α from which ϵ_2 continues. The scheduler σ from α should behave according to σ_1 whenever α is reached in ϵ_1 , and according to σ_2 , with the appropriate argument, whenever α is reached in ϵ_2 . The schedules from α must be weighted by the probabilities with which ϵ_1 and ϵ_2 lead to α . The term $N(\alpha)$ is a normalization factor whose computation is just technical. In the formal proof we have to show that σ is indeed a scheduler (in other words, the value of N is correct), and that σ generates the representation of $\mu_1 \xrightarrow{\beta_1\beta_2} \mu_3$. These proofs are mainly detailed algebraic manipulations. An interesting point of the proof is (4), which expresses the probability of a finite execution fragment in the concatenation of ϵ_1 and ϵ_2 in terms of the probabilities of finite execution fragments of ϵ_1 and ϵ_2 .

We show that σ is a scheduler. That is, for each finite execution fragment α , $\sigma(\alpha)(D) \leq 1$, or equivalently, $\epsilon_1(C_\alpha)\sigma_1(\alpha)(D) + \sum_{\alpha' \leq \alpha} \epsilon_1(\alpha')\epsilon_2(C_{\alpha \triangleright \alpha'} | C_{\alpha' \triangleright \alpha'})\sigma_2(\alpha \triangleright \alpha')(D) \leq N(\alpha)$, where the left-hand term is the numerator of the definition of σ applied to D . Since σ_2 is a scheduler, we know that $\sigma_2(\alpha \triangleright \alpha')(D) \leq 1$. Also, observe that $\epsilon_1(C_\alpha)\sigma_1(\alpha)(D) = \epsilon_1(C_\alpha - \{\alpha\})$. Thus, it suffices to show that $\epsilon_1(C_\alpha - \{\alpha\}) + \sum_{\alpha' \leq \alpha} \epsilon_1(\alpha')\epsilon_2(C_{\alpha \triangleright \alpha'} | C_{\alpha' \triangleright \alpha'}) \leq N(\alpha)$. We separate from the sum the term with $\alpha' = \alpha$ and observe that $\epsilon_2(C_{\alpha \triangleright \alpha} | C_{\alpha \triangleright \alpha}) \leq 1$. The new inequality, which suffices for our purposes, is $\epsilon_1(C_\alpha - \{\alpha\}) + \epsilon_1(\{\alpha\}) + \sum_{\alpha' < \alpha} \epsilon_1(\alpha')\epsilon_2(C_{\alpha \triangleright \alpha'} | C_{\alpha' \triangleright \alpha'}) \leq N(\alpha)$. However, since $\epsilon_1(C_\alpha - \{\alpha\}) + \epsilon_1(\{\alpha\}) = \epsilon_1(C_\alpha)$, the inequality above is $N(\alpha) \leq N(\alpha)$, which is trivially true.

Let ϵ be the probabilistic execution fragment generated by σ and μ_1 . We show that ϵ represents $\mu_1 \xrightarrow{\beta_1\beta_2} \mu_3$ in three steps. First we show by induction on the length of a finite execution fragment α that $\epsilon(C_\alpha) = N(\alpha)$. The base case is trivial since $N(q) = \epsilon_1(C_q)$ by definition of N , and $\epsilon_1(C_q) = \epsilon(C_q) = \mu_1(q)$ by definition of ϵ_1 and ϵ since both measures are generated by μ_1 . For the inductive step, let $\alpha = \alpha'a'q'$. If $\epsilon(C_{\alpha'}) = 0$, then by the definition of the measure of a cone, in (1), $\epsilon(C_\alpha) = 0$. We show that $N(\alpha) = 0$ as well. By induction, since $\epsilon(C_{\alpha'}) = 0$, $N(\alpha') = 0$. By definition of N , $\epsilon_1(C_{\alpha'}) = 0$, and $\epsilon_1(\alpha'')\epsilon_2(C_{\alpha' \triangleright \alpha''} | C_{\alpha'' \triangleright \alpha''}) = 0$ for each $\alpha'' < \alpha'$. Then, by the definition of a conditional measure and of a measure of a cone, $\epsilon_1(C_\alpha) = 0$ and $\epsilon_1(\alpha'')\epsilon_2(C_{\alpha \triangleright \alpha''} | C_{\alpha'' \triangleright \alpha''}) = 0$ for each $\alpha'' < \alpha'$. By $\epsilon_1(C_{\alpha'}) = 0$, also $\epsilon_1(\alpha') = 0$, and hence $\epsilon_1(\alpha')\epsilon_2(C_{\alpha \triangleright \alpha'} | C_{\alpha' \triangleright \alpha'}) = 0$. We have shown that all the terms of the definition of $N(\alpha)$ are 0, and thus $N(\alpha) = 0$ as needed. If $\epsilon(C_{\alpha'}) > 0$, then by expanding σ with its definition in (1), the definition of the measure of a cone, we get that $\epsilon(C_\alpha)/\epsilon(C_{\alpha'})$ equals

$$\sum_{tr \in D(\alpha)} \frac{\epsilon_1(C_{\alpha'})\sigma_1(\alpha')(tr) + \sum_{\alpha'' \leq \alpha'} \epsilon_1(\alpha'')\epsilon_2(C_{\alpha' \triangleright \alpha''} | C_{\alpha'' \triangleright \alpha''})\sigma_2(\alpha' \triangleright \alpha'')(tr)}{N(\alpha')} \mu_{tr}(q').$$

By induction, $\epsilon(C_{\alpha'}) = N(\alpha')$, and thus the two terms can be simplified in the equation above. Then, by rearranging terms algebraically, we get

$$\begin{aligned} \epsilon(C_\alpha) &= \sum_{tr \in D(\alpha)} \epsilon_1(C_{\alpha'})\sigma_1(\alpha')(tr)\mu_{tr}(q') \\ &+ \sum_{\alpha'' \leq \alpha'} \epsilon_1(\alpha'') \sum_{tr \in D(\alpha)} \epsilon_2(C_{\alpha' \triangleright \alpha''} | C_{\alpha'' \triangleright \alpha''})\sigma_2(\alpha' \triangleright \alpha'')(tr)\mu_{tr}(q'). \end{aligned}$$

By the definition of the measure of a cone, the first term above is $\epsilon_1(C_\alpha)$. With a similar argument, after applying the definition of a conditional measure and distinguishing the cases where $\epsilon_2(C_{\alpha''\triangleright\alpha''}) = 0$, the second term above is $\sum_{\alpha'' \leq \alpha'} \epsilon_1(\alpha'')\epsilon_2(C_{\alpha\triangleright\alpha''} \mid C_{\alpha''\triangleright\alpha''})$. Thus, we get

$$\epsilon(C_\alpha) = \epsilon_1(C_\alpha) + \sum_{\alpha'' \leq \alpha'} \epsilon_1(\alpha'')\epsilon_2(C_{\alpha\triangleright\alpha''} \mid C_{\alpha''\triangleright\alpha''}).$$

Then it is enough to observe that the right-hand side of the equation above is $N(\alpha)$ since $\alpha'' \leq \alpha'$ iff $\alpha'' < \alpha$.

Second we show that for each finite execution fragment α ,

$$(4) \quad \epsilon(\alpha) = \sum_{\alpha' < \alpha} \epsilon_1(\alpha')\epsilon_2(\alpha \triangleright \alpha' \mid C_{\alpha'\triangleright\alpha'}).$$

Observe that, by the definition of the cone σ -field, $\epsilon(\alpha) = \epsilon(C_\alpha) - \sum_{a \in A, q \in Q} \epsilon(C_{\alpha a q})$. By replacing the ϵ measures of cones with the definition of N in the equation above, we get

$$\begin{aligned} \epsilon(\alpha) &= \epsilon_1(C_\alpha) + \sum_{\alpha' < \alpha} \epsilon_1(\alpha')\epsilon_2(C_{\alpha\triangleright\alpha'} \mid C_{\alpha'\triangleright\alpha'}) \\ &\quad - \sum_{a \in A, q \in Q} \left(\epsilon_1(C_{\alpha a q}) + \sum_{\alpha' < \alpha a q} \epsilon_1(\alpha')\epsilon_2(C_{\alpha a q \triangleright \alpha'} \mid C_{\alpha'\triangleright\alpha'}) \right). \end{aligned}$$

We now use the terms $\epsilon_1(C_\alpha)$ and $-\sum_{a \in A, q \in Q} \epsilon_1(C_{\alpha a q})$ in the equation above to derive $\epsilon_1(\alpha)$, and similarly we use part of the other two terms to derive the following:

$$\epsilon(\alpha) = \epsilon_1(\alpha) + \sum_{\alpha' < \alpha} \epsilon_1(\alpha')\epsilon_2(\alpha \triangleright \alpha' \mid C_{\alpha'\triangleright\alpha'}) - \sum_{a \in A, q \in Q} \epsilon_1(\alpha)\epsilon_2(C_{\alpha a q \triangleright \alpha} \mid C_{\alpha \triangleright \alpha}).$$

Observe that the third term in the equation above is $\epsilon_1(\alpha)\epsilon_2((C_{\alpha \triangleright \alpha} - \{\alpha \triangleright \alpha\}) \mid C_{\alpha \triangleright \alpha})$. Thus, by adding and subtracting the term $\epsilon_1(\alpha)\epsilon_2(\alpha \triangleright \alpha \mid C_{\alpha \triangleright \alpha})$ and rearranging algebraically, we get

$$\epsilon(\alpha) = \sum_{\alpha' \leq \alpha} \epsilon_1(\alpha')\epsilon_2(\alpha \triangleright \alpha' \mid C_{\alpha'\triangleright\alpha'}) + \epsilon_1(\alpha)(1 - \epsilon_2(C_{\alpha \triangleright \alpha} \mid C_{\alpha \triangleright \alpha})).$$

If $\epsilon_1(\alpha) = 0$, then (4) follows trivially. If $\epsilon_1(\alpha) > 0$, then, since ϵ_1 represents $\mu_1 \xrightarrow{\beta_1} \mu_2$, $\mu_2(q) > 0$, where q is $lstate(\alpha)$. Since ϵ_2 is generated by σ_2 and μ_2 , $\epsilon_2(C_q) = \mu_2(q) > 0$. By definition of \triangleright , $\alpha \triangleright \alpha = q$, and thus $\epsilon_2(C_{\alpha \triangleright \alpha}) > 0$. By the definition of the conditional measure, $\epsilon_2(C_{\alpha \triangleright \alpha} \mid C_{\alpha \triangleright \alpha}) = 1$, which leads again to (4).

Third we show that ϵ represents $\mu_1 \xrightarrow{\beta_1 \beta_2} \mu_3$. Let α be such that $\epsilon(\alpha) > 0$. By (4) there exists a prefix α' of α such that $\epsilon_1(\alpha') > 0$ and $\epsilon_2(\alpha' \triangleright \alpha) > 0$. Then $trace(\alpha') = \beta_1$ and $trace(\alpha \triangleright \alpha') = \beta_2$. Thus, $trace(\alpha) = \beta_1 \beta_2$. Let μ be $lstate(\epsilon)$. We are left to show that $\mu = \mu_3$. Consider a state q of Q . By the definition of an image measure and (4), $\mu(q) = \sum_{\alpha \mid lstate(\alpha)=q} \sum_{\alpha' \leq \alpha} \epsilon_1(\alpha')\epsilon_2(\alpha \triangleright \alpha' \mid C_{\alpha'\triangleright\alpha'})$. The sum above can be restructured as follows:

$$\mu(q) = \sum_{\alpha'} \sum_{\alpha \mid fstate(\alpha)=lstate(\alpha'), lstate(\alpha)=q} \epsilon_1(\alpha')\epsilon_2(\alpha \mid C_{\alpha'\triangleright\alpha'}).$$




FIG. 3.1. Trace distribution inclusion is not preserved by composition (without communication).

We further partition the sums over $lstate(\alpha')$, thus getting

$$\mu(q) = \sum_{q' \in Q} \sum_{\alpha' | lstate(\alpha')=q'} \sum_{\alpha | fstate(\alpha)=q', lstate(\alpha)=q} \epsilon_1(\alpha') \epsilon_2(\alpha | C_{q'}).$$

Observe that, since $\mu_2 = lstate(\epsilon_1)$, if $\mu_2(q') = 0$, then there is no α' with last state q' such that $\epsilon_1(\alpha') > 0$. Thus we can restrict the first sum to those q' such that $\mu_2(q') > 0$. Also, if $\mu_2(q') > 0$, we have already concluded before that $\epsilon_2(C_{\alpha' \triangleright \alpha'}) = \mu_2(q')$. Thus, we get

$$\mu(q) = \sum_{q' \in Q | \mu_2(q') > 0} \sum_{\alpha' | lstate(\alpha')=q'} \sum_{\alpha | fstate(\alpha)=q', lstate(\alpha)=q} \epsilon_1(\alpha') \epsilon_2(\alpha) / \mu_2(q').$$

Observe that the two inner sums can be exchanged. By definition of μ_2 it follows that $\sum_{\alpha' | lstate(\alpha')=q'} \epsilon_1(\alpha') = \mu_2(q')$. Thus, we get

$$\mu(q) = \sum_{q' \in Q | \mu_2(q') > 0} \sum_{\alpha | fstate(\alpha)=q', lstate(\alpha)=q} \epsilon_2(\alpha).$$

Following the same argument that we used to restrict the sum over q' , we can remove such restriction, and thus we can remove the most external sum, leading to $\mu(q) = \sum_{\alpha | lstate(\alpha)=q} \epsilon_2(\alpha)$, which is the definition of $\mu_3(q)$. \square

3.4. Composition. Two PAs, \mathcal{P}_1 and \mathcal{P}_2 , are *compatible* if $H_1 \cap A_2 = A_1 \cap H_2 = \emptyset$. The (*parallel*) *composition* of two compatible PAs \mathcal{P}_1 and \mathcal{P}_2 , denoted by $\mathcal{P}_1 \parallel \mathcal{P}_2$, is the PA $\mathcal{P} = (Q_1 \times Q_2, (\bar{q}_1, \bar{q}_2), E_1 \cup E_2, H_1 \cup H_2, D)$, where D is the set of triples $(q, a, \mu_1 \times \mu_2)$ such that, for $i \in \{1, 2\}$,

$$a \in A_i \Rightarrow (\pi_i(q), a, \mu_i) \in D_i \quad \text{and} \quad a \notin A_i \Rightarrow \mu_i = \delta(\pi_i(q)).$$

Let ϵ be a probabilistic execution (fragment) of $\mathcal{P}_1 \parallel \mathcal{P}_2$ and let $i \in \{1, 2\}$. Define $\pi_i(\epsilon)$, the i th projection of ϵ , to be the image measure under π_i of ϵ . It is easy to verify that the projection function is measurable. When convenient, we denote a projection by $\epsilon[\mathcal{P}_i]$, where \mathcal{P}_i is the PA that appears in the i th position.

PROPOSITION 3.7. *Let \mathcal{P}_1 and \mathcal{P}_2 be compatible PAs, and let ϵ be a probabilistic execution (fragment) of $\mathcal{P}_1 \parallel \mathcal{P}_2$. Then for each $i \in \{1, 2\}$, $\pi_i(\epsilon)$ is a probabilistic execution (fragment) of \mathcal{P}_i .*

Proof. The proof follows by Propositions 4.3.4 and 4.3.5 of [32]. \square

The trace distribution preorder is not preserved by composition [34, 37], as is shown by the following example.




FIG. 3.2. Trace distribution inclusion is not preserved by composition (with communication).

Example 3.2 (failure of compositionality). Consider the two (nondeterministic) automata \mathcal{P}_1 and \mathcal{P}_2 of Figure 3.1. The two automata are trace equivalent, and it is easy to see that they are also trace distribution equivalent. Now consider the compositions $\mathcal{P}_1 \parallel \mathcal{C}$ and $\mathcal{P}_2 \parallel \mathcal{C}$, where \mathcal{C} is the PA of Figure 3.1 and we assume that the actions of \mathcal{C} are not shared with \mathcal{P}_1 and \mathcal{P}_2 . It is possible to build a probabilistic execution of $\mathcal{P}_1 \parallel \mathcal{C}$ as follows: first a is scheduled followed by d ; then e or f is scheduled, depending on the outcome state of the transition labeled by d ; finally, b or c is scheduled, depending on whether e or f was scheduled. Formally, we consider the probabilistic execution induced by the deterministic memoryless scheduler specified by the following partial function, where a transition is denoted by the unique action labeling it:

| Q_1 | Q_C | $D_{\mathcal{P}_1 \parallel \mathcal{C}}$ |
|-------|-------|---|
| s_1 | t_1 | a |
| s_2 | t_1 | d |
| s_2 | t_2 | e |
| s_2 | t_3 | f |
| s_2 | t_4 | b |
| s_2 | t_5 | c |

Thus, in the resulting trace distribution there is a total correlation between e, b and f, c , respectively. The same trace distribution cannot be obtained from $\mathcal{P}_2 \parallel \mathcal{C}$ because after scheduling the transition labeled by a we are already bound to b or c , and thus the occurrence of b or c cannot be correlated to e or f in this case.

Example 3.2 may appear pathological since, in the probabilistic execution of $\mathcal{P}_1 \parallel \mathcal{C}$ that correlates the choices between e and f and between b and c , a nondeterministic choice of \mathcal{P}_1 is resolved based on information that is not available to \mathcal{P}_1 . This may lead us to propose a naive solution to the nonpreservation of trace distribution inclusion by parallel composition, where we require that each PA in a parallel composition be able to resolve its nondeterministic choices based on local knowledge only. However, a more elaborate example shows that this naive idea also does not work.

Example 3.3 (failure of compositionality). Consider the two automata \mathcal{P}_1 and \mathcal{P}_2 of Figure 3.2, which are essentially the automata of Example 3.2 where self-loop transitions labeled by e and f are added to each state. In this case the context \mathcal{C} synchronizes with \mathcal{P}_1 and \mathcal{P}_2 on actions e and f , and \mathcal{P}_1 is able to learn which of e or f occurs, thus determining the correlation with b and c based on local knowledge only.

The solution of resolving nondeterminism based on local knowledge is adopted in [10] for a probabilistic extension of reactive modules; however, the idea of [10] cannot be extended easily to PAs because of key structural differences in the models: in PAs there is a total interleaving of the transitions taken by different PAs in a parallel composition, while in probabilistic reactive modules there are several independent *atoms* that are not forced to interleave. A direct adaptation of the idea of [10] to PAs would require drastic modifications of the model that go beyond the scope of this paper: transitions would have to be labeled by sets of actions and be structured in such a way that each action affects different parts of the state.

An alternative approach, followed in [32] and adopted in this paper, consists of defining a new *trace distribution precongruence* relation, denoted by \leq_{DC} , as the coarsest precongruence (for parallel composition) that is included in the trace distribution preorder \leq_D , and finding alternative characterizations of \leq_{DC} . It is known from [32] that there exists a simple context, called the *principal context*, that is sufficiently powerful to distinguish all PAs that are not in the trace distribution precongruence relation; alternatively, a testing scenario is proposed in [33].

In this paper we characterize \leq_{DC} in terms of probabilistic simulation relations. Another simple alternative characterization of \leq_{DC} that is useful for our study is given by the following proposition.

PROPOSITION 3.8. *Let \mathcal{P}_1 and \mathcal{P}_2 be PAs. Then $\mathcal{P}_1 \leq_{DC} \mathcal{P}_2$ iff for every PA \mathcal{C} that is compatible with both \mathcal{P}_1 and \mathcal{P}_2 , $\mathcal{P}_1 \parallel \mathcal{C} \leq_D \mathcal{P}_2 \parallel \mathcal{C}$.*

Proof. Define relation \sqsubseteq such that $\mathcal{P}_1 \sqsubseteq \mathcal{P}_2$ iff for every PA \mathcal{C} that is compatible with both \mathcal{P}_1 and \mathcal{P}_2 , $\mathcal{P}_1 \parallel \mathcal{C} \leq_D \mathcal{P}_2 \parallel \mathcal{C}$.

Let $\mathcal{P}_1 \leq_{DC} \mathcal{P}_2$ and let \mathcal{C} be a PA compatible with both \mathcal{P}_1 and \mathcal{P}_2 . Since \leq_{DC} is a precongruence by definition, then $\mathcal{P}_1 \parallel \mathcal{C} \leq_{DC} \mathcal{P}_2 \parallel \mathcal{C}$. Since, again by definition, \leq_{DC} is included in \leq_D , then $\mathcal{P}_1 \parallel \mathcal{C} \leq_D \mathcal{P}_2 \parallel \mathcal{C}$. Thus, $\mathcal{P}_1 \sqsubseteq \mathcal{P}_2$, which implies that \leq_{DC} is included in \sqsubseteq .

Conversely, observe that \sqsubseteq is reflexive and transitive, and thus a preorder relation. Observe also that, by using a trivial context \mathcal{C} with no external actions and no transitions, \sqsubseteq is included in \leq_D . Finally, using the associativity of parallel composition, observe that \sqsubseteq is preserved by parallel composition, and thus is a precongruence. This means that \sqsubseteq is a precongruence included in \leq_D . Since \leq_{DC} is the coarsest precongruence included in \leq_D , we get that \sqsubseteq is included in \leq_{DC} . \square

3.5. Simulation relations. The definitions of forward simulation and weak forward simulation in section 2 can be extended naturally to PAs [34]. However, Segala has shown [31] that the resulting simulations are not complete for \leq_{DC} , and has defined new candidate simulations. These new simulations relate states to probability measures on states.

In order to define the new simulations formally, we need two new concepts. First we show how to *lift* a relation between sets to a relation between measures over sets [18]. Let $R \subseteq X \times Y$. The *lifting* of R is a relation $R' \subseteq \text{Disc}(X) \times \text{Disc}(Y)$ such that $\mu_X R' \mu_Y$ iff there is a function $w : X \times Y \rightarrow [0, 1]$ that satisfies the following:

1. If $w(x, y) > 0$, then $x R y$.
2. For each $x \in X$, $\sum_{y \in Y} w(x, y) = \mu_X(x)$.
3. For each $y \in Y$, $\sum_{x \in X} w(x, y) = \mu_Y(y)$.

We abuse notation slightly and denote the lifting of a relation R by R as well. Second, we define a *flattening* operation that converts a measure μ in $\text{Disc}(\text{Disc}(X))$ into a measure $\text{flatten}(\mu)$ in $\text{Disc}(X)$. Namely, we define $\text{flatten}(\mu) = \sum_{\rho \in \text{supp}(\mu)} \mu(\rho)\rho$.




FIG. 3.3. A forward probabilistic simulation between two PAs.

We now define simulations for PAs. A relation $R \subseteq Q_1 \times Disc(Q_2)$ is a *probabilistic forward simulation* (resp., *weak probabilistic forward simulation*) from PA \mathcal{P}_1 to PA \mathcal{P}_2 iff $E_1 = E_2$ and both of the following hold:

1. $\bar{q}_1 R \delta(\bar{q}_2)$.
2. For each pair q_1, μ_2 such that $q_1 R \mu_2$ and each transition $q_1 \xrightarrow{a} \mu'_1$, there exists a measure $\xi'_2 \in Disc(Disc(Q_2))$ such that $\mu'_1 R \xi'_2$ and such that $\mu_2 \xrightarrow{a} flatten(\xi'_2)$ (resp., $\mu_2 \xrightarrow{a} flatten(\xi'_2)$) is a hyper-transition (resp., a weak hyper-transition) of \mathcal{P}_2 .

We write $\mathcal{P}_1 \leq_{PF} \mathcal{P}_2$ (resp., $\mathcal{P}_1 \leq_{wPF} \mathcal{P}_2$) whenever there is a probabilistic forward simulation (resp., a weak probabilistic forward simulation) from \mathcal{P}_1 to \mathcal{P}_2 .

Example 3.4 (forward probabilistic simulation). Figure 3.3 gives an example of two PAs that are in the kernel of probabilistic forward simulation. However, there would be no simulation from \mathcal{P}_1 to \mathcal{P}_2 if we did not allow states to be related to measures over states. The probabilistic forward simulation R from \mathcal{P}_1 to \mathcal{P}_2 relates each state of \mathcal{P}_1 with the Dirac measure over its primed version of \mathcal{P}_2 , relates s_1 with the uniform measure over s'_3 and s'_4 , and relates s_2 with the uniform measure over s'_5 and s'_6 . The transition from s_0 can be simulated from s'_0 by scheduling the only transition enabled. Indeed, the target measure $\mathcal{U}(s'_3, s'_4, s'_5, s'_6)$ is the flattening of $\mathcal{U}(\mathcal{U}(s'_3, s'_4), \mathcal{U}(s'_5, s'_6))$, and it is easy to check that $\mathcal{U}(s_1, s_2) R \mathcal{U}(\mathcal{U}(s'_3, s'_4), \mathcal{U}(s'_5, s'_6))$.

Note that a forward simulation between nondeterministic automata is a probabilistic forward simulation between the two automata viewed as PAs, as described next.

PROPOSITION 3.9. *Let \mathcal{A}_1 and \mathcal{A}_2 be nondeterministic automata. Then*

1. $\mathcal{A}_1 \leq_F \mathcal{A}_2$ iff $\mathcal{A}_1 \leq_{PF} \mathcal{A}_2$, and
2. $\mathcal{A}_1 \leq_{wF} \mathcal{A}_2$ iff $\mathcal{A}_1 \leq_{wPF} \mathcal{A}_2$.

Proof. The left-to-right inclusions are easy since, given a (weak) forward simulation R from \mathcal{A}_1 to \mathcal{A}_2 , it is routine to check that the relation $R' \triangleq \{(q_1, \delta(q_2)) \mid q_1 R q_2\}$ is a (weak) probabilistic forward simulation from \mathcal{A}_1 to \mathcal{A}_2 .

For the converse implication, let R be a (weak) probabilistic forward simulation from \mathcal{A}_1 to \mathcal{A}_2 . Define a relation $R' \triangleq \{(q_1, q_2) \mid \exists \mu q_1 R \mu, q_2 \in \text{supp}(\mu)\}$. We show that R' is a (weak) forward simulation from \mathcal{A}_1 to \mathcal{A}_2 .

The start condition is trivial since $\bar{q}_1 R \delta(\bar{q}_2)$, and thus $\bar{q}_1 R' \bar{q}_2$. For the step condition, let $q_1 R' q_2$, and let $q_1 \xrightarrow{a} q'_1$. By the definition of R' , there exists a measure

μ such that $q_1 R \mu$ and $q_2 \in \text{supp}(\mu)$. Since R is a (weak) forward simulation, there exists a hyper-transition $\mu \xrightarrow{a} \mu'$ (a weak hyper-transition $\mu \xrightarrow{a} \mu'$), where μ' is the flattening of some measure μ'' such that $\delta(q'_1) R \mu''$. By the definition of hyper-transition, there is a combined transition $q_2 \xrightarrow{a} \mu_2$ (a weak combined transition $q_2 \xrightarrow{a} \mu_2$) such that $\text{supp}(\mu_2) \subseteq \text{supp}(\mu')$. For the strong case, let $q_2 \xrightarrow{a} q'_2$ be one of the transitions of D_2 that are combined in $q_2 \xrightarrow{a} \mu_2$. Then, $q'_2 \in \text{supp}(\mu_2)$. For the weak case, consider a scheduler σ that generates $q_2 \xrightarrow{a} \mu_2$, and build a new scheduler σ' that on input α stops (does not return any transition) if $\sigma(\alpha)$ stops with some nonzero probability and chooses any transition in $\text{supp}(\sigma(\alpha))$ that reduces the distance from a stopping point otherwise. This leads to a weak transition $q_2 \xrightarrow{a} q'_2$, where $q'_2 \in \text{supp}(\mu_2)$. We now show that $q'_1 R' q'_2$, which suffices. Since $q'_2 \in \text{supp}(\mu_2)$, and since $\text{supp}(\mu_2) \subseteq \text{supp}(\mu')$, then $q'_2 \in \text{supp}(\mu')$. Since $\mu' = \text{flatten}(\mu'')$, then q'_2 is also in the support of some measure $\rho \in \text{supp}(\mu'')$. Thus, $q'_1 R \rho$, and, by the definition of R' , $q'_1 R' q'_2$, as needed. \square

PROPOSITION 3.10. *Let \mathcal{P}_1 and \mathcal{P}_2 be PAs. Then the following hold:*

1. *If $\mathcal{P}_1 \leq_{PF} \mathcal{P}_2$, then $\mathcal{P}_1 \leq_{wPF} \mathcal{P}_2$.*
2. *If $H_1 = H_2 = \emptyset$, then $\mathcal{P}_1 \leq_{PF} \mathcal{P}_2$ iff $\mathcal{P}_1 \leq_{wPF} \mathcal{P}_2$.*
3. *If $\mathcal{P}_1 \leq_{wPF} \mathcal{P}_2$, then $\mathcal{P}_1 \leq_{DC} \mathcal{P}_2$.*

Proof. The first item follows from the fact that a combined transition is a special case of a weak combined transition; the second item follows from the fact that in the absence of internal actions a weak combined transition is a combined transition. The proof of the third item is quite involved, and we refer the reader to Proposition 8.7.1 of [32]. The main idea is to use the weak probabilistic forward simulation from \mathcal{P}_1 to \mathcal{P}_2 to build, for each probabilistic execution of \mathcal{P}_1 , a corresponding probabilistic execution of \mathcal{P}_2 with the same trace distribution. \square

3.6. Tree-structured PAs. A *path* of a PA \mathcal{P} is a finite sequence $\gamma = q_0 a_1 \mu_1 q_1 a_2 \mu_2 q_2 \dots q_n$ of alternating states, actions, and distribution over states, starting with the start state of \mathcal{P} such that for each nonfinal i , $q_i \xrightarrow{a_{i+1}} \mu_{i+1}$ and $q_{i+1} \in \text{supp}(\mu_{i+1})$. We write $\text{lstate}(\gamma)$ to denote q_n and $\text{paths}(\mathcal{P})$ for the set of all paths of \mathcal{P} . We say that \mathcal{P} is *tree-structured* if each state is reached via a unique path. Tree-structured PAs are characterized uniquely by the property that all states are reachable, the start state does not occur in the target of any transition, and each of the other states occurs in the target of exactly one transition. Tree-structured nondeterministic automata are also characterized uniquely by this property, albeit for a different notion of transition.

If a PA is tree-structured, then its underlying nondeterministic automaton is also tree-structured. The following example shows that the converse does not hold.

Example 3.5 (non-tree-structured PAs). Figure 3.4 shows a PA that is not tree-structured, as state q' can be reached via two different paths. The underlying nondeterministic automaton is tree-structured, however, since the only way to reach state q' is via the execution qaq' .

The *unfolding* of a PA \mathcal{P} , denoted by $\text{Unfold}(\mathcal{P})$, is the tree-structured PA \mathcal{Q} obtained from \mathcal{P} by unfolding its transition graph into a tree. Formally,

- $Q_{\mathcal{Q}} = \text{paths}(\mathcal{P})$,
- $\bar{q}_{\mathcal{Q}} = \bar{q}_{\mathcal{P}}$,
- $E_{\mathcal{Q}} = E_{\mathcal{P}}$,
- $H_{\mathcal{Q}} = H_{\mathcal{P}}$, and
- $D_{\mathcal{Q}} = \{(\gamma, a, \mu) \mid (\exists \mu')[(\text{lstate}(\gamma), a, \mu') \in D_{\mathcal{P}} \wedge (\forall q \in \text{supp}(\mu'))[\mu'(q) = \mu(\gamma a \mu' q)]]\}$.




FIG. 3.4. The PA on the left is not tree-structured even though its underlying nondeterministic automaton on the right is.

PROPOSITION 3.11. $\mathcal{P} \equiv_{PF} \text{Unfold}(\mathcal{P})$.

Proof. It is easy to check that the relation R , where $\alpha R \delta(q)$ iff $\text{lstate}(\alpha) = q$ is a probabilistic forward simulation from $\text{Unfold}(\mathcal{P})$ to \mathcal{P} and that the “inverse” of R , that is, the relation R' such that $q R' \delta(\alpha)$ iff $\text{lstate}(\alpha) = q$, is a probabilistic forward simulation from \mathcal{P} to $\text{Unfold}(\mathcal{P})$. \square

PROPOSITION 3.12. $\mathcal{P} \equiv_{DC} \text{Unfold}(\mathcal{P})$.

Proof. The proof follows by Proposition 3.11 and Proposition 3.10, parts 1 and 3. \square

3.7. Truncations and continuations. We now define two simple constructions on probabilistic execution fragments that will be useful for our proofs. Specifically, we define the truncation of a probabilistic execution fragment, which is the result of stopping the computation at some designated points, and the continuation of a probabilistic execution fragment, which represents the rest of a probabilistic execution fragment after some finite execution fragment has occurred.

Let ϵ be a probabilistic execution fragment of a PA \mathcal{P} , generated by some scheduler σ , and let Θ be a set of finite execution fragments of \mathcal{P} . Define the *truncation of ϵ at Θ* to be the same as ϵ except that no transition is scheduled from all the Θ places, that is, the probabilistic execution fragment ϵ' , with the same start state as ϵ , generated by a new scheduler σ' such that $\sigma'(\alpha) = \sigma(\alpha)$ if $\alpha \notin \Theta$ and $\sigma'(\alpha)(D) = 0$ if $\alpha \in \Theta$.

PROPOSITION 3.13. *The definition of truncation of a probabilistic execution fragment ϵ is independent of the choice of the inducing scheduler.*

Proof. Let μ be the first state of ϵ , and let σ_1, σ_2 be two schedulers that, together with μ , induce ϵ . Let Θ be a set of finite execution fragments of \mathcal{P} , and let σ'_1, σ'_2 be the schedulers built from σ_1, σ_2 , respectively, according to the definition of truncation. Let ϵ_1, ϵ_2 be the induced probabilistic execution fragments, and suppose by contradiction that $\epsilon_1 \neq \epsilon_2$. Then there exists a finite execution α such that $\epsilon_1(C_\alpha) \neq \epsilon_2(C_\alpha)$. Consider such a finite execution α of minimum length. Observe that $|\alpha| > 0$ since $\epsilon(C_q) = \epsilon_1(C_q) = \epsilon_2(C_q) = \mu(q)$ for each state $q \in Q$. Thus, $\alpha = \alpha' a' q'$ for some α', a', q' , where $\epsilon_1(C_{\alpha'}) = \epsilon_2(C_{\alpha'})$. We distinguish two cases.

Case 1. If $\alpha' \in \Theta$, then, by the definitions of σ'_1 and σ'_2 , $\sigma'_1(\alpha')(D) = \sigma'_2(\alpha')(D) = 0$. Thus, $\epsilon_1(C_\alpha) = \epsilon_2(C_\alpha) = 0$, a contradiction.

Case 2. If $\alpha' \notin \Theta$, then, by the definitions of σ'_1 and σ'_2 , $\sigma'_1(\alpha') = \sigma_1(\alpha')$ and

$\sigma'_2(\alpha') = \sigma_2(\alpha')$. Then,

$$\begin{aligned}
\epsilon_1(C_\alpha) &= \epsilon_1(C_{\alpha'}) \sum_{tr \in D(\alpha')} \sigma'_1(\alpha')(tr) \mu_{tr}(q') \quad (\text{by (1)}) \\
&= \epsilon_2(C_{\alpha'}) \sum_{tr \in D(\alpha')} \sigma_1(\alpha')(tr) \mu_{tr}(q') \\
&\quad (\text{by } \sigma'_1(\alpha') = \sigma_1(\alpha'), \epsilon_1(C_{\alpha'}) = \epsilon_2(C_{\alpha'})) \\
&= \epsilon_2(C_{\alpha'}) \sum_{tr \in D(\alpha')} \sigma_2(\alpha')(tr) \mu_{tr}(q') \quad (\text{by } \sigma_1 \text{ and } \sigma_2 \text{ induce } \epsilon) \\
&= \epsilon_2(C_{\alpha'}) \sum_{tr \in D(\alpha')} \sigma'_2(\alpha')(tr) \mu_{tr}(q') \quad (\text{by } \sigma'_2(\alpha') = \sigma_2(\alpha')) \\
&= \epsilon_2(C_\alpha) \quad (\text{by (1)}),
\end{aligned}$$

again a contradiction. \square

Let ϵ be a probabilistic execution fragment of a PA \mathcal{P} , generated by a scheduler σ , and let α be a finite execution fragment with $fstate(\alpha) \in \text{supp}(fstate(\epsilon))$. Define $\epsilon \triangleright \alpha$, the *continuation of ϵ after prefix α* , to be the probabilistic execution fragment generated by the following scheduler σ' from $lstate(\alpha)$:

$$\sigma'(\alpha') = \begin{cases} \sigma(\alpha \frown \alpha') & \text{if } fstate(\alpha') = lstate(\alpha), \\ 0 & \text{otherwise,} \end{cases}$$

where by 0 we denote the identically 0 function.

PROPOSITION 3.14. *The definition of $\epsilon \triangleright \alpha$ is independent of the choice of the inducing scheduler.*

Proof. The proof proceeds in a manner similar to that of Proposition 3.13. Let μ be the first state of ϵ , and let σ_1, σ_2 be two schedulers that, together with μ , induce ϵ . Let q' be $lstate(\alpha)$. Let σ'_1, σ'_2 be the schedulers built from σ_1, σ_2 , respectively, according to the definition of $\epsilon \triangleright \alpha$. Let ϵ_1, ϵ_2 be the induced probabilistic execution fragments from q' , and suppose by contradiction that $\epsilon_1 \neq \epsilon_2$. Then there exists a finite execution α' such that $\epsilon_1(C_{\alpha'}) \neq \epsilon_2(C_{\alpha'})$. Consider such a finite execution α' of minimum length. Observe that $|\alpha'| > 0$ since $\epsilon(C_{q'}) = \epsilon_1(C_{q'}) = \epsilon_2(C_{q'}) = 1$ and, for each state $q'' \neq q'$, $\epsilon(C_{q''}) = \epsilon_1(C_{q''}) = \epsilon_2(C_{q''}) = 0$. Thus, $\alpha' = \alpha'' \alpha'' q''$ for some α'', α'', q'' , where $\epsilon_1(C_{\alpha''}) = \epsilon_2(C_{\alpha''})$. We distinguish two cases.

Case 1. If $fstate(\alpha'') \neq q'$, then, by the definitions of ϵ_1 and ϵ_2 , $\epsilon_1(C_{\alpha'}) = \epsilon_2(C_{\alpha'}) = 0$, a contradiction.

Case 2. If $fstate(\alpha'') = q'$, then, by the definitions of σ'_1 and σ'_2 , $\sigma'_1(\alpha'') = \sigma_1(\alpha \frown \alpha'')$ and $\sigma'_2(\alpha'') = \sigma_2(\alpha \frown \alpha'')$. Then,

$$\begin{aligned}
\epsilon_1(C_{\alpha'}) &= \epsilon_1(C_{\alpha''}) \sum_{tr \in D(\alpha'')} \sigma'_1(\alpha'')(tr) \mu_{tr}(q'') \quad (\text{by (1)}) \\
&= \epsilon_2(C_{\alpha''}) \sum_{tr \in D(\alpha'')} \sigma_1(\alpha \frown \alpha'')(tr) \mu_{tr}(q'') \\
&\quad (\text{by } \sigma'_1(\alpha'') = \sigma_1(\alpha \frown \alpha'') \text{ and } \epsilon_1(C_{\alpha''}) = \epsilon_2(C_{\alpha''})) \\
&= \epsilon_2(C_{\alpha''}) \sum_{tr \in D(\alpha'')} \sigma_2(\alpha \frown \alpha'')(tr) \mu_{tr}(q'') \quad (\text{by } \sigma_1 \text{ and } \sigma_2 \text{ induce } \epsilon) \\
&= \epsilon_2(C_{\alpha''}) \sum_{tr \in D(\alpha'')} \sigma'_2(\alpha'')(tr) \mu_{tr}(q'') \quad (\text{by } \sigma'_2(\alpha'') = \sigma_2(\alpha \frown \alpha'')) \\
&= \epsilon_2(C_{\alpha'}) \quad (\text{by (1)}),
\end{aligned}$$

again a contradiction. \square

The following proposition relates the continuation of ϵ after some prefix α with ϵ itself. In practice it states that $\epsilon \triangleright \alpha$ is closely related to $\epsilon \upharpoonright C_\alpha$.

PROPOSITION 3.15. *Let ϵ be a probabilistic execution fragment of a PA \mathcal{P} , and let α be a finite execution fragment of \mathcal{P} . Then, for each finite execution α' with $lstate(\alpha) = fstate(\alpha')$, $\epsilon(C_{\alpha \frown \alpha'}) = \epsilon(C_\alpha) \cdot (\epsilon \triangleright \alpha)(C_{\alpha'})$.*

Proof. The proof follows easily by induction on the length of α' from the definition of the probability of a cone. \square

4. Tester automata and observer schedulers. The proofs of our completeness results rely on a special context for a PA, which we call its *tester PA*. The tester automaton, $tester(\mathcal{P})$, of a PA \mathcal{P} can observe the states that \mathcal{P} goes through and the transitions that are scheduled during a probabilistic execution. This information is revealed by means of externally visible transitions of $tester(\mathcal{P})$ with the help of a specific scheduler, called the *observer*, which synchronizes \mathcal{P} with its tester. In this section we present the constructions of the tester and observer and prove some results about the resulting trace distributions.

Informally, the tester of a PA \mathcal{P} is a PA \mathcal{C} whose states include a distinguished start state, all the states of \mathcal{P} , and all the transitions of \mathcal{P} . Automaton \mathcal{C} has a special transition from its own start state, $\bar{q}_{\mathcal{C}}$, to the start state of \mathcal{P} , $\bar{q}_{\mathcal{P}}$, labeled by $\bar{q}_{\mathcal{P}}$. Also, from every state q of \mathcal{P} , \mathcal{C} has a uniform transition labeled by ch (“choose”) to the set of transitions of \mathcal{P} that begin in state q . Finally, for every transition tr of \mathcal{P} and every state q in the support of μ_{tr} , \mathcal{C} has a transition labeled by q from tr to q .

DEFINITION 4.1. *The tester PA of a PA \mathcal{P} , denoted by $tester(\mathcal{P})$, is a PA $\mathcal{C} = (Q_{\mathcal{C}}, \bar{q}_{\mathcal{C}}, E_{\mathcal{C}}, H_{\mathcal{C}}, D_{\mathcal{C}})$, where*

- $Q_{\mathcal{C}} = \{\bar{q}_{\mathcal{C}}\} \cup Q_{\mathcal{P}} \cup D_{\mathcal{P}}$,
- $E_{\mathcal{C}} = Q_{\mathcal{P}} \cup \{ch\}$,
- $H_{\mathcal{C}} = \emptyset$, and
- $D_{\mathcal{C}} = \{(\bar{q}_{\mathcal{C}}, \bar{q}_{\mathcal{P}}, \delta(\bar{q}_{\mathcal{P}}))\} \cup \{(q, ch, \mathcal{U}(\{tr \in D_{\mathcal{P}} \mid source(tr) = q\})) \mid q \in Q_{\mathcal{P}} \wedge q \rightarrow\} \cup \{(tr, q, \delta(q)) \mid tr \in D_{\mathcal{P}}, q \in supp(\mu_{tr})\}$.

Observe that the tester of an ordinary nondeterministic automaton enables at most one transition from each state, and dually, the tester of an automaton that enables at most one transition from each state is a nondeterministic automaton. This observation, together with the results that we prove later in the paper, imply that a fully probabilistic context is enough to observe the branching structure of a nondeterministic automaton.

PROPOSITION 4.2. *The following hold:*

1. *The tester of a nondeterministic automaton is fully probabilistic; that is, it enables at most one transition from each state.*
2. *The tester of a fully probabilistic automaton is a nondeterministic automaton; that is, it contains only transitions whose target measures are Dirac.*

Proof. For the first item, observe that the only states of $tester(\mathcal{P})$ that may enable more than one transition are of the form $tr \in D_{\mathcal{P}}$, which enable one transition for each state in $supp(\mu_{tr})$; however, the size of $supp(\mu_{tr})$ is 1 in a nondeterministic automaton.

For the second item, observe that the only states of $tester(\mathcal{P})$ that may enable non-Dirac transitions are of the form $q \in Q_{\mathcal{P}}$, which may enable a transition labeled by ch to a uniform measure over the set of transitions enabled from q in \mathcal{P} ; however, there is at most one transition enabled from q in a fully probabilistic automaton. \square

We assume without loss of generality that a PA \mathcal{P} and its tester do not have any actions in common (otherwise we can simply rename states of \mathcal{P} to achieve our goal), and thus \mathcal{P} and its tester are compatible.

Since $tester(\mathcal{P})$ and \mathcal{P} share no actions, merely composing $tester(\mathcal{P})$ with \mathcal{P} does not ensure that $tester(\mathcal{P})$ faithfully emulates the behavior of \mathcal{P} . However, an appropriate scheduler can synchronize the two automata and ensure such an emulation,

which will be sufficient for our purposes. Given a PA \mathcal{P} , we define a specific scheduler σ for $\mathcal{P} \parallel \text{tester}(\mathcal{P})$, called the *observer* of \mathcal{P} , that synchronizes the two automata so that the internal structure of \mathcal{P} is visible in the trace. Specifically, the scheduler σ starts by scheduling the transition of $\text{tester}(\mathcal{P})$ from the start state of $\text{tester}(\mathcal{P})$ to the start state of \mathcal{P} , leading to state (\bar{q}, \bar{q}) , which is of the form (q, q) . Then σ repeats the following as long as $q \rightarrow$:

1. Schedule the *ch* transition of $\text{tester}(\mathcal{P})$, thus choosing a transition *tr* of \mathcal{P} .
2. Schedule transition *tr* of \mathcal{P} , leading \mathcal{P} to a new state q' .
3. Schedule the transition of $\text{tester}(\mathcal{P})$ labeled by the state q' , resulting in the state (q', q') , which is again of the form (q, q) .

DEFINITION 4.3. *The observer of a PA \mathcal{P} , which we denote by $\text{observer}(\mathcal{P})$, is a deterministic (almost memoryless) scheduler for $\mathcal{P} \parallel \text{tester}(\mathcal{P})$ that bases its decisions on the last state and sometimes the last action of its argument according to the following table. Here, \mathcal{C} denotes $\text{tester}(\mathcal{P})$, q is any state such that $q \rightarrow$, *tr* is any transition in $D_{\mathcal{P}}$, and q' is any state in μ_{tr} .*

| $Q_{\mathcal{P}}$ | $Q_{\text{tester}(\mathcal{P})}$ | Last action | $D_{\mathcal{P} \parallel \text{tester}(\mathcal{P})}$ |
|-------------------------|--|---------------|---|
| $\bar{q}_{\mathcal{P}}$ | $\bar{q}_{\text{tester}(\mathcal{P})}$ | | $\bar{q}_{\mathcal{P}}$ -labeled transition of $\text{tester}(\mathcal{P})$ |
| q | q | | <i>ch</i> -labeled transition of $\text{tester}(\mathcal{P})$ |
| q | <i>tr</i> | <i>ch</i> | transition <i>tr</i> of \mathcal{P} |
| q' | <i>tr</i> | not <i>ch</i> | q' -labeled transition of $\text{tester}(\mathcal{P})$ |

Scheduler $\text{observer}(\mathcal{P})$ and start state $(\bar{q}_{\mathcal{P}}, \bar{q}_{\text{tester}(\mathcal{P})})$ induce a trace distribution for $\mathcal{P} \parallel \text{tester}(\mathcal{P})$ where all states and external actions of \mathcal{P} appear explicitly.

DEFINITION 4.4. *The observation of a PA \mathcal{P} , denoted by $\text{observation}(\mathcal{P})$, is the trace distribution induced by $\text{observer}(\mathcal{P})$ and $(\bar{q}_{\mathcal{P}}, \bar{q}_{\text{tester}(\mathcal{P})})$.*

Remark 4.5. The *ch*-labeled transitions of a tester are defined to lead to uniform measures over states of $\text{tester}(\mathcal{P})$ that represent transitions of \mathcal{P} . This is well defined since we have assumed that nondeterministic and probabilistic automata are finite branching. From the technical point of view, however, the proofs of this paper rely on the fact that the transitions labeled by *ch* assign nonzero probability (not necessarily the same probability) to each one of the options that are available in a nondeterministic choice. Thus, it would be possible to remove the finite branching restriction from the definition of automata and modify the definition of a tester automaton so that, whenever there are countably many transitions from a state q , the corresponding *ch*-labeled transition of the tester assigns a nonuniform measure to the transitions enabled from q , for example, a Poisson distribution after enumerating all possible transitions enabled from q . We have chosen not to deal with countable branching automata in this paper because it would complicate proofs without adding much insight.

We state and prove some properties of $\text{observation}(\mathcal{P})$. The first property, given in (5), says that the cone of traces beginning with the starting state of \mathcal{P} has probability 1. The second property, (6), says that for any state q of \mathcal{P} from which some transition is enabled and for each finite trace β of $\mathcal{P} \parallel \text{tester}(\mathcal{P})$, the probability of the cone of traces beginning with βq is the same as the probability of the cone beginning with $\beta q \text{ch}$; that is, once βq occurs, the probability that *ch* follows is 1. The third property, (7), says that for any state q of \mathcal{P} and for each finite trace β of $\mathcal{P} \parallel \text{tester}(\mathcal{P})$, the probability of the cone of traces beginning with $\beta q \text{ch}$ is the same as the sum of the probabilities of the cones beginning with $\beta q \text{ch} \beta'$, where β' represents one single step of \mathcal{P} from q ; that is, once *ch* occurs, one of the transitions of \mathcal{P} that are enabled from q is exposed. The right-hand side of (7) consists of two parts dealing with external

and internal transitions, respectively.

PROPOSITION 4.6. *The trace distribution $\eta = \text{observation}(\mathcal{P})$ induced by the observer of a PA \mathcal{P} satisfies the following three properties, for all finite traces β of $\mathcal{P} \parallel \text{tester}(\mathcal{P})$ and for all states q of \mathcal{P} :*

$$\begin{aligned}
 (5) \quad & \eta(C_{\bar{q}}) = 1, \\
 (6) \quad & q \rightarrow \implies \eta(C_{\beta q}) = \eta(C_{\beta q \text{ ch}}), \\
 (7) \quad & \eta(C_{\beta q \text{ ch}}) = \sum_{(a, q') | a \in E, q \xrightarrow{a} q'} \eta(C_{\beta q \text{ ch } a q'}) + \sum_{q' | (\exists a) a \in H, q \xrightarrow{a} q'} \eta(C_{\beta q \text{ ch } q'}).
 \end{aligned}$$

Proof. Equation (5) follows from the fact that *observer*(\mathcal{P}) schedules action \bar{q} immediately. Equation (6) follows from the fact that, after scheduling action q , thus leading to a state of the form (q, q) , *observer*(\mathcal{P}) immediately schedules action *ch* if q enables at least one transition. Equation (7) follows from the fact that, after scheduling *ch*, *observer*(\mathcal{P}) schedules one of the transitions of \mathcal{P} that are enabled from q , say $q \xrightarrow{a} \mu$, followed by a transition of *tester*(\mathcal{P}) labeled by a state in *supp*(μ). \square

The following technical properties will be needed in the proofs of section 6. The first property, in (8), says that the probability of observing a state q' reachable in a tree-structured PA \mathcal{P} with a single transition, say *tr*, from another state q is the probability of observing q , divided by the number of transitions enabled from q , and multiplied by the probability of reaching q' in *tr*, the only transition that may lead to q' since \mathcal{P} is tree-structured. Indeed, q' can be observed only if q is observed (probability of observing q), the transition *tr* is chosen (factor $1/k$ since transitions are chosen uniformly), and the chosen transition leads to q' . The second property, (9), is similar to the first one, where the probability of observing q and scheduling the transition *tr* is replaced by the probability of observing any state in the target of *tr*.

PROPOSITION 4.7. *Let \mathcal{P} be a tree-structured PA, and let η be *observation*(\mathcal{P}). Let $tr = (q, a, \mu)$ be a transition of \mathcal{P} . Let k be the number of transitions that are enabled from q in \mathcal{P} , and let q' be a state in *supp*(μ). Then the following properties hold:*

$$\begin{aligned}
 (8) \quad & \eta(\diamond q') = \frac{\eta(\diamond q)}{k} \mu(q'), \\
 (9) \quad & \eta(\diamond q') = \left(\sum_{q'' \in \text{supp}(\mu)} \eta(\diamond q'') \right) \mu(q').
 \end{aligned}$$

Proof. Let σ be the observer of \mathcal{P} , and let ϵ_σ be the probabilistic execution induced by σ . Since \mathcal{P} is tree-structured, the set Θ_q contains a single execution α . Indeed, by the definition of tree-structured, there is only one execution in \mathcal{P} ending with state q , and σ simply interleaves this execution with transitions labeled by *ch*, by the names of the transitions of \mathcal{P} that are needed to reach q , and by the names of the states that are reached. Similarly, $\Theta_{q'}$ contains a single execution α' .

Once state q is reached, σ schedules action *ch*, reaching state *tr* of *tester*(\mathcal{P}) with probability $1/k$. Then, σ schedules transition *tr*, reaching state q' in \mathcal{P} with probability $\mu(q')$, and finally σ schedules the transition of *tester*(\mathcal{P}) labeled by q' . Thus, $\epsilon_\sigma(C_{\alpha'}) = \epsilon_\sigma(C_\alpha)(1/k)\mu(q')$. Then (8) follows by (2).

By summing over $\text{supp}(\mu)$ in (8), we get

$$(10) \quad \sum_{q'' \in \text{supp}(\mu)} \eta(\diamond q'') = \frac{\eta(\diamond q)}{k} \sum_{q'' \in \text{supp}(\mu)} \mu(q'').$$

Observe that $\sum_{q'' \in \text{supp}(\mu)} \mu(q'') = 1$. Hence, (10) simplifies to

$$(11) \quad \sum_{q'' \in \text{supp}(\mu)} \eta(\diamond q'') = \frac{\eta(\diamond q)}{k}.$$

Substitution of (11) into (8) gives us (9), as needed. \square

5. Characterizations of \leq_{DC} for nondeterministic automata. In this section, we present our characterization theorems for \leq_{DC} for nondeterministic automata: Theorem 5.2 characterizes \leq_{DC} in terms of \leq_F , for nondeterministic automata without internal actions, and Theorem 5.4 characterizes \leq_{DC} in terms of \leq_{wF} , for arbitrary nondeterministic automata. In each case, we prove the result first for tree-structured nondeterministic automata and then extend it to the non-tree-structured case via unfolding. The interesting direction for each of these results is the completeness direction, showing that $\mathcal{A}_1 \leq_{DC} \mathcal{A}_2$ implies the existence of a simulation relation from \mathcal{A}_1 to \mathcal{A}_2 .

The strategy that we use to prove our completeness results is also applied in many other full abstraction results; see, for example, [5, 14]. By Proposition 3.8, $\mathcal{A}_1 \leq_{DC} \mathcal{A}_2$ implies that $\mathcal{A}_1 \leq_D \mathcal{A}_2$ for all contexts \mathcal{C} . Thus it suffices to construct a specific context \mathcal{C} with the property that the trace distributions of $\mathcal{A}_1 \parallel \mathcal{C}$ contain all information about \mathcal{A}_1 that is preserved by the simulation preorder. More specifically, we compose \mathcal{A}_1 with the context $\mathcal{C} = \text{tester}(\mathcal{A}_1)$ and consider just a single trace distribution of the composed system, namely $\text{observation}(\mathcal{A}_1)$, the one generated by $\text{observer}(\mathcal{A}_1)$. We show, for any other nondeterministic automaton \mathcal{A}_2 , that if the composition $\mathcal{A}_2 \parallel \text{tester}(\mathcal{A}_1)$ generates the trace distribution $\text{observation}(\mathcal{A}_1)$, then \mathcal{A}_2 actually simulates \mathcal{A}_1 in a strong sense. Namely, whenever \mathcal{A}_1 reaches some state q_1 , \mathcal{A}_2 can reach a corresponding state q_2 from which it generates the same trace distribution. The formalities of the proof are intricate, in part because states of \mathcal{A}_1 also show up as states of $\text{tester}(\mathcal{A}_1)$ and within the trace distribution of $\text{observation}(\mathcal{A}_1)$. In the proof we try to be very explicit about the roles of states of \mathcal{A}_1 , but we also warn the reader to be alert to this potential source of confusion.

5.1. Nondeterministic automata without internal actions. We begin by considering nondeterministic automata without internal actions. We first consider tree-structured nondeterministic automata.

PROPOSITION 5.1. *Let $\mathcal{A}_1, \mathcal{A}_2$ be nondeterministic automata without internal actions such that \mathcal{A}_1 is tree-structured. Then $\mathcal{A}_1 \leq_{DC} \mathcal{A}_2$ implies $\mathcal{A}_1 \leq_F \mathcal{A}_2$.*

Proof. Assume that $\mathcal{A}_1 \leq_{DC} \mathcal{A}_2$. Let \mathcal{C} be $\text{tester}(\mathcal{A}_1)$ and η be $\text{observation}(\mathcal{A}_1)$, that is, the trace distribution of $\mathcal{A}_1 \parallel \mathcal{C}$ induced by the scheduler $\text{observer}(\mathcal{A}_1)$. Since $\mathcal{A}_1 \leq_{DC} \mathcal{A}_2$ implies $\mathcal{A}_1 \parallel \mathcal{C} \leq_D \mathcal{A}_2 \parallel \mathcal{C}$, Proposition 3.8 implies that η is also a trace distribution of $\mathcal{A}_2 \parallel \mathcal{C}$. That is, there exists a probabilistic execution ϵ of $\mathcal{A}_2 \parallel \mathcal{C}$, induced by some scheduler σ_2 , such that $\text{tdist}(\epsilon) = \eta$.

For each state q_1 in Q_1 , let Θ_{q_1} be the set of finite executions of $\mathcal{A}_2 \parallel \mathcal{C}$ whose last transition is labeled by q_1 . For each state q_2 of \mathcal{A}_2 , let Θ_{q_1, q_2} be the set of executions in Θ_{q_1} whose last state is the pair (q_2, q_1) .

Define a relation R on $Q_1 \times Q_2$ as follows: $q_1 R q_2$ iff there exists a finite execution α in Θ_{q_1, q_2} such that $\epsilon(C_\alpha) > 0$. We claim that R is a forward simulation from \mathcal{A}_1 to \mathcal{A}_2 .

For the start condition, we must show that $\bar{q}_1 R \bar{q}_2$. Consider the start state (\bar{q}_2, \bar{q}_C) of $\mathcal{A}_2 \parallel \mathcal{C}$. Since there are no internal actions in \mathcal{A}_2 or \mathcal{C} , and since, by (5) from Proposition 4.6, $\eta(C_{\bar{q}_1}) = 1$, the only action that is scheduled initially by σ_2 is \bar{q}_1 , leading to state (\bar{q}_2, \bar{q}_1) . Thus, the finite execution $\alpha = (\bar{q}_2, \bar{q}_C)\bar{q}_1(\bar{q}_2, \bar{q}_1)$ is an element of $\Theta_{\bar{q}_1, \bar{q}_2}$ such that $\epsilon(C_\alpha) > 0$, as needed.

For the step condition, assume $q_1 R q_2$ and let $q_1 \xrightarrow{a}_1 q'_1$ be a transition of \mathcal{A}_1 , which we denote by tr for convenience. We exhibit a matching transition $q_2 \xrightarrow{a}_2 q'_2$.

By the definition of R , there exists a finite execution α in Θ_{q_1, q_2} such that $\epsilon(C_\alpha) > 0$. Since Θ_{q_1, q_2} is a subset of Θ_{q_1} , by definition of Θ_{q_1} , $trace(\alpha) = \beta q_1$ for some finite trace β . Therefore, $\eta(C_{\beta q_1}) > 0$. Since q_1 enables at least one transition in \mathcal{A}_1 , specifically transition tr , (6) from Proposition 4.6 implies that $\eta(C_{\beta q_1 ch}) = \eta(C_{\beta q_1})$. Then, since \mathcal{A}_2 and \mathcal{C} have no internal actions, σ_2 schedules action ch from α with probability 1.

By the definition of $tester(\mathcal{A}_1)$, the transition labeled by ch leaving from state q_1 of \mathcal{C} leads to state tr with probability > 0 . Hence, $\epsilon(C_{\alpha ch(q_2, tr)}) > 0$. By (7) from Proposition 4.6, where only the first term of the right-hand side is nonzero due to the absence of internal actions, $\eta(C_{\beta q_1 ch}) = \sum_{(a, q') | a \in E, q_1 \xrightarrow{a} q'} \eta(C_{\beta q_1 ch a q'})$. Hence, σ_2 must extend $\alpha ch(q_2, tr)$ with two steps labeled by an action and a state of \mathcal{A}_1 , respectively, where the action and the state are compatible with one of the transitions of \mathcal{A}_1 that are enabled from q_1 . Since state tr of \mathcal{C} enables only action q'_1 , and since, by the tree-structure of \mathcal{A}_1 , a is uniquely determined by q'_1 , the action and state scheduled by σ_2 are a and q'_1 . Therefore, there exists a state q'_2 of \mathcal{A}_2 such that the execution $\alpha' = \alpha ch(q_2, tr)a(q'_2, tr)q'_1(q'_2, q'_1)$ is an execution in $\Theta_{q'_1, q'_2}$ such that $\epsilon(C_{\alpha'}) > 0$. Then $q'_1 R q'_2$ and $q_2 \xrightarrow{a}_2 q'_2$, as needed. \square

Now we present our result for general (non-tree-structured) nondeterministic automata without internal actions.

THEOREM 5.2. *Let $\mathcal{A}_1, \mathcal{A}_2$ be nondeterministic automata without internal actions. Then $\mathcal{A}_1 \leq_{DC} \mathcal{A}_2$ iff $\mathcal{A}_1 \leq_F \mathcal{A}_2$.*

Proof. First we prove soundness of forward simulations:

$$\begin{aligned} \mathcal{A}_1 \leq_F \mathcal{A}_2 &\Rightarrow \mathcal{A}_1 \leq_{PF} \mathcal{A}_2 && \text{(Proposition 3.9.1)} \\ &\Rightarrow \mathcal{A}_1 \leq_{wPF} \mathcal{A}_2 && \text{(Proposition 3.10.1)} \\ &\Rightarrow \mathcal{A}_1 \leq_{DC} \mathcal{A}_2 && \text{(Proposition 3.10.3).} \end{aligned}$$

Next we establish completeness:

$$\begin{aligned} \mathcal{A}_1 \leq_{DC} \mathcal{A}_2 &\Rightarrow Unfold(\mathcal{A}_1) \leq_F \mathcal{A}_1 \leq_{DC} \mathcal{A}_2 && \text{(Proposition 2.4)} \\ &\Rightarrow Unfold(\mathcal{A}_1) \leq_{DC} \mathcal{A}_1 \leq_{DC} \mathcal{A}_2 && \text{(as in soundness proof)} \\ &\Rightarrow Unfold(\mathcal{A}_1) \leq_{DC} \mathcal{A}_2 && (\leq_{DC} \text{ is transitive)} \\ &\Rightarrow Unfold(\mathcal{A}_1) \leq_F \mathcal{A}_2 && \text{(Proposition 5.1)} \\ &\Rightarrow \mathcal{A}_1 \leq_F Unfold(\mathcal{A}_1) \leq_F \mathcal{A}_2 && \text{(Proposition 2.4)} \\ &\Rightarrow \mathcal{A}_1 \leq_F \mathcal{A}_2 && (\leq_F \text{ is transitive}). \quad \square \end{aligned}$$

5.2. Nondeterministic automata with internal actions. Next we extend the results of section 5.1 to nondeterministic automata that may include internal actions. The proofs are analogous to those in section 5.1. The difference is that, in several places in the proof of Proposition 5.3, we need to reason about multistep

extensions of executions instead of single-step extensions. Again, we begin with tree-structured nondeterministic automata.

PROPOSITION 5.3. *Let $\mathcal{A}_1, \mathcal{A}_2$ be nondeterministic automata such that \mathcal{A}_1 is tree-structured. Then $\mathcal{A}_1 \leq_{DC} \mathcal{A}_2$ implies $\mathcal{A}_1 \leq_{wF} \mathcal{A}_2$.*

Proof. Assume that $\mathcal{A}_1 \leq_{DC} \mathcal{A}_2$. Let \mathcal{C} be *tester*(\mathcal{A}_1) and η be *observation*(\mathcal{A}_1), that is, the trace distribution of $\mathcal{A}_1 \parallel \mathcal{C}$ induced by the scheduler *observer*(\mathcal{A}_1). Define the scheduler σ_2 , the probabilistic execution ϵ , and the Θ sets as in the proof of Proposition 5.1.

The definition of R is slightly different: $q_1 R q_2$ iff there exists a state q'_2 such that $q_2 \Rightarrow q'_2$ and there exists $\alpha \in \Theta_{q_1, q'_2}$ such that $\epsilon(C_\alpha) > 0$. We claim that R is a weak forward simulation from \mathcal{A}_1 to \mathcal{A}_2 .

For the start condition, we must show that $\bar{q}_1 R \bar{q}_2$. By Item 1 of Proposition 4.6, $\eta(C_{\bar{q}_1}) = 1$. This means that there exists a finite execution fragment α of $\mathcal{A}_2 \parallel \mathcal{C}$ with trace \bar{q}_1 that ends with action \bar{q}_1 such that $\epsilon(C_\alpha) > 0$. By definition of \mathcal{C} , the last state of α is (q_2, \bar{q}_1) for some state q_2 satisfying $\bar{q}_2 \Rightarrow q_2$. By definition of R , $\bar{q}_1 R \bar{q}_2$, as needed.

For the step condition, assume $q_1 R q_2$ and let $q_1 \xrightarrow{a} q'_1$ be a transition of \mathcal{A}_1 , which we denote by tr . We exhibit a matching weak transition $q_2 \xrightarrow{a} q'_2$.

By definition of R , there exists a state q''_2 of \mathcal{A}_2 such that $q_2 \Rightarrow q''_2$, and there exists a finite execution fragment α in Θ_{q_1, q''_2} such that $\epsilon(C_\alpha) > 0$. Since Θ_{q_1, q''_2} is a subset of Θ_{q_1} , by definition of Θ_{q_1} , $trace(\alpha) = \beta q_1$ for some finite trace β . Therefore, $\eta(C_{\beta q_1}) > 0$. Since q_1 enables at least one transition in \mathcal{A}_1 , specifically transition tr , (6) from Proposition 4.6 implies that $\eta(C_{\beta q_1 ch}) = \eta(C_{\beta q_1})$. Thus, there exists an execution fragment α' of $\mathcal{A}_2 \parallel \mathcal{C}$ with trace ch such that $\epsilon(C_{\alpha \wedge \alpha'}) > 0$. Furthermore, since, by definition of $\mathcal{C} = \text{tester}(\mathcal{A}_1)$, the transition of \mathcal{C} labeled by ch that leaves from state q_1 leads to state tr with nonzero probability, we can assume that the last state of α' is of the form (q', tr) for some state q' of \mathcal{A}_2 .

Recall from above that $\eta(C_{\beta q_1 ch}) > 0$. By (7) from Proposition 4.6, $\eta(C_{\beta q_1 ch}) = \sum_{(a, q') | a \in E, q \xrightarrow{a} q'} \eta(C_{\beta q_1 ch a q'}) + \sum_{q' | (\exists a) a \in H, q_1 \xrightarrow{a} q'} \eta(C_{\beta q_1 ch q'})$. Hence, σ_2 must extend $\alpha \wedge \alpha'$ in such a way that the first or the first two external actions are compatible with one of the transitions of \mathcal{A}_1 that are enabled from q_1 . (The number of external actions depends on whether the compatible transition of \mathcal{A}_1 is labeled by an internal or external action.) Since state tr of \mathcal{C} enables only action q'_1 , and since, by the tree-structure of \mathcal{A}_1 , a is uniquely determined by q'_1 , the first or first two external actions of $\mathcal{A}_2 \parallel \mathcal{C}$ scheduled by σ_2 are either q'_1 or $a q'_1$, depending on whether a is internal or external. Thus, there exists an execution fragment α'' of $\mathcal{A}_2 \parallel \mathcal{C}$, with trace $trace(a q'_1)$, such that $\epsilon(C_{\alpha \wedge \alpha' \wedge \alpha''}) > 0$. Furthermore, we can assume that the last transition of α'' is labeled by q'_1 (simply truncate α'' otherwise).

Let (q'_2, q'_1) be the last state of α'' . Then, $\alpha \wedge \alpha' \wedge \alpha'' \in \Theta_{q'_1, q'_2}$, thus showing that $q'_1 R q'_2$. It remains to show that $q_2 \xrightarrow{a} q'_2$. For this, it suffices to recall that $q_2 \Rightarrow q''_2$ and observe that $q''_2 \xrightarrow{a} q'_2$ since the execution fragment $(\alpha' \wedge \alpha'') \upharpoonright \mathcal{A}_2$ has trace $trace(a)$, first state q''_2 , and last state q'_2 . \square

Using the same approach as before, we may eliminate the assumption in Proposition 5.3 that \mathcal{A}_1 is tree-structured.

THEOREM 5.4. *Let $\mathcal{A}_1, \mathcal{A}_2$ be nondeterministic automata. Then $\mathcal{A}_1 \leq_{DC} \mathcal{A}_2$ iff $\mathcal{A}_1 \leq_{wF} \mathcal{A}_2$.*

Proof. This proof is analogous to that of Theorem 5.2. First we prove soundness

of weak forward simulations:

$$\begin{aligned} \mathcal{A}_1 \leq_{wF} \mathcal{A}_2 &\Rightarrow \mathcal{A}_1 \leq_{wPF} \mathcal{A}_2 \quad (\text{Proposition 3.9.2}) \\ &\Rightarrow \mathcal{A}_1 \leq_{DC} \mathcal{A}_2 \quad (\text{Proposition 3.10.3}). \end{aligned}$$

Now we prove completeness:

$$\begin{aligned} \mathcal{A}_1 \leq_{DC} \mathcal{A}_2 &\Rightarrow \text{Unfold}(\mathcal{A}_1) \leq_F \mathcal{A}_1 \leq_{DC} \mathcal{A}_2 \quad (\text{Proposition 2.4}) \\ &\Rightarrow \text{Unfold}(\mathcal{A}_1) \leq_{DC} \mathcal{A}_1 \leq_{DC} \mathcal{A}_2 \\ &\quad (\text{as in proof Theorem 5.2, soundness part}) \\ &\Rightarrow \text{Unfold}(\mathcal{A}_1) \leq_{DC} \mathcal{A}_2 \quad (\leq_{DC} \text{ is transitive}) \\ &\Rightarrow \text{Unfold}(\mathcal{A}_1) \leq_{wF} \mathcal{A}_2 \quad (\text{Proposition 5.3}) \\ &\Rightarrow \mathcal{A}_1 \leq_F \text{Unfold}(\mathcal{A}_1) \leq_{wF} \mathcal{A}_2 \quad (\text{Proposition 2.4}) \\ &\Rightarrow \mathcal{A}_1 \leq_{wF} \text{Unfold}(\mathcal{A}_1) \leq_{wF} \mathcal{A}_2 \quad (\text{Proposition 2.3.1}) \\ &\Rightarrow \mathcal{A}_1 \leq_{wF} \mathcal{A}_2 \quad (\leq_{wF} \text{ is transitive}). \quad \square \end{aligned}$$

6. Characterizations of \leq_{DC} for PAs. Now we present our characterization theorems for \leq_{DC} for PAs: Theorem 6.3 characterizes \leq_{DC} in terms of \leq_{PF} , for PAs without internal actions, and Theorem 6.5 characterizes \leq_{DC} in terms of \leq_{wPF} , for arbitrary PAs. Again, we give the results first for tree-structured PAs and extend them by unfolding. As before, the interesting direction is the completeness direction, showing that $\mathcal{P}_1 \leq_{DC} \mathcal{P}_2$ implies the existence of a simulation relation from \mathcal{P}_1 to \mathcal{P}_2 . Our proofs of completeness for PAs are analogous to those for nondeterministic automata.

6.1. PAs without internal actions. We first consider tree-structured PAs.

PROPOSITION 6.1. *Let $\mathcal{P}_1, \mathcal{P}_2$ be PAs without internal actions such that \mathcal{P}_1 is tree-structured. Then $\mathcal{P}_1 \leq_{DC} \mathcal{P}_2$ implies $\mathcal{P}_1 \leq_{PF} \mathcal{P}_2$.*

Proof. Assume that $\mathcal{P}_1 \leq_{DC} \mathcal{P}_2$. Let \mathcal{C} be $\text{tester}(\mathcal{P}_1)$ and η be $\text{observation}(\mathcal{P}_1)$, that is, the trace distribution of $\mathcal{P}_1 \parallel \mathcal{C}$ induced by the scheduler $\text{observer}(\mathcal{P}_1)$. Define the scheduler σ_2 , the probabilistic execution ϵ , and the Θ sets as in the proof of Proposition 5.1.

Define a relation R as follows: $q_1 R \mu_2$ iff $\sum_{\alpha \in \Theta_{q_1}} \epsilon(C_\alpha) > 0$ and for each state $q_2 \in Q_2$,

$$(12) \quad \mu_2(q_2) = \frac{\sum_{\alpha \in \Theta_{q_1, q_2}} \epsilon(C_\alpha)}{\sum_{\alpha \in \Theta_{q_1}} \epsilon(C_\alpha)}.$$

That is, the measure μ_2 describes probabilities of the various Θ_{q_1, q_2} 's relative to Θ_{q_1} . Note that the equation above is well defined since, by the tree-structure of \mathcal{P}_1 , all the cones represented by Θ_{q_1} are disjoint, and thus $\sum_{\alpha \in \Theta_{q_1}} \epsilon(C_\alpha) \leq 1$. We claim that R is a probabilistic forward simulation from \mathcal{P}_1 to \mathcal{P}_2 .

Before proving that R is a probabilistic forward simulation we make several observations:

1. Relation R is a function from Q_1 to $\text{Disc}(Q_2)$. Indeed, if $\sum_{\alpha \in \Theta_{q_1}} \epsilon(C_\alpha) > 0$, then there exists exactly one measure that satisfies (12). Furthermore, given the construction of η and the fact that \mathcal{P}_1 is tree-structured (i.e., all states are reachable), every state q_1 of Q_1 occurs with some positive probability in η . Thus, since η is induced by ϵ , $\sum_{\alpha \in \Theta_{q_1}} \epsilon(C_\alpha) > 0$ for all states q_1 of Q_1 .
2. If $q_1 R \mu_2$, then, for each state $q_2 \in Q_2$ and each execution $\alpha \in \Theta_{q_1, q_2}$,

$$(13) \quad \epsilon(C_\alpha) > 0 \Rightarrow q_2 \in \text{supp}(\mu_2).$$

That is, the execution α occurs with nonzero probability in ϵ only if μ_2 assigns nonzero probability to q_2 . This property is a direct consequence of (12).

3. For each transition $q_1 \xrightarrow{a} \mu'_1$ of \mathcal{P}_1 , the following equation holds:

$$(14) \quad \mu'_1(q'_1) = \frac{\sum_{\alpha \in \Theta_{q'_1}} \epsilon(C_\alpha)}{\sum_{q \in \text{supp}(\mu'_1), \alpha \in \Theta_q} \epsilon(C_\alpha)}$$

That is, the relative probabilities of the states of $\text{supp}(\mu'_1)$ in ϵ are given by μ'_1 . This result follows by instantiating (9) from Proposition 4.7 with $q_1 \xrightarrow{a} \mu'_1$ to derive the probability of a state q'_1 in the support of μ'_1 , and by replacing the diamond expressions according to (2) from Proposition 3.3.

4. For each transition $tr = q_1 \xrightarrow{a} \mu'_1$ of \mathcal{P}_1 , the following equation holds:

$$(15) \quad \sum_{\alpha \in \Theta_{q_1}} \epsilon(C_\alpha) = k \sum_{q \in \text{supp}(\mu'_1), \alpha \in \Theta_q} \epsilon(C_\alpha),$$

where k is the number of transitions of \mathcal{P}_1 enabled from q_1 . That is, the probability of reaching q_1 in ϵ is k times the probability of reaching q_1 and scheduling tr . Informally, transition tr is scheduled only if state q_1 is reached and the outcome of the following transition labeled by ch is tr , which happens with probability $1/k$. The reason why $\sum_{q \in \text{supp}(\mu'_1), \alpha \in \Theta_q} \epsilon(C_\alpha)$ is the probability of reaching q_1 and scheduling tr is that states from $\text{supp}(\mu'_1)$ can occur only after q_1 has occurred and tr is reached (see the definition of tester automaton and of observer of a tester automaton), and furthermore states from $\text{supp}(\mu'_1)$ occur with probability 1 once tr is reached (see (7) from Proposition 4.6).

This follows by instantiating (8) from Proposition 4.7 with tr , replacing the diamond expressions according to (2) from Proposition 3.3, summing over $\text{supp}(\mu'_1)$, observing that $\sum_{q'_1 \in \text{supp}(\mu'_1)} \mu'_1(q'_1) = 1$, and deriving $\sum_{\alpha \in \Theta_{q_1}} \epsilon(C_\alpha)$ from the resulting equation.

We are now ready to show that R is a probabilistic forward simulation. For the start condition, we must show that $\bar{q}_1 R \delta(\bar{q}_2)$.

Consider the start state (\bar{q}_2, \bar{q}_c) of $\mathcal{P}_2 \parallel \mathcal{C}$. Since there are no internal actions in \mathcal{P}_2 or \mathcal{C} , and since, by (5) from Proposition 4.6, $\eta(C_{\bar{q}_1}) = 1$, the only action that is scheduled initially by σ_2 is \bar{q}_1 , leading to state (\bar{q}_2, \bar{q}_1) with probability 1. Thus, the finite execution $\alpha = (\bar{q}_2, \bar{q}_c)\bar{q}_1(\bar{q}_2, \bar{q}_1)$ is an element of $\Theta_{\bar{q}_1, \bar{q}_2}$ such that $\epsilon(C_\alpha) = 1$, and, by definition of R , $\bar{q}_1 R \delta(\bar{q}_2)$, as needed.

For the step condition, assume that $q_1 R \mu_2$ and let $q_1 \xrightarrow{a} \mu'_1$ be a transition of \mathcal{P}_1 , which we denote by tr . We must exhibit a probability measure $\xi'_2 \in \text{Disc}(\text{Disc}(Q_2))$ and a hyper-transition $\mu_2 \xrightarrow{a} \mu''_2$, matching the given transition, where $\mu''_2 = \text{flatten}(\xi'_2)$ and $\mu'_1 R \xi'_2$. We do this by deriving a transition tr_α of \mathcal{P}_2 for each execution α of Θ_{q_1} and by combining the tr_α 's appropriately into transitions tr_q , for each state $q \in \text{supp}(\mu_2)$, that are the basis for the required hyper-transition. The tr_α transitions are derived from η ; the construction considers only those α 's for which $\epsilon(C_\alpha) > 0$. The other α 's can be treated arbitrarily.

Consider an execution α of Θ_{q_1} such that $\epsilon(C_\alpha) > 0$. By property (13), $\alpha \in \Theta_{q_1, q_2}$ for some state q_2 in $\text{supp}(\mu_2)$. Since Θ_{q_1, q_2} is a subset of Θ_{q_1} , by definition of Θ_{q_1} , $\text{trace}(\alpha) = \beta q_1$ for some finite trace β . Therefore, $\eta(C_{\beta q_1}) > 0$. Since q_1 enables at least one transition in \mathcal{P}_1 , specifically transition tr , (6) from Proposition 4.6 implies

that $\eta(C_{\beta q_1 ch}) = \eta(C_{\beta q_1})$. Then, since \mathcal{P}_2 and \mathcal{C} have no internal actions, σ_2 schedules action ch from α with probability 1.

By definition of $\mathcal{C} = \text{tester}(\mathcal{P}_1)$, the transition labeled by ch leaving from state q_1 of \mathcal{C} leads to state tr with probability > 0 . Hence, $\epsilon(C_{\alpha ch(q_2, tr)}) > 0$. By (7) from Proposition 4.6, where only the first term of the right-hand side is nonzero due to the absence of internal actions, $\eta(C_{\beta q_1 ch}) = \sum_{(a, q') | a \in E, q_1 \xrightarrow{a} q'} \eta(C_{\beta q_1 ch a q'})$. Hence, σ_2 must extend $\alpha ch(q_2, tr)$ with two steps labeled by an action and a state of \mathcal{P}_1 , respectively, where the action and the state are compatible with one of the transitions of \mathcal{P}_1 that are enabled from q_1 . Since state tr of \mathcal{C} enables only actions in $\text{supp}(\mu'_1)$, and since, by the tree-structure of \mathcal{P}_1 , a is uniquely determined by μ'_1 , the action that is scheduled is a , and the state that is scheduled is a state in $\text{supp}(\mu'_1)$. Thus, $\sigma_2(\alpha ch(q_2, tr))$ returns a probability measure over transitions labeled by a . This measure identifies a combined transition of \mathcal{P}_2 labeled by a that leaves from q_2 , which we denote by tr_α .

Now, using the tr_α transitions, we define a combined transition from each state in the support of μ_2 . Namely, for each state $q \in \text{supp}(\mu_2)$, let tr_q be the combined transition of \mathcal{P}_2 defined by

$$(16) \quad tr_q \triangleq \sum_{\alpha \in \Theta_{q_1, q}} \frac{\epsilon(C_\alpha)}{\sum_{\alpha' \in \Theta_{q_1, q}} \epsilon(C_{\alpha'})} tr_\alpha.$$

Informally, each element of $\Theta_{q_1, q}$ is an execution that contributes to the emulation of transition $q_1 \xrightarrow{a} \mu'_1$ from q . Equation (16) computes tr_q , the overall contribution to the emulation from q , by averaging over all elements of $\Theta_{q_1, q}$. We could prove that $\Theta_{q_1, q}$ contains only one element α' such that $\epsilon(C_{\alpha'}) > 0$ and simplify (16) accordingly. However, this simplification is not necessary for the proof. Now we define the measure $\mu''_2 \in \text{Disc}(Q_2)$:

$$(17) \quad \mu''_2 \triangleq \sum_{q \in \text{supp}(\mu_2)} \mu_2(q) \mu_{tr_q}.$$

Then, by construction, $\mu_2 \xrightarrow{a} \mu''_2$ is a hyper-transition of \mathcal{P}_2 .

It remains to define a probability measure $\xi'_2 \in \text{Disc}(\text{Disc}(Q_2))$ such that $\mu''_2 = \text{flatten}(\xi'_2)$ and $\mu'_1 R \xi'_2$.

For each $q \in \text{supp}(\mu'_1)$, let μ_q be the unique measure such that $q R \mu_q$. We can identify μ_q because R is a function. Define $\xi'_2 \in \text{Disc}(\text{Disc}(Q_2))$ such that, for each $q \in \text{supp}(\mu'_1)$, $\xi'_2(\mu_q) = \sum_{q' \in \text{supp}(\mu'_1) | \mu_{q'} = \mu_q} \mu'_1(q')$. Then $\mu'_1 R \xi'_2$ by the definition of ξ'_2 .

It remains to show that $\mu''_2 = \text{flatten}(\xi'_2)$, that is, that $\mu''_2 = \sum_{\rho \in \text{supp}(\xi'_2)} \xi'_2(\rho) \rho$. From the definition of ξ'_2 and of the flatten operator, it suffices to show that for every $q_2 \in Q_2$,

$$(18) \quad \mu''_2(q_2) = \sum_{q \in \text{supp}(\mu'_1)} \mu'_1(q) \mu_q(q_2).$$

To prove (18) we first claim that the following equation is valid for each pair of states q_1, q_2 of \mathcal{P}_1 and \mathcal{P}_2 , respectively, if k denotes the number of transitions of \mathcal{P}_1 that are enabled from q_1 :

$$(19) \quad \sum_{\alpha \in \Theta_{q_1}} \epsilon(C_\alpha) \mu_{tr_\alpha}(q_2) = k \sum_{q \in \text{supp}(\mu'_1), \alpha \in \Theta_{q, q_2}} \epsilon(C_\alpha).$$

Informally, the left-hand side of (19) represents the probability of scheduling q_1 and then reaching q_2 according to the transition tr_α , without considering the outcome of the transition labeled by ch . The right-hand side, on the other hand, computes the probability of scheduling q_1 , scheduling ch and reaching μ'_1 , and then scheduling tr_α and reaching q_2 . State μ'_1 is reached by ch with probability $1/k$, which justifies the k factor in the right-hand side.

To prove (19), consider an execution $\alpha \in \Theta_{q,q_2}$, where $q \in \text{supp}(\mu'_1)$. Since q always occurs after q_1 , execution α can be split into $\alpha' \frown \alpha''$, where $\alpha' \in \Theta_{q_1}$. Furthermore, $\text{trace}(\alpha'') = ch\ aq$, and since there are no internal actions in \mathcal{P}_2 and \mathcal{C} , α is the unique extension of α' that is in Θ_{q,q_2} . In particular,

$$\alpha'' = (q', q_1) \ ch \ (q', tr)a(q_2, tr)q(q_2, q)$$

for some state q' of \mathcal{P}_2 , and $\epsilon(C_\alpha) = \epsilon(C_{\alpha'})(1/k)\mu_{tr_{\alpha'}}(q_2)$. Thus, each summand on the right-hand side of (19) has a corresponding summand on the left-hand side that differs by a factor of k , and the correspondence relation is an injection. If the correspondence is not a bijection, then the α terms that are left out on the left-hand side are such that $\mu_{tr_\alpha}(q_2) = 0$ (otherwise an extension in Θ_{q,q_2} for some q exists). This suffices.

We now consider the left-hand side of (18). Consider the definition of μ''_2 given by (17). By expanding $\mu_2(q)$ according to the definition of μ_2 given by (12), and expanding $\mu_{tr}(q_2)$ according to the definition of μ_{tr} given by (16), we obtain

$$\mu''_2(q_2) = \sum_{q \in \text{supp}(\mu_2)} \frac{\sum_{\alpha \in \Theta_{q_1,q}} \epsilon(C_\alpha)}{\sum_{\alpha \in \Theta_{q_1}} \epsilon(C_\alpha)} \frac{\sum_{\alpha \in \Theta_{q_1,q}} \epsilon(C_\alpha) \mu_{tr_\alpha}(q_2)}{\sum_{\alpha \in \Theta_{q_1,q}} \epsilon(C_\alpha)}.$$

By cross-simplifying the top-leftmost and bottom-rightmost factors, and by factoring the left denominator out of the sum, we obtain

$$\mu''_2(q_2) = \frac{\sum_{q \in \text{supp}(\mu_2)} \sum_{\alpha \in \Theta_{q_1,q}} \epsilon(C_\alpha) \mu_{tr_\alpha}(q_2)}{\sum_{\alpha \in \Theta_{q_1}} \epsilon(C_\alpha)}.$$

By property (13), we can rewrite the numerator as follows:

$$\mu''_2(q_2) = \frac{\sum_{\alpha \in \Theta_{q_1}} \epsilon(C_\alpha) \mu_{tr_\alpha}(q_2)}{\sum_{\alpha \in \Theta_{q_1}} \epsilon(C_\alpha)}.$$

By multiplying numerator and denominator by k , applying (19) to the numerator, and applying (15) to the denominator, we obtain

$$(20) \quad \mu''_2(q_2) = \frac{\sum_{q \in \text{supp}(\mu'_1), \alpha \in \Theta_{q,q_2}} \epsilon(C_\alpha)}{\sum_{q \in \text{supp}(\mu'_1), \alpha \in \Theta_q} \epsilon(C_\alpha)}.$$

We now consider the right-hand side of (18). By applying (14) and (12) to the two factors of the right-hand side of (18), and by simplifying common factors algebraically, we obtain

$$(21) \quad \sum_{q \in \text{supp}(\mu'_1)} \mu'_1(q) \mu_q(q_2) = \frac{\sum_{q \in \text{supp}(\mu'_1), \alpha \in \Theta_{q,q_2}} \epsilon(C_\alpha)}{\sum_{q \in \text{supp}(\mu'_1), \alpha \in \Theta_q} \epsilon(C_\alpha)}.$$

Now (18) follows by direct combination of (20) and (21). \square

Interestingly, the probabilistic forward simulation that we constructed in the above proof is functional. Functional simulations are usually called *refinement mappings* [21, 26]. Write $\mathcal{P}_1 \leq_{PR} \mathcal{P}_2$ if there exists a functional probabilistic forward simulation from \mathcal{P}_1 to \mathcal{P}_2 . Then we can state the following new proposition, which is a probabilistic version of Proposition 3.12 in [26].

PROPOSITION 6.2. *Let $\mathcal{P}_1, \mathcal{P}_2$ be PAs without internal actions such that \mathcal{P}_1 is tree-structured. Then $\mathcal{P}_1 \leq_{PF} \mathcal{P}_2$ iff $\mathcal{P}_1 \leq_{PR} \mathcal{P}_2$.*

Proof. It is enough to observe that each state q_1 of \mathcal{P}_1 occurs with some positive probability in the trace distribution η of the proof of Proposition 6.1. \square

As usual, we may eliminate the assumption that \mathcal{P} is tree-structured.

THEOREM 6.3. *Let $\mathcal{P}_1, \mathcal{P}_2$ be PAs without internal actions. Then $\mathcal{P}_1 \leq_{DC} \mathcal{P}_2$ iff $\mathcal{P}_1 \leq_{PF} \mathcal{P}_2$.*

Proof. First we prove the soundness of probabilistic forward simulations:

$$\begin{aligned} \mathcal{P}_1 \leq_{PF} \mathcal{P}_2 &\Rightarrow \mathcal{P}_1 \leq_{wPF} \mathcal{P}_2 && \text{(Proposition 3.10.1)} \\ &\Rightarrow \mathcal{P}_1 \leq_{DC} \mathcal{P}_2 && \text{(Proposition 3.10.3).} \end{aligned}$$

Now we prove completeness:

$$\begin{aligned} \mathcal{P}_1 \leq_{DC} \mathcal{P}_2 &\Rightarrow \text{Unfold}(\mathcal{P}_1) \leq_{DC} \mathcal{P}_1 \leq_{DC} \mathcal{P}_2 && \text{(Proposition 3.12)} \\ &\Rightarrow \text{Unfold}(\mathcal{P}_1) \leq_{DC} \mathcal{P}_2 && (\leq_{DC} \text{ is transitive)} \\ &\Rightarrow \text{Unfold}(\mathcal{P}_1) \leq_{PF} \mathcal{P}_2 && \text{(Proposition 6.1)} \\ &\Rightarrow \mathcal{P}_1 \leq_{PF} \text{Unfold}(\mathcal{P}_1) \leq_{PF} \mathcal{P}_2 && \text{(Proposition 3.11)} \\ &\Rightarrow \mathcal{P}_1 \leq_{PF} \mathcal{P}_2 && (\leq_{PF} \text{ is transitive). } \quad \square \end{aligned}$$

6.2. PAs with internal actions. Again, we start with tree-structured PAs.

PROPOSITION 6.4. *Let $\mathcal{P}_1, \mathcal{P}_2$ be PAs with \mathcal{P}_1 tree-structured. Then $\mathcal{P}_1 \leq_{DC} \mathcal{P}_2$ implies $\mathcal{P}_1 \leq_{wPF} \mathcal{P}_2$.*

Proof. Assume that $\mathcal{P}_1 \leq_{DC} \mathcal{P}_2$. Define the tester PA \mathcal{C} of \mathcal{P}_1 , the observer σ_1 , the trace distribution η , the scheduler σ_2 , the probabilistic execution ϵ , and the Θ sets as in the proof of Proposition 5.1. Define relation R according to (12) as in the proof of Proposition 6.1. Observe that formulas (13), (14), and (15) hold for the same reasons as before. Define a new relation R' as follows: $q_1 R' \mu_2$ iff there exists a measure μ'_2 such that $\mu_1 \Rightarrow \mu'_2$ and $q_1 R \mu'_2$. Observe that trivially $R \subseteq R'$. We show that R' is a weak probabilistic forward simulation from \mathcal{P}_1 to \mathcal{P}_2 .

For the start condition, we must show that $\bar{q}_1 R' \delta(\bar{q}_2)$. By Item 1 of Proposition 4.6, $\eta(C_{\bar{q}_1}) = 1$. This means that $\sum_{\alpha \in \Theta_{\bar{q}_1} | \text{trace}(\alpha) = \bar{q}_1} \epsilon(C_\alpha) = 1$. Let $\Theta'_{\bar{q}_1}$ be the set of elements of $\Theta_{\bar{q}_1}$ with trace \bar{q}_1 , and let ϵ' be the truncation of ϵ to $\Theta'_{\bar{q}_1}$. Then ϵ' assigns probability 1 to the set of finite execution fragments with trace \bar{q}_1 . Furthermore, observing that all elements of $\Theta'_{\bar{q}_1}$ are not prefixes of each other, we derive $\epsilon'(C_\alpha) = \epsilon'(\{\alpha\})$ for each $\alpha \in \Theta'_{\bar{q}_1}$. Finally, observing that each element of $\Theta_{\bar{q}_1} - \Theta'_{\bar{q}_1}$ is not a prefix of any element of $\Theta'_{\bar{q}_1}$, we derive $\epsilon'(C_\alpha) = \epsilon'(\{\alpha\}) = 0$ for each $\alpha \in \Theta_{\bar{q}_1} - \Theta'_{\bar{q}_1}$. Let μ'_2 be $\text{lstate}(\epsilon')$. By the definition of a weak hypertransition, $\delta(\bar{q}_2) \Rightarrow \mu'_2$. We show that $\bar{q}_1 R \mu'_2$, which suffices. Consider a state q_2 of \mathcal{P}_2 . By definition of μ'_2 , definition of $\Theta_{\bar{q}_1, q_2}$, and the fact that $\text{supp}(\epsilon') \subseteq \Theta_{\bar{q}_1}$, $\mu'_2(q_2) = \sum_{\alpha \in \Theta_{\bar{q}_1, q_2}} \epsilon'(\{\alpha\})$. Then the result follows immediately by observing that this equation corresponds to (12) since $\epsilon'(\{\alpha\}) = \epsilon'(C_\alpha)$ when $\alpha \in \Theta_{\bar{q}_1}$ and since $\sum_{\alpha \in \Theta_{\bar{q}_1}} \epsilon(C_\alpha) = 1$.

For the step condition, assume that $q_1 R' \mu_2$, and let $q_1 \xrightarrow{a}_1 \mu'_1$ be a transition of \mathcal{P}_1 , which we denote by tr . By definition of R' , there exists a measure μ'_2 such

that $\mu_2 \Rightarrow \mu'_2$ and $q_1 R \mu'_2$. We now show that there exists a measure ξ'_2 and a measure $\mu''_2 = \text{flatten}(\xi'_2)$ such that $\mu_1 R \xi'_2$ and $\mu'_2 \xrightarrow{a} \mu''_2$. Then $\mu_1 R' \xi'_2$, and by Proposition 3.6, $\mu_2 \xrightarrow{a} \mu''_2$.

The proof of existence of ξ'_2 and μ''_2 proceeds exactly as in the case of Proposition 6.1 except for the definition of the tr_α transitions. Thus, in the rest of the proof we construct the tr_α 's and prove that (19) still holds.

We introduce a special *conditional* construction that is needed for the definition of the tr_α 's. Let \mathcal{C}_{tr} be the same as \mathcal{C} except that the transition $q_1 \xrightarrow{ch} \mu$, where μ is uniquely determined by q_1 , is replaced by $q_1 \xrightarrow{ch} \delta(tr)$. Given a scheduler σ for $\mathcal{P}_2 \parallel \mathcal{C}$, define the scheduler $\sigma \mid tr$ for $\mathcal{P}_2 \parallel \mathcal{C}_{tr}$ that is the same as σ except that transition $q_1 \xrightarrow{ch} \delta(tr)$ of \mathcal{C}_{tr} is chosen whenever σ chooses $q_1 \xrightarrow{ch} \mu$. Given a probabilistic execution fragment ϵ' of $\mathcal{P}_2 \parallel \mathcal{C}$, generated by some scheduler σ , define $\epsilon' \mid tr$ to be the result of $\sigma \mid tr$ applied to $\mathcal{P} \parallel \mathcal{C}_{tr}$ from the start state of ϵ' . The intuition behind $\epsilon' \mid tr$ is that we study ϵ' under the condition that tr is the outcoming state of \mathcal{C} whenever $q_1 \xrightarrow{ch} \mu$ is scheduled. Then, the following two properties are valid:

1. $(\epsilon' \mid tr) \upharpoonright \mathcal{P}_2$ is a probabilistic execution fragment of \mathcal{P}_2 .
2. For each finite execution fragment α of $\mathcal{P}_2 \parallel \mathcal{C}$ where state tr occurs and such that $fstate(\alpha)$ is not of the form (\cdot, tr) , $(\epsilon' \mid tr)(C_\alpha) = k\epsilon(C_\alpha)$, where k is the size of $supp(\mu)$.

The first item follows immediately from Proposition 3.7, given that $\epsilon' \mid tr$ is a probabilistic execution fragment of $\mathcal{P}_2 \parallel \mathcal{C}_{tr}$. The second item follows directly from the definition of probability of a cone, since in ϵ' the probability associated with the edge $q \text{ ch } (\cdot, tr)$ is $1/k$ while in $\epsilon' \mid tr$ the probability of the same edge is 1.

We now define the tr_α 's. Consider an execution α of Θ_{q_1} such that $\epsilon(C_\alpha) > 0$. Let ϵ^1 be the truncation of ϵ at all the points in $\cup_{q \in supp(\mu'_1)} \Theta_q$, which is a probabilistic execution of $\mathcal{P}_2 \parallel \mathcal{C}$ by definition. Let ϵ^1_α be $\epsilon^1 \triangleright \alpha$, which is a probabilistic execution fragment of $\mathcal{P}_2 \parallel \mathcal{C}$ by definition. Finally, let ϵ^2_α be $(\epsilon^1_\alpha \mid tr) \upharpoonright \mathcal{P}_2$, which is a probabilistic execution fragment of \mathcal{P}_2 by Property 1.

By definition of Θ_{q_1} , $trace(\alpha) = \beta q_1$ for some finite trace β . Therefore, $\eta(C_{\beta q_1}) > 0$. Since q_1 enables at least one transition in \mathcal{P}_1 , specifically transition tr , (6) from Proposition 4.6 implies that $\eta(C_{\beta q_1 \text{ ch}}) = \eta(C_{\beta q_1})$. Thus, action ch occurs as the first external action with probability 1 in μ^1_α .

By (7) from Proposition 4.6, if the occurrence of action ch leads \mathcal{C} to state tr , then an action in $supp(\mu'_1)$ occurs eventually in ϵ with probability 1, leading \mathcal{C} to a state in $supp(\mu'_1)$, which is a truncation point according to the definition of ϵ^1 . Thus, the probability of termination in $\epsilon^1_\alpha \mid tr$ is 1, as well as the probability of termination in ϵ^2_α ; that is, ϵ^2_α assigns probability 1 to the set of finite executions. Furthermore, given that action a is uniquely determined by μ'_1 (\mathcal{P}_1 is tree-structured), again by (7) from Proposition 4.6 all finite executions α' with $\epsilon^2_\alpha(\alpha') > 0$ have trace $trace(a)$. Thus, ϵ^2_α is a representation of a weak combined transition labeled by a from $lstate(\alpha) \upharpoonright \mathcal{P}_2$. Denote such transition by tr_α .

We are left to show that (19) still holds. That is,

$$\sum_{\alpha \in \Theta_{q_1}} \epsilon(C_\alpha) \mu_{tr_\alpha}(q_2) = k \sum_{q \in supp(\mu'_1), \alpha \in \Theta_{q, q_2}} \epsilon(C_\alpha).$$

We consider first the term $\mu_{tr_\alpha}(q_2)$. From the definition of tr_α and of weak

combined transition we get

$$\mu_{tr_\alpha}(q_2) = \sum_{\alpha' | lstate(\alpha')=q_2} \epsilon_\alpha^2(\alpha').$$

By applying the definition of projection, and using the fact that $\epsilon_\alpha^1 | tr$ assigns probability 1 to the set of finite executions, we get

$$\mu_{tr_\alpha}(q_2) = \sum_{\alpha' | lstate(\alpha' \upharpoonright \mathcal{P}_2)=q_2} (\epsilon_\alpha^1 | tr)(\alpha').$$

Given that the truncation points of ϵ^1 are all at the $\cup_{q \in \text{supp}(\mu'_1)} \Theta_q$ points, the only finite executions α' that have nonzero probability are such that $\alpha \wedge \alpha'$ is in some set Θ_q . Furthermore, given that no execution in $\cup_{q \in \text{supp}(\mu'_1)} \Theta_q$ is a prefix of another (our PAs are tree-structured and all actions in $\text{supp}(\mu'_1)$ occur in different branches), the probabilities of the finite executions can be replaced by the probabilities of their cones, thus getting

$$\mu_{tr_\alpha}(q_2) = \sum_{q \in \text{supp}(\mu'_1)} \sum_{\alpha' | \alpha \wedge \alpha' \in \Theta_{q,q_2}} (\epsilon_\alpha^1 | tr)(C_{\alpha'}).$$

By property 2 we can get rid of the conditional on tr by introducing a k factor, thus getting

$$(22) \quad \mu_{tr_\alpha}(q_2) = \sum_{q \in \text{supp}(\mu'_1)} \sum_{\alpha' | \alpha \wedge \alpha' \in \Theta_{q,q_2}} k \epsilon_\alpha^1(C_{\alpha'}).$$

By replacing $\mu_{tr_\alpha}(q_2)$ according to (22) in the left-hand side of (19) and by rearranging terms algebraically, we obtain

$$\sum_{\alpha \in \Theta_{q_1}} \epsilon(C_\alpha) \mu_{tr_\alpha}(q_2) = k \sum_{q \in \text{supp}(\mu'_1)} \sum_{\alpha \in \Theta_{q_1}} \sum_{\alpha' | \alpha \wedge \alpha' \in \Theta_{q,q_2}} \epsilon(C_\alpha) \epsilon_\alpha^1(C_{\alpha'}).$$

By using that $\epsilon_\alpha^1 = \epsilon^1 \triangleright \alpha$ (by definition) and Proposition 3.15, the two probabilities in the equation above can be grouped into $\epsilon(C_{\alpha \wedge \alpha'})$. By observing that all elements in Θ_{q,q_2} , with $q \in \text{supp}(\mu'_1)$, have a prefix in Θ_{q_1} , the intermediate sum can be removed, thus getting

$$\sum_{\alpha \in \Theta_{q_1}} \epsilon(C_\alpha) \mu_{tr_\alpha}(q_2) = k \sum_{q \in \text{supp}(\mu'_1)} \sum_{\alpha \in \Theta_{q,q_2}} \epsilon(C_\alpha),$$

which is (19), as needed. \square

THEOREM 6.5. *Let $\mathcal{P}_1, \mathcal{P}_2$ be PAs. Then $\mathcal{P}_1 \leq_{DC} \mathcal{P}_2$ iff $\mathcal{P}_1 \leq_{wPF} \mathcal{P}_2$.*

Proof. Soundness of weak probabilistic forward simulations follows immediately from Proposition 3.10. Completeness is established as follows:

$$\begin{aligned} \mathcal{P}_1 \leq_{DC} \mathcal{P}_2 &\Rightarrow \text{Unfold}(\mathcal{P}_1) \leq_{DC} \mathcal{P}_1 \leq_{DC} \mathcal{P}_2 \quad (\text{Proposition 3.12}) \\ &\Rightarrow \text{Unfold}(\mathcal{P}_1) \leq_{DC} \mathcal{P}_2 \quad (\leq_{DC} \text{ is transitive}) \\ &\Rightarrow \text{Unfold}(\mathcal{P}_1) \leq_{wPF} \mathcal{P}_2 \quad (\text{Proposition 6.4}) \\ &\Rightarrow \mathcal{P}_1 \leq_{PF} \text{Unfold}(\mathcal{P}_1) \leq_{wPF} \mathcal{P}_2 \quad (\text{Proposition 3.11}) \\ &\Rightarrow \mathcal{P}_1 \leq_{wPF} \text{Unfold}(\mathcal{P}_1) \leq_{wPF} \mathcal{P}_2 \quad (\text{Proposition 3.10}) \\ &\Rightarrow \mathcal{P}_1 \leq_{wPF} \mathcal{P}_2 \quad (\leq_{wPF} \text{ is transitive}). \quad \square \end{aligned}$$

7. Concluding remarks. We have characterized the trace distribution precongruence for nondeterministic and probabilistic automata, with and without internal actions, in terms of four kinds of simulation relations, \leq_F , \leq_{wF} , \leq_{PF} , and \leq_{wPF} . In particular, this shows that probabilistic contexts are capable of observing all the distinctions that can be expressed using these simulation relations. Our main technical contribution is the definition of special contexts, called testers, that, under the action of an appropriate scheduler, can reveal the branching structure of a PA via a trace distribution. Some technical improvements are possible. For example, our finite branching restriction can be relaxed to countable branching, simply by replacing uniform distributions in the tester automata by other distributions such as exponential distributions. In that case, however, calculations become more complicated. We have also considered nondeterministic and probabilistic automata with countably many states and actions. Again, this restriction can be relaxed at the cost of complicating the definition of the σ -field of execution fragments: the generators would be arbitrary unions of cones, and the measure of a union of cones would be just the sum of the measures of each single cone. Indeed, discrete transitions and discrete schedulers ensure that there are at most countably many cones with nonzero measure.

Although in this paper we reach a point where we have a full understanding of trace distribution precongruence as a branching relation, a natural question is whether it is possible to define linear probabilistic extensions of language inclusion. A potential approach is to consider ordinary traces paired with their maximal or minimal probabilities under all schedulers, but the induced preorder relations do not appear to be interesting. Another approach, followed in [19], is to extend classical testing preorders by considering the maximal and minimal probabilities of success of a test; however, even in such case the resulting precongruence is characterized in terms of simulation relations that, although weaker than the relations studied in this paper, are still branching relations. Other approaches ensure compositionality of trace distribution inclusion by restricting parallel composition so that the nondeterminism of each component is resolved based only on externally visible behavior of the other components. This approach is investigated in [10] in a synchronous model. In [8, 7], an asynchronous switched probabilistic input/output automaton model (PIOA) is presented, which uses a token structure to eliminate global nondeterministic choices. This token structure ensures that, at any point in time, there is at most one active component in a system and this unique component determines the next active component. Thus, global scheduling is performed jointly by all local schedulers, which have access to local information only. A notion of switched probabilistic systems is defined, which are switched PIOAs paired with sets of *acceptable* I/O schedulers. A trace-style semantics for switched probabilistic systems is given, using the notion of likelihood assignments. This semantics is shown to be compositional with respect to a parallel operator that combines local I/O schedulers into a joint I/O scheduler. Thus, the approach of [8, 7] can be characterized as *schedule-and-compose*, where local nondeterministic choices are resolved before the components are placed in parallel. In [7] also a similar strategy is pursued, but without the token structure. Instead, several axioms are imposed on the reactive and generative transition structures, so that branching occurs only when it is meant to be globally visible (i.e., the branches carry different visible action labels). These axioms capture a local-oblivious assumption on adversaries, which is well known in the area of randomized consensus [9, 3]. The model is proven to be compositional with respect to a schedule-and-compose operator similar to that in [8].

Acknowledgment. We are grateful to Erik de Vink for numerous detailed observations and suggestions for improvements.

REFERENCES

- [1] S. AGGARWAL, *Time Optimal Self-Stabilizing Spanning Tree Algorithms*, Master's thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, 1994; also available as Technical Report MIT/LCS/TR-632.
- [2] S. ANDOVA AND T. WILLEMSE, *Branching bisimulation for probabilistic systems: Characteristics and decidability*, Theoret. Comput. Sci., 356 (2006), pp. 325–355.
- [3] Y. AUMANN AND M. A. BENDER, *Efficient low-contention asynchronous consensus with the value-oblivious adversary scheduler*, Distrib. Comput., 17 (2005), pp. 191–207.
- [4] F. BARTELS, A. SOKOLOVA, AND E. DE VINK, *A hierarchy of probabilistic system types*, Theoret. Comput. Sci., 327 (2004), pp. 3–22.
- [5] J. A. BERGSTRA, J. W. KLOP, AND E.-R. OLDEROG, *Readies and failures in the algebra of communicating processes*, SIAM J. Comput., 17 (1988), pp. 1134–1177.
- [6] L. CHEUNG, *Randomized wait-free consensus using an atomicity assumption*, in Principles of Distributed Systems (Proceedings of OPODIS 2005, Pisa, Italy), J. H. Anderson, G. Prencipe, and R. Wattenhofer, eds., Lecture Notes in Comput. Sci. 3974, Springer, New York, 2006, pp. 47–60.
- [7] L. CHEUNG, *Reconciling Nondeterministic and Probabilistic Choices*, Ph.D. thesis, Faculty of Science, Radboud University, Nijmegen, The Netherlands, 2006.
- [8] L. CHEUNG, N. A. LYNCH, R. SEGALA, AND F. W. VAANDRAGER, *Switched PIOA: Parallel composition via distributed scheduling*, Theoret. Comput. Sci., 365 (2006), pp. 83–108.
- [9] B. CHOR, A. ISRAELI, AND M. LI, *Wait-free consensus using asynchronous hardware*, SIAM J. Comput., 23 (1994), pp. 701–712.
- [10] L. DE ALFARO, T. A. HENZINGER, AND R. JHALA, *Compositional methods for probabilistic systems*, in Proceedings of CONCUR 01, Aalborg, Denmark, 2001, K. G. Larsen and M. Nielsen, eds., Lecture Notes in Comput. Sci. 2154, Springer, New York, 2001, pp. 351–365.
- [11] W. FELLER, *An Introduction to Probability Theory and Its Applications. Volume 1*, John Wiley & Sons, New York, 1950.
- [12] R. J. VAN GLABBEEK, *The linear time-branching time spectrum I. The semantics of concrete, sequential processes*, in Handbook of Process Algebra, J. A. Bergstra, A. Ponse, and S. A. Smolka, eds., North-Holland, Amsterdam, 2001, pp. 3–99.
- [13] R. J. VAN GLABBEEK, S. A. SMOLKA, AND B. STEFFEN, *Reactive, generative, and stratified models of probabilistic processes*, Inform. and Control, 121 (1995), pp. 59–80.
- [14] J. F. GROOTE AND F. W. VAANDRAGER, *Structured operational semantics and bisimulation as a congruence*, Inform. and Comput., 100 (1992), pp. 202–260.
- [15] H. A. HANSSON, *Time and Probability in Formal Design of Distributed Systems*, Real-Time Safety Critical Systems 1, Elsevier, New York, 1994.
- [16] A. HINTON, M. Z. KWIATKOWSKA, G. NORMAN, AND D. PARKER, *Prism: A tool for automatic verification of probabilistic systems*, in TACAS, H. Hermanns and J. Palsberg, eds., Lecture Notes in Comput. Sci. 3920, Springer-Verlag, New York, 2006, pp. 441–444.
- [17] C. A. R. HOARE, *Communicating Sequential Processes*, Prentice-Hall, Englewood Cliffs, NJ, 1985.
- [18] B. JONSSON AND K. G. LARSEN, *Specification and refinement of probabilistic processes*, in Proceedings of the 6th Annual Symposium on Logic in Computer Science, Amsterdam, IEEE Press, Piscataway, NJ, 1991, pp. 266–277.
- [19] B. JONSSON AND W. YI, *Testing preorders for probabilistic processes can be characterized by simulations*, Theoret. Comput. Sci., 282 (2002), pp. 33–51.
- [20] M. Z. KWIATKOWSKA AND G. NORMAN, *Verifying randomized byzantine agreement*, in FORTE, D. Peled and M. Y. Vardi, eds., Lecture Notes in Comput. Sci. 2529, Springer-Verlag, New York, 2002, pp. 194–209.
- [21] L. LAMPORT, *Specifying concurrent program modules*, ACM Trans. Prog. Lang. Syst., 5 (1983), pp. 190–222.
- [22] K. G. LARSEN AND A. SKOU, *Bisimulation through probabilistic testing*, Inform. and Comput., 94 (1991), pp. 1–28.
- [23] N. A. LYNCH, I. SAIAS, AND R. SEGALA, *Proving time bounds for randomized distributed algorithms*, in Proceedings of the 13th Annual ACM Symposium on the Principles of Distributed Computing, Los Angeles, CA, 1994, ACM, New York, 1994, pp. 314–323.

- [24] N. A. LYNCH, R. SEGALA, AND F. W. VAANDRAGER, *Compositionality for probabilistic automata*, in Proceedings of the 14th International Conference on Concurrency Theory (CONCUR 2003), Marseille, France, R. Amadio and D. Lugiez, eds., Lecture Notes in Comput. Sci. 2761, Springer-Verlag, New York, 2003, pp. 208–221.
- [25] N. A. LYNCH AND M. R. TUTTLE, *An introduction to input/output automata*, CWI Quarterly, 2 (1989), pp. 219–246.
- [26] N. A. LYNCH AND F. W. VAANDRAGER, *Forward and backward simulations, I: Untimed systems*, Inform. and Comput., 121 (1995), pp. 214–233.
- [27] R. MILNER, *Communication and Concurrency*, Prentice–Hall, Englewood Cliffs, NJ, 1989.
- [28] G. NORMAN, *Analysing randomized distributed algorithms*, in Validation of Stochastic Systems, C. Baier, B. R. Haverkort, H. Hermanns, J.-P. Katoen, and M. Siegle, eds., Lecture Notes in Comput. Sci. 2925, Springer-Verlag, New York, 2004, pp. 384–418.
- [29] A. PHILIPPOU, I. LEE, AND O. SOKOLSKY, *Weak bisimulation for probabilistic systems*, in Proceedings of CONCUR 2000, University Park, PA, C. Palamidessi, ed., Lecture Notes in Comput. Sci. 1877, Springer-Verlag, New York, 2000, pp. 334–349.
- [30] A. POGOSYANTS, R. SEGALA, AND N. A. LYNCH, *Verification of the randomized consensus algorithm of Aspnes and Herlihy: A case study*, Distributed Computing, 13 (2000), pp. 155–186.
- [31] R. SEGALA, *Compositional trace-based semantics for probabilistic automata*, in Proceedings of CONCUR 95, Philadelphia, PA, I. Lee and S. A. Smolka, eds., Lecture Notes in Comput. Sci. 962, Springer-Verlag, New York, 1995, pp. 234–248.
- [32] R. SEGALA, *Modeling and Verification of Randomized Distributed Real-Time Systems*, Ph.D. thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, 1995; also available as Technical Report MIT/LCS/TR-676.
- [33] R. SEGALA, *Testing probabilistic automata*, in Proceedings of CONCUR 96, Pisa, Italy, 1996, U. Montanari and V. Sassone, eds., Lecture Notes in Comput. Sci. 1119, Springer, New York, 1996, pp. 299–314.
- [34] R. SEGALA AND N. A. LYNCH, *Probabilistic simulations for probabilistic processes*, Nordic J. Comput., 2 (1995), pp. 250–273.
- [35] R. SEGALA AND A. TURRINI, *Comparative analysis of bisimulation relations on alternating and nonalternating probabilistic models*, in Proceedings of the Second International Conference on the Quantitative Evaluation of Systems (QEST 2005), Torino, Italy, IEEE Computer Society, Piscataway, NJ, 2005, pp. 44–53.
- [36] A. SOKOLOVA AND E. P. DE VINK, *Probabilistic automata: System types, parallel composition, and comparison*, in Validation of Stochastic Systems—A Guide to Current Research, C. Baier, B. R. Haverkort, H. Hermanns, J.-P. Katoen, and M. Siegle, eds., Lecture Notes in Comput. Sci. 2925, Springer-Verlag, New York, 2004, pp. 1–43.
- [37] M. I. A. STOELINGA, *Alea Jacta Est: Verification of Probabilistic, Real-Time, and Parametric Systems*, Ph.D. thesis, Faculty of Science, University of Nijmegen, Nijmegen, The Netherlands, 2002.
- [38] M. I. A. STOELINGA, *An introduction to probabilistic automata*, Bull. European Assoc. Theoret. Comput. Sci., 78 (2002), pp. 176–198.
- [39] M. I. A. STOELINGA AND F. W. VAANDRAGER, *Root contention in IEEE 1394*, in Proceedings of the 5th International AMAST Workshop on Formal Methods for Real-Time and Probabilistic Systems, Bamberg, Germany, J.-P. Katoen, ed., Lecture Notes in Comput. Sci. 1601, Springer-Verlag, New York, 1999, pp. 53–74.
- [40] M. Y. VARDI, *Automatic verification of probabilistic concurrent finite-state programs*, in Proceedings of the 26th IEEE Symposium on Foundations of Computer Science, Portland, OR, 1985, IEEE Press, Piscataway, NJ, pp. 327–338.