Fully Abstract Characterizations of Testing Preorders for Probabilistic Processes

Shoji Yuen*
Dept. of Information Engineering
Nagoya University
Chikusa, Nagoya 464-01
Japan
yuen@nuie.nagoya-u.ac.jp

Rance Cleaveland** Zeynep Dayar
Dept. of Computer Science
North Carolina State University
Raleigh, NC 27695
USA
rance@csc.ncsu.edu
gr-zd@druid.csc.ncsu.edu

Scott A. Smolka***
Dept. of Computer Science
SUNY at Stony Brook
Stony Brook, NY 11794-4400
USA
sas@cs.sunysb.edu

Abstract. We present alternative characterizations of the testing preorders for probabilistic processes proposed in [CSZ92]. For a given probabilistic process, the characterization takes the form of a mapping from *probabilistic traces* to the interval $[0, 1]$, where a probabilistic trace is an alternating sequence of actions and probability distributions over actions.

Our results, like those of [CSZ92], pertain to divergence-free probabilistic processes, and are presented in two stages: probabilistic tests without internal τ-transitions are considered first, followed by probabilistic tests with τ-transitions. In each case, we show that our alternative characterization is *fully abstract* with respect to the corresponding testing preorder, thereby resolving an open problem in [CSZ92]. In the second case, we use the alternative characterization to show that the testing preorder is actually an equivalence relation.

Finally, we give *proof techniques*, derived from the alternative characterizations, for establishing preorder relationships between probabilistic processes. The utility of these techniques is demonstrated by means of some simple examples.

1 Introduction

Communicating systems often exhibit behavior that is probabilistic or statistical in nature. For example, one may observe that a faulty communication link drops a message 2% of the time or that a site in a network is down with probability 0.05. Analyzing such behavior can thus lead to characterizations of the relative reliability of systems.

* Research performed while on leave at SUNY Stony Brook.

** Research supported in part by NSF/DARPA grant CCR-9014775, NSF grant CCR-9120995, ONR Young Investigator Award N00014-92-J-1582, and NSF Young Investigator Award CCR-9257963.

*** Research supported in part by NSF Grants CCR-9120995 and CCR-9208585, and AFOSR Grant F49620-93-1-0250DEF.

In [CSZ92], Cleaveland et al. proposed a testing preorder for probabilistic processes based on the probabilities with which a process passes tests. The basic idea is the following: probabilistic processes $\mathcal{P}$ and $\mathcal{Q}$ are related by the preorder if for all probabilistic tests $\mathcal{T}$, the probability by which $\mathcal{P}$ passes $\mathcal{T}$ is no greater than the probability by which $\mathcal{Q}$ passes $\mathcal{T}$.

The testing preorder of [CSZ92] ($\sqsubseteq$) can be viewed as an extension of the testing preorder of DeNicola and Hennessy [DH83, Hen88] for nondeterministic processes ($\sqsubseteq_{test}$) as follows. Let h be the homomorphism that maps probabilistic processes and tests onto nondeterministic processes and tests by erasing probabilities. Then, for probabilistic process $\mathcal{P}$ and probabilistic test $\mathcal{T}$, if $\mathcal{P}$ passes $\mathcal{T}$ with probability 1 (respectively a non-zero probability) in the probabilistic setting, then $h(\mathcal{P})$ *must* (respectively *may*) pass $h(\mathcal{T})$ in the nondeterministic setting. It follows that $\mathcal{P} \sqsubseteq \mathcal{Q}$ implies $h(\mathcal{P}) \sqsubseteq_{test} h(\mathcal{Q})$.

While intuitively appealing, the operational definition of $\sqsubseteq$ can be difficult to reason about. In particular, establishing that one process is related to another requires a consideration of the behavior of both processes in the context of all possible tests. Accordingly, it would be useful to have an alternative, more denotational characterization of $\sqsubseteq$ to ease the task of establishing relationships between probabilistic processes. Moreover, this characterization should be *fully abstract*, relating probablistic processes in exactly the same manner as $\sqsubseteq$.

In this paper, we present such a fully abstract characterization of $\sqsubseteq$, thereby resolving an open problem in [CSZ92]. For a given probabilistic process, our characterization takes the form of a mapping from *probabilistic traces* to the interval $[0, 1]$, where a probabilistic trace is a sequence of pairs, with the first component of each pair being a probability distribution over actions and the second being an action.

Our results, like those of [CSZ92], pertain to divergence-free probabilistic processes, and are presented in two stages: probabilistic tests without internal τ-transitions are considered first, followed by probabilistic tests with τ-transitions. This two-tier presentation allows us to isolate the impact of τ-transitions in tests on the theory. In both cases, we show that our alternative characterization is fully abstract with respect to the corresponding testing preorder.

Using the alternative characterizations, we also show that the preorder in which tests may perform τ's is a strict refinement of the one in which tests are τ-free and, more importantly, that the finer preorder is actually an equivalence relation, a fact that went unnoticed in [CSZ92]. Finally, we give *proof techniques*, derived from the alternative characterizations, that further simplify the task of establishing the preorders. The utility of these techniques is demonstrated by means of some simple examples.

In terms of related work, Wu et al. [WSS94] propose a testing equivalence for *probabilistic I/O automata* and present a fully abstract characterization of the equivalence. In I/O automata [LT87], a distinction is made between *input* actions, which come from the environment and are always enabled, and *output* actions, which are locally controlled. Accordingly, in probabilistic I/O automata, separate probability distributions are associated with each input action and a single distribution with all locally controlled actions. Moreover, a *delay parameter*, which plays a key role in probabilistic I/O automata composition, is associated with each automaton state. Probabilistic I/O automata are thus a kind of hybrid of the reactive and generative models of probabilistic processes [vGSST90]. In contrast, processes and tests in our model are strictly generative and our approach is grounded in the testing theory of [DH83, Hen88].

Yi and Larsen [YL92] add an internal probabilistic choice operator to CCS in order to model processes that exhibit both nondeterministic and probabilistic behavior. They then develop a theory of probabilistic testing in which tests are processes that can perform a distinguished action to report success. In this framework a process passes a test with a set of probabilities, with each probability representing a possible resolution of the nondeterministic (i.e. nonprobabilistic) behavior of the process and the test. They also define testing preorders based on DeNicola and Hennessy's [DH83] testing preorder.

In [Chr90], Christoff also considers testing preorders for probabilistic processes and their fully abstract characterizations. His approach, however, differs from our own on a number of key points. Christoff's tests are deterministic and there is no notion of a success state; instead he defines testing preorders based on the probabilities induced by the interaction of a process and a test on L^*, where L is a set of observable actions. In our setting, tests are just probabilistic processes with a set of success states, and the testing preorder is defined in terms of the probability of a process successfully passing a test. Christoff's fully abstract models are defined in terms of "probabilistic functions" that map $(2^L - \emptyset)^* \times L^*$ to $[0,1]$. In our case, instead of using elements of $(2^L - \emptyset)^*$ we use sequences of probability distributions.

In [LS92a], Larsen and Skou introduce probabilistic bisimulation and a characterization of the probabilistic bisimulation as a form of probabilistic testing. In [LS92b], they define a calculus of reactive processes and an accompanying probabilistic modal logic. The latter is used to investigate compositional verification of probabilistic processes. Complete axiomatizations of both the calculus and logic are given. [HJ90] axiomatizes a calculus of *alternating processes* whose behavior alternates between nondeterministic and probabilistic states.

The structure of the rest of this paper is as follows. Section 2 reviews the testing theory of probabilistic processes put forth in [CSZ92]. Section 3 presents the alternative characterization of the testing preorder for tests without τ-transitions, and a proof technique derived from the characterization. Section 4 extends the results of the previous section to tests with τ-transitions and proves that the testing preorder is an equivalence relation. Section 5 concludes and contains our directions for future work.

2 Testing Preorders for Probabilistic Processes

In this section, we briefly review the testing theory of probabilistic processes from [CSZ92].

Definition 1 (Probabilistic Process). A *probabilistic process* is a 4-tuple (P, Act, μ, p_0) where

- P is a countable set of *states*;
- Act is a countable set of *observable actions* not containing the distinguished *silent action*, τ (we write Act_τ for $Act \cup \{\tau\}$);
- $\mu : P \times Act_\tau \times P \rightarrow [0,1]$, the *transition probability distribution function*, satisfies: for all $p \in P$, $\sum_{p' \in P, \alpha \in Act_\tau} \mu(p, \alpha, p') \in \{0,1\}$;
- $p_0 \in P$ is the *start state.*

We write **PP** for the class of all probabilistic processes; $p \xrightarrow{\alpha,\pi} p'$ if $\pi = \mu(p, \alpha, p') > 0$; and $\mu(p, \alpha)$ for $\sum_{p' \in P} \mu(p, \alpha, p')$. We refer to $p \xrightarrow{\alpha,\pi} p'$ as a *transition* of $\mathcal{P}$.

We require that probabilistic processes be *finite-branching* and *divergence-free*. A probabilistic process $\mathcal{P} = (P, Act, \mu, p_0)$ is finite-branching if for all $p \in P$ the set $\{\langle \alpha, p \rangle \mid \mu(p, \alpha, p') > 0\}$ is finite, while $\mathcal{P}$ is divergence-free if for no $p \in P$ is there an infinite sequence of states $p_0, p_1, \ldots$ with $p = p_0$ and $\mu(p_i, \tau, p_{i+1}) > 0$.

If A is a set then we call $d : A \to [0, 1]$ a *distribution over* A if $\sum_{a \in A} d(a) \in \{0, 1\}$; note the (slight) departure from traditional terminology. We denote the set of all distributions over A by D_A; thus, D_{Act_τ} represents the set of all distributions over Act_τ. If p is a state in process (P, Act, μ, p_0) then we write $\mu(p)$ for the distribution that satisfies $\mu(p)(\alpha) = \mu(p, \alpha)$ for all $\alpha \in Act_\tau$.

A *probabilistic test* is defined as a probabilistic process with a set of *successful* states.

Definition 2 (Probabilistic Test). A *probabilistic test* is a 5-tuple (T, Act, μ, t_0, S), where (T, Act, μ, t_0) is a probabilistic process and $S \subseteq T$. We write $\mathbf{T}$ for the class of all probabilistic tests.

Herewith, we restrict probabilistic tests to be *finite* in the sense that there exists no infinite sequence of transitions from any state. We also drop the adjective "probabilistic" when no confusion arises.

The preorder studied in this paper relates processes on the basis of the probabilities with which they pass tests. To define this probability formally, we introduce *interaction systems*.

Definition 3 (Interaction System). Let $\mathcal{P} = (P, Act, \mu_P, p_0)$ be a process and $\mathcal{T} = (T, Act, \mu_T, t_0, S)$ be a test. The *interaction system* for $\mathcal{P}$ and $\mathcal{T}$, $\mathcal{P} \| \mathcal{T}$, is of the form $(P \times T, Act, \mu_I, \langle p_0, t_0 \rangle)$ where μ_I is defined as follows. First let the *normalization function* $\nu : P \times T \to [0, 1]$ be given by:

$$\nu(p, t) = \Big(\sum_{\alpha \in Act} \mu_P(p, \alpha) \cdot \mu_T(t, \alpha) \Big) + \mu_P(p, \tau) + \mu_T(t, \tau) - \mu_P(p, \tau) \cdot \mu_T(t, \tau).$$

Then:

$$\mu_I(\langle p, t \rangle, \alpha, \langle p', t' \rangle) = \begin{cases} 0 & \text{if } \nu(p,t) = 0 \\ \frac{\mu_P(p,\alpha,p') \cdot (1 - \mu_T(t,\tau))}{\nu(p,t)} & \text{if } t = t',\ \nu(p,t) \neq 0, \text{ and } \alpha = \tau \\ \frac{\mu_P(t,\alpha,t') \cdot (1 - \mu_P(p,\tau))}{\nu(p,t)} & \text{if } p = p',\ \nu(p,t) \neq 0, \text{ and } \alpha = \tau \\ \frac{\mu_P(p,\alpha,p') \cdot \mu_T(t,\alpha,t')}{\nu(p,t)} & \text{otherwise} \end{cases}$$

Note that an interaction system has the same form as a process, and hence all the notation defined for processes may also be applied to interaction systems. Intuitively, we wish to view processes and tests as *reactive systems*; thus, if in the course of applying a test to a process an interaction between the two is possible, then some interaction is guaranteed to take place. The definition of interaction system captures this notion: the normalization factor may be seen as the total probability that an interaction happens, and this probability is used to condition the probabilities of individual interactions.

An interaction system can also be interpreted in terms of a coin-flipping game. $\mathcal{P}$ and $\mathcal{T}$, in states p and t, flip Act-labeled coins weighted according to $\mu_P(p)$ and $\mu_T(t)$, respectively. If the actions, which can be τ, indicated by the coins are same, then the test and the process synchronize on the action and the interaction system performs a transition. If the coin of the process indicates τ and the coin of the test does not, then the process does its τ independently, and vice versa. Finally, if the actions are different, and neither is τ, then the coins are flipped again.

In order to define the probability with which a process passes a test, we first introduce the notion of a computation in an interaction system.

Definition 4. Let $\mathcal{P} = (P, Act, \mu_P, p_0)$ be a process and $\mathcal{T} = (T, Act, \mu_T, t_0, S)$ a test, with $\mathcal{P}\|\mathcal{T}$ the associated interaction system.

- We call $\langle p, t \rangle \in P \times T$ a *configuration.* A configuration $\langle p, t \rangle$ is *successful* if $t \in S$.
- A *computation* from a configuration $\langle p, t \rangle$ is a maximal sequence of the form
$$\langle p, t \rangle \xrightarrow{\alpha_1, \pi_1} \langle p_1, t_1 \rangle \xrightarrow{\alpha_2, \pi_2} \cdots \xrightarrow{\alpha_n, \pi_n} \langle p_n, t_n \rangle$$
where no configuration $\langle p_i, t_i \rangle$ is successful for $0 \leq i < n$. Such a computation is *successful* if $\langle p_n, t_n \rangle$ is successful. For a configuration $\langle p, t \rangle$, let $\mathcal{C}_{\langle p,t \rangle}$ be the set of all computations beginning with $\langle p, t \rangle$, and let $\mathcal{S}_{\langle p,t \rangle} \subseteq \mathcal{C}_{\langle p,t \rangle}$ represent the set of all such successful computations. In addition, we use $\mathcal{C}_{\mathcal{P}\|\mathcal{T}}$ and $\mathcal{S}_{\mathcal{P}\|\mathcal{T}}$ to denote the sets $\mathcal{C}_{\langle p_0,t_0 \rangle}$ and $\mathcal{S}_{\langle p_0,t_0 \rangle}$, respectively.
- Let C be the above computation from $\langle p, t \rangle$. The the sequence F of transitions given by
$$F = \langle p, t \rangle \xrightarrow{\alpha_1, \pi_1} \langle p_1, t_1 \rangle \xrightarrow{\alpha_2, \pi_2} \cdots \xrightarrow{\alpha_m, \pi_m} \langle p_m, t_m \rangle$$
is said to be a *prefix* of C, written $F \preceq C$, if $m \leq n$. (Note that F is not necessarily a computation.)
- Let C be a computation. Then $Pr(C) \in [0, 1]$ may be defined inductively as follows.
$$Pr(\langle p, t \rangle) = 1$$
$$Pr(\langle p, t \rangle \xrightarrow{\alpha, \pi} C') = \pi \cdot Pr(C')$$
If $\mathcal{C} \subseteq \mathcal{C}_{\langle p,t \rangle}$ for some $\langle p, t \rangle$ then we define $Pr(\mathcal{C}) = \sum_{C \in \mathcal{C}} Pr(C)$.

Intuitively, $Pr(C)$ records the probability that the process and test engage in computation C, starting from the initial configuration of C. It is straightforward to establish that for any configuration $\langle p, t \rangle$, Pr is a probability distribution over $\mathcal{C}_{\langle p,t \rangle}$. The probability with which a process passes a test then becomes the cumulative probability of the successful computations in the interaction system.

Definition 5. Let $\mathcal{P} = (P, Act, \mu_P, p_0)$ be a process and $\mathcal{T} = (T, Act, \mu_T, t_0, S)$ be a test. For $\pi \in [0, 1]$ we define the following predicates:

- $p \; pass_\pi \; t$ if $Pr(\mathcal{S}_{p\|t}) \geq \pi$
- $\mathcal{P} \; pass_\pi \; \mathcal{T}$ if $p_0 \; pass_\pi \; t_0$

$\mathcal{P}\ pass_\pi\ \mathcal{T}$ holds if $\mathcal{P}$ has successful computations from $\langle\, p_0, t_0 \,\rangle$ in $\mathcal{P}||\mathcal{T}$ with total probability at least π. We now may define the testing preorder with respect to a set of tests as follows.

Definition 6 (Testing Preorder for Probablistic Processes). Let $\mathcal{E}$ be a class of probabilistic tests, and let $\mathcal{P} = (P, Act, \mu_P, p_0)$ and $\mathcal{Q} = (Q, Act, \mu_Q, q_0)$. Then $\mathcal{P} \sqsubseteq^{\mathcal{E}} \mathcal{Q}$ if for all $\mathcal{T} \in \mathcal{E}$ and $\pi \in [0,1]$, $\mathcal{P}\ pass_\pi\ \mathcal{T}$ implies $\mathcal{Q}\ pass_\pi\ \mathcal{T}$.

Intuitively, $\mathcal{P} \sqsubseteq^{\mathcal{E}} \mathcal{Q}$ means that $\mathcal{P}$ is "more reliable" with respect to $\mathcal{E}$ than $\mathcal{Q}$. In [CSZ92], $\mathcal{E}$ is taken to be the class of all tests; in this paper, we first consider a restricted class of tests in which no τ's are present before examining all tests.

3 Alternative Characterization for Tests without τ

In this section, we investigate the testing preorder for tests without τ-transitions. A test $\mathcal{T} = (T, Act, \mu_T, t_0, S)$ is called *τ-free* if $\mu_T(t, \tau, t') = 0$ for all $t, t' \in T$. We write $\mathbf{T}_0$ for the set of all τ-free tests. For notational convenience, we denote the preorder for τ-free tests as $\sqsubseteq_0$ instead of $\sqsubseteq^{\mathbf{T}_0}$.

3.1 The Alternative Characterization

The alternative characterization for $\sqsubseteq_0$ relies on *probabilistic traces* and the probabilities with which processes may execute probabilistic traces.

Definition 7 (Probabilistic Trace). A *probabilistic trace* over *Act* is an element of $(D_{Act} \times Act)^*$, where in each $\langle\, d, \alpha \,\rangle$ the set $\{\beta \mid d(\beta) > 0\}$ is finite. We use Tr_{Act} to denote the set of all probabilistic traces over *Act*.

A probabilistic trace consists of a sequence of pairs, with the first component of each pair being a probability distribution over actions and the second being an action. Intuitively, a trace represents a possible "interaction session" with a probabilistic process; each distribution defines the probabilities with which different actions are enabled by the environment, while the action component specifies the action that the environment and process agree to execute. Given such a trace, our goal is to define the probability that the process "executes" the trace to completion. Before doing this, however, we introduce some useful notation on traces.

Definition 8. Let $s = \langle\, d_1, \alpha_1 \,\rangle \cdots \langle\, d_i, \alpha_i \,\rangle \cdots \langle\, d_n, \alpha_n \,\rangle$ be an element of Tr_{Act}. Then $\delta(s) = d_1 \cdots d_i \cdots d_n$ and $\lambda(s) = \alpha_1 \cdots \alpha_i \cdots \alpha_n$, and $\delta(s,i) = d_i$ and $\lambda(s,i) = \alpha_i$. The length of s is denoted $|s|$; here $|s| = n$. We also use ε to represent the empty probabilistic trace.

We now turn to defining the probability with which a process exhibits a given probabilistic trace. We begin by introducing probabilistic analogues of *weak transitions*. Function $W_{\mathcal{P},\varepsilon} : P \times P \times D_{Act} \to [0,1]$ is intended to reflect the probability with which, beginning in one state, process $\mathcal{P}$ may silently evolve via some number of τ-transitions to another state, given a particular probability distribution describing the actions the environment is enabling. To define this function, we first give a function $W_{\mathcal{P}} : P \times N \times P \times D_{Act} \to [0,1]$, which describes the probability that $\mathcal{P}$ evolves from one state to another using the specified number of τ-transitions.

$$W_{\mathcal{P}}(p,0,p',d) \;=\; \begin{cases} 1 & \text{if } p = p' \\ 0 & \text{otherwise} \end{cases}$$

$$W_{\mathcal{P}}(p,i+1,p',d) = \begin{cases} \frac{1}{\nu(p,d)} \cdot \sum_{p'' \in P} \mu_P(p,\tau,p'') \cdot W_{\mathcal{P}}(p'',i,p',d) \\ \qquad \text{if } \nu(p,d) \neq 0 \\ 0 \quad \text{otherwise} \end{cases}$$

where $\nu(p,d) = (\sum_{\alpha \in Act} \mu_P(p,\alpha) \cdot d(\alpha)) + \mu_P(p,\tau)$.

$W_{\mathcal{P},\varepsilon}(p,p',d) = \sum_{i=0}^{\infty} W_{\mathcal{P}}(p,i,p',d)$

Note that the fact that $\mathcal{P}$ is divergence-free ensures that $W_{\mathcal{P},\varepsilon}(p,p',d)$ is well-defined.

We now introduce the analogue of a weak visible transition. Function $N_{\mathcal{P}} : P \times Act \times P \times D_{Act}$ represents the probability that in one state, process $\mathcal{P}$ may perform some internal transitions, and then the given action, to arrive at another state, given a probability distribution describing the actions enabled by the environment.

$$N_{\mathcal{P}}(p,\alpha,p',d) = \sum_{\{q \mid \nu(q,d)>0\}} W_{\mathcal{P},\varepsilon}(p,q,d) \cdot \frac{d(\alpha)}{\nu(q,d)} \cdot \mu_P(q,\alpha,p')$$

We define a function $M_{\mathcal{P}}$ for process $\mathcal{P}$ that, given a state p and a probabilistic trace s, returns the probability of the process "accepting" s starting at p.

Definition 9. Function $M_{\mathcal{P}} : P \times Tr_{Act} \to [0,1]$ may be defined as follows.

$$M_{\mathcal{P}}(p,\varepsilon) = 1$$
$$M_{\mathcal{P}}(p,\langle\, d,\alpha \,\rangle s') = \sum_{p' \in P} N_{\mathcal{P}}(p,\alpha,p',d) \cdot M_{\mathcal{P}}(p',s')$$

The alternative characterization for τ-free tests, written $\ll_0$, is now given in terms of $M_{\mathcal{P}}$.

Definition 10. $\mathcal{P} \ll_0 \mathcal{Q}$ if for all $s \in Tr_{Act}$ $M_{\mathcal{P}}(p_0,s) \leq M_{\mathcal{Q}}(q_0,s)$.

3.2 Full Abstraction

The rest of this section is devoted to showing that $\ll_0$ coincides with $\sqsubseteq_0$, in other words, that $\ll_0$ is *fully abstract* with respect to $\sqsubseteq_0$. We break this proof into two pieces. We first establish that $\ll_0$ is an *abstraction* of $\sqsubseteq_0$; then we show that $\ll_0$ is *adequate* for $\sqsubseteq_0$.

3.2.1 Abstraction

The theorem we wish to establish in this section is the following.

Theorem 1 (Abstraction). *$\mathcal{P} \sqsubseteq_0 \mathcal{Q}$ implies $\mathcal{P} \ll_0 \mathcal{Q}$.*

To prove this result, we show how to construct a probabilistic test from a probabilistic trace with the property that the probability with which a process passes the test is the same as the probability with which it executes the sequence. Given a sequence s, the definition of this test $\mathcal{T}(s)$ is the following.

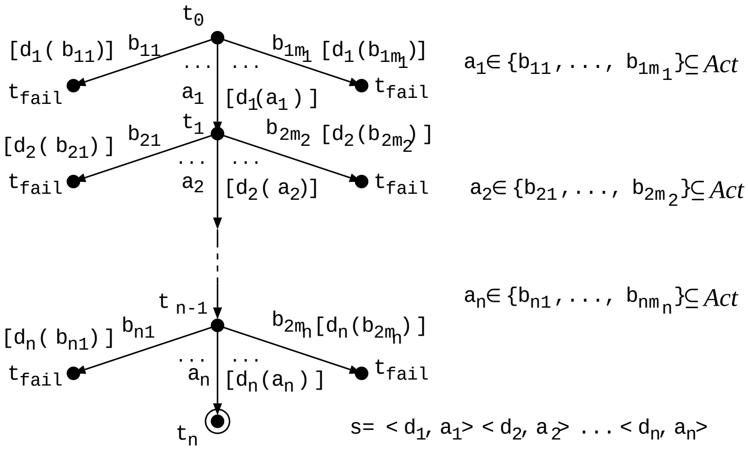

Fig. 1. Structure of $\mathcal{T}(s)$

Definition 11. For $s \in Tr_{Act}$, the test $\mathcal{T}(s)$ is defined by $(T, Act, \mu_s, t_0, \{t_{|s|}\})$ where $T = \{t_i | 0 \leq i \leq |s|\} \cup \{t_{fail}\}$ and

$$\mu_s(t, \alpha, t') = \begin{cases} \delta(s,i)(\alpha) & \text{if } t = t_{i-1},\ t' = t_i \text{ and } \alpha = \lambda(s,i) \\ \delta(s,i)(\alpha) & \text{if } t = t_{i-1},\ t' = t_{fail} \text{ and } \alpha \neq \lambda(s,i) \\ 0 & \text{otherwise} \end{cases}$$

Intuitively, $\mathcal{T}(s)$ has a "spine" labeled by the actions of s which is the unique path leading to the only successful state. The paths that are not on the spine lead to the failure state. Figure 1 depicts a typical $\mathcal{T}(s)$. We are now able to prove the following.

Lemma 1. *Let* $\mathcal{P} = (P, Act, \mu_P, p_0)$*. Then for all* $s \in Tr_{Act}$*,* $Pr(\mathcal{S}_{\mathcal{P}\|\mathcal{T}(s)}) = M_{\mathcal{P}}(p_0, s)$.

Proof. We prove, by induction on the length of s, a stronger result; namely, that for all $p \in P$, $Pr(\mathcal{S}_{p\|\mathcal{T}(s)}) = M_{\mathcal{P}}(p, s)$.

Given this lemma, the proof of Theorem 1 follows in a straightforward manner. Let s be a probabilistic trace; we want to show that $M_{\mathcal{P}}(p_0, s) \leq M_{\mathcal{Q}}(q_0, s)$. This follows from the fact that since $\mathcal{P} \sqsubseteq_0 \mathcal{Q}$, $Pr(\mathcal{S}_{\mathcal{P}\|\mathcal{T}(s)}) \leq Pr(\mathcal{S}_{\mathcal{Q}\|\mathcal{T}(s)})$. Then Lemma 1 gives the desired result.

3.2.2 Adequacy

We now wish to establish that $\ll_0$ is adequate for $\sqsubseteq_0$; that is, our task is to prove the following theorem.

Theorem 2 (Adequacy). $\mathcal{P} \ll_0 \mathcal{Q}$ *implies* $\mathcal{P} \sqsubseteq_0 \mathcal{Q}$.

To prove this result, we in essence argue that in order to understand how a process interacts with tests, it suffices to consider tests of a very restricted form—namely, tests that can be constructed from probabilistic traces. We refer to these tests as *essential* tests, since they expose all relevant aspects of process behavior.

In order to see why essential tests suffice to characterize processes with respect to $\sqsubseteq_0$, consider the factors that determine the probability with which a process passes an arbitrary test. In general, such a test has some number of "successful transition sequences" that lead from the start state of the test to an accepting state and do not pass through an accepting state along the way. When a process interacts with such a test, each successful computation exercises exactly one of these successful sequences; by summing up the probabilities of these computations, we may speak of the probability with which a process exercises a given successful sequence. Then the sum of the probabilities over all such sequences yields the probability with which the process passes the test. Thus, if we can construct an appropriate essential test for each such successful path, we may characterize the behavior of a process with respect to the test by looking only at the relevant essential tests.

To formalize these intuitions, let $\mathcal{P} = (P, Act, \mu_P, p_0)$ be a process and $\mathcal{T} = (T, Act, \mu_T, t_0, S)$ be a test, and let F be of the form:

$$\langle p, t \rangle \xrightarrow{\alpha_1,\, \pi_1} \langle p_1, t_1 \rangle \xrightarrow{\alpha_2,\, \pi_2} \cdots \xrightarrow{\alpha_m,\, \pi_m} \langle p_m, t_m \rangle$$

We use $F\lceil_{\mathcal{T}}$ to denote the transition sequence in $\mathcal{T}$ that interacts with a transition sequence in $\mathcal{P}$ to produce F. Thus, $F\lceil_{\mathcal{T}}$ is of the form:

$$t \xrightarrow{\alpha_1,\, \rho_1} t_1 \xrightarrow{\alpha_2,\, \rho_2} \cdots \xrightarrow{\alpha_m,\, \rho_m} t_m$$

Suppose that $F_{\mathcal{T}}$ is a transition sequence in $\mathcal{T}$ that starts from state t. Then we define $\mathcal{C}_{\langle p,t \rangle}\lceil_{F_{\mathcal{T}}} = \{F \mid F\lceil_{\mathcal{T}} = F_{\mathcal{T}} \text{ and } F \preceq C \text{ for some } C \in \mathcal{C}_{\langle p,t \rangle}\}$. Intuitively, $\mathcal{C}_{\langle p,t \rangle}\lceil_{F_{\mathcal{T}}}$ is the set of computation prefixes from $\langle p, t \rangle$ that exercise the transition sequence $F_{\mathcal{T}}$ in $\mathcal{T}$.

Now suppose that $F_{\mathcal{T}}$ is the transition sequence $t \xrightarrow{\alpha_1,\, \rho_1} t_1 \cdots \xrightarrow{\alpha_m,\, \rho_m} t_m$ in test $\mathcal{T}$, and let $F_{\mathcal{T}}(i)$ denote the ith state of $F_{\mathcal{T}}$, i.e. $F_{\mathcal{T}}(0) = t$ and $F_{\mathcal{T}}(i) = t_i$ for $i > 0$. We may define a probabilistic trace $tr(F_{\mathcal{T}}) = \langle d_1, \alpha_1 \rangle \cdots \langle d_m, \alpha_m \rangle$, where $d_i = \mu_T(F_{\mathcal{T}}(i-1))$ is the probability distribution on actions induced by $\mathcal{T}$'s transition relation at state $F_{\mathcal{T}}(i)$. Note that the fact that $\mathcal{T}$ is finite-branching ensures that each d_i returns 0 for all but a finite number of actions.

In general, because of potential nondeterminism in $\mathcal{T}$, $d_i(\alpha_i)$ will exceed $\mu_T(t_{i-1}, \alpha, t_i)$, and hence the probabilistic trace does not accurately reflect the probabilities in the transition sequence. However, the next lemma shows a strong relationship between the two.

Lemma 2. *Let $\mathcal{P} = (P, Act, \mu_P, p_0)$ be a process and $\mathcal{T} = (T, Act, \mu_T, t_0, S)$ be a test, and let $F_{\mathcal{T}}$ be a transition sequence in $\mathcal{T}$ starting from t_0. Then $Pr(\mathcal{C}_{\mathcal{P}\|\mathcal{T}}\lceil_{F_{\mathcal{T}}}) = v \cdot M_{\mathcal{P}}(p_0, s)$ where $s = tr(F_{\mathcal{T}})$ and $v = \prod_{1 \le i \le |F_{\mathcal{T}}|} \frac{\mu_T(t_{i-1}, \alpha_i, t_i)}{\mu_T(t_{i-1}, \alpha_i)}$ such that $t_i = F_{\mathcal{T}}(i)$ with $\alpha_i = \lambda(s, i)$.*

Proof. We prove a stronger fact; namely $Pr(\mathcal{C}_{\langle p,t \rangle}\lceil_{F_{\mathcal{T}}}) = v \cdot M_{\mathcal{P}}(p, s)$ for all $p \in P$ and $u_0 \in T$, where $v = \prod_{1 \le i \le |F_{\mathcal{T}}|} \frac{\mu_T(u_{i-1}, \beta_i, u_i)}{\mu_T(u_{i-1}, \beta_i)}$ with $u_i = F_{\mathcal{T}}(i)$ and $\beta_i = \lambda(s, i)$. This follows by induction on the length of $F_{\mathcal{T}}$.

In the statement of this lemma, note that v and s only depend on $F_{\mathcal{T}}$. Thus, once $F_{\mathcal{T}}$ is fixed then so are v and s. This fact and Lemma 1 enable us now to prove the following.

Lemma 3. *Let $F_{\mathcal{T}}$ be a non-empty transition sequence in $\mathcal{T}$ starting from t_0. Then $\mathcal{P} \ll_0 \mathcal{Q}$ implies $Pr(\mathcal{C}_{\mathcal{P}\|\mathcal{T}}\lceil_{F_{\mathcal{T}}}) \leq Pr(\mathcal{C}_{\mathcal{Q}\|\mathcal{T}}\lceil_{F_{\mathcal{T}}})$.*

The proof of the adequacy theorem now follows from the following fact. Fix a test $\mathcal{T}$, and let $\mathcal{T}_S$ be the set of successful transition sequences in $\mathcal{T}$. Then $Pr(\mathcal{S}_{\mathcal{P}\|\mathcal{T}}) = \sum_{F_{\mathcal{T}}\in\mathcal{T}_S} Pr(\mathcal{C}_{\mathcal{P}\|\mathcal{T}}\lceil_{F_{\mathcal{T}}})$ and $Pr(\mathcal{S}_{\mathcal{Q}\|\mathcal{T}}) = \sum_{F_{\mathcal{T}}\in\mathcal{T}_S} Pr(\mathcal{C}_{\mathcal{Q}\|\mathcal{T}}\lceil_{F_{\mathcal{T}}})$. Since $\mathcal{P} \ll_0 \mathcal{Q}$, Lemma 3 ensures that for each $F_{\mathcal{T}} \in \mathcal{T}_S$, $Pr(\mathcal{C}_{\mathcal{P}\|\mathcal{T}}\lceil_{F_{\mathcal{T}}}) \leq Pr(\mathcal{C}_{\mathcal{Q}\|\mathcal{T}}\lceil_{F_{\mathcal{T}}})$, and it therefore follows that $Pr(\mathcal{S}_{\mathcal{P}\|\mathcal{T}}) \leq Pr(\mathcal{S}_{\mathcal{Q}\|\mathcal{T}})$.

3.3 A Proof Technique

We close this section by giving a result that eases the task of showing $\mathcal{P} \ll_0 \mathcal{Q}$. Let $traces(\mathcal{P})$ be the set of the sequences of the observable actions (without probabilities) that $\mathcal{P}$ can execute from its start state. The following is immediate.

Lemma 4. *If $\lambda(s) \notin traces(\mathcal{P})$, then $M_{\mathcal{P}}(p_0, s) = 0$*

It also follows that if $\mathcal{P} \ll_0 \mathcal{Q}$ then any trace of $\mathcal{P}$ is a trace of $\mathcal{Q}$.

Lemma 5. *$\mathcal{P} \ll_0 \mathcal{Q}$ implies $traces(\mathcal{P}) \subseteq traces(\mathcal{Q})$*

Now suppose that $S \subseteq Tr_{Act}$. We write $\mathcal{P} \ll_0^S \mathcal{Q}$ if $M_{\mathcal{P}}(p_0, s) \leq M_{\mathcal{Q}}(q_0, s)$ for all $s \in S$. The next theorem shows that to check whether $\mathcal{P} \ll_0 \mathcal{Q}$ holds, we can limit the probabilistic traces that need checking as follows.

Theorem 3. *Let $S_{\mathcal{P}} = \{s \in Tr_{Act} \mid \lambda(s) \in traces(\mathcal{P})\}$. Then $\mathcal{P} \ll_0 \mathcal{Q}$ if and only if $\mathcal{P} \ll_0^{S_{\mathcal{P}}} \mathcal{Q}$ and $traces(\mathcal{P}) \subseteq traces(\mathcal{Q})$.*

On the basis of this theorem and the full abstractness result, then, we have the following proof technique for showing that $\mathcal{P} \sqsubseteq_0 \mathcal{Q}$.

- Check if $traces(\mathcal{P}) \subseteq traces(\mathcal{Q})$.
- Check if $M_{\mathcal{P}}(p_0, s) \leq M_{\mathcal{Q}}(q_0, s)$ for $s \in Tr$ such that $\lambda(s) \in traces(\mathcal{Q})$.

Example 1. We prove that $\mathcal{P}_1 \sqsubseteq_0 \mathcal{Q}_1$, where $\mathcal{P}_1$ and $\mathcal{Q}_1$ are depicted in the following.

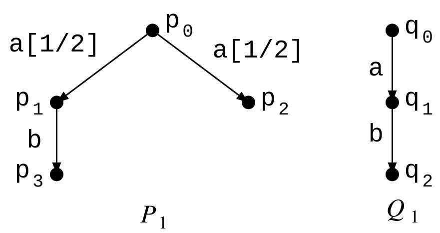

Since $traces(\mathcal{P}_1) = traces(\mathcal{Q}_1) = \{\epsilon, a, ab\}$, we must check the probabilistic traces $\langle d_1, a\rangle$ and $\langle d_1, a\rangle\langle d_2, b\rangle$.

(Case 1): $s = \langle d_1, a\rangle$. Then $M_{\mathcal{P}_1}(p_0, s) = M_{\mathcal{Q}_1}(q_0, s) = 1$

(Case 2): $s = \langle d_1, a\rangle\langle d_2, b\rangle$. Then $M_{\mathcal{P}_1}(p_0, s) = \frac{1}{2} \leq 1 = M_{\mathcal{Q}_1}(q_0, s)$

Example 2. We establish that $\mathcal{P}_2 \sqsubseteq_0 \mathcal{Q}_2$.

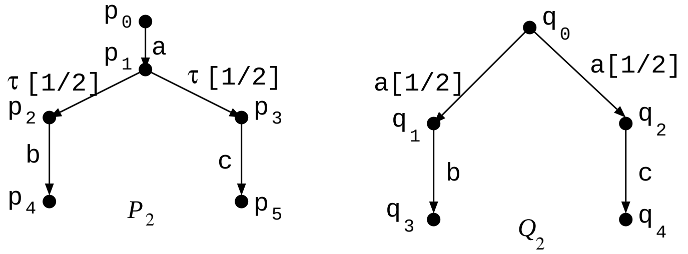

Since $traces(\mathcal{P}_2) = traces(\mathcal{Q}_2) = \{\varepsilon, a, ab, ac\}$, we check the following cases.

(Case 1): $s = \langle d_1, a\rangle$. $M_{\mathcal{P}_2}(p_0, \langle d_1, a\rangle) = M_{\mathcal{Q}_2}(q_0, \langle d_1, a\rangle) = 1$

(Case 2): $s = \langle d_1, a\rangle\langle d_2, b\rangle$.

$$M_{\mathcal{P}_2}(p_0, \langle d_1, a\rangle\langle d_2, b\rangle) = M_{\mathcal{Q}_2}(q_0, \langle d_1, a\rangle\langle d_2, b\rangle) = \frac{1}{2}$$

(Case 3): $s = \langle d_1, a\rangle\langle d_2, c\rangle$. Similar to Case 2.

We can show $\mathcal{Q}_2 \sqsubseteq_0 \mathcal{P}_2$ in a similar way. Consequently, we have that $\mathcal{P}_2$ and $\mathcal{Q}_2$ are equivalent.

4 Alternative Characterization for General Tests

We next investigate $\sqsubseteq^{\mathbf{T}}$, the testing preorder induced by the general tests, which may include transitions labeled by τ. (In the sequel, we simply write $\sqsubseteq^{\mathbf{T}}$ as $\sqsubseteq$.) As the set of tests is now larger, one might expect this new preorder to be strictly finer than $\sqsubseteq_0$, and this is indeed the case. For instance, in Example 1, the processes are related with respect to τ-free tests, but they are not related if τ-transitions are possible in tests. Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be the following tests, where $t_{1,2}$ and $t_{2,2}$ are the successful states.

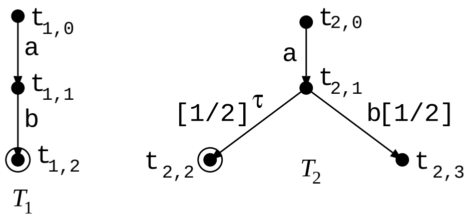

$Pr(\mathcal{S}_{\mathcal{P}_1\|\mathcal{T}_1}) = \frac{1}{2}$ and $Pr(\mathcal{S}_{\mathcal{Q}_1\|\mathcal{T}_1}) = 1$, but $Pr(\mathcal{S}_{\mathcal{P}_1\|\mathcal{T}_2}) = \frac{3}{4}$ and $Pr(\mathcal{S}_{\mathcal{Q}_1\|\mathcal{T}_2}) = \frac{1}{2}$, showing these processes are not related.

4.1 The Alternative Characterization

In order to define the alternative characterization of $\sqsubseteq$ we first introduce τ's into probabilistic traces.

Definition 12 (Probabilistic Trace with τ). A sequence $s \in (D_{Act_\tau} \times Act_\tau)^*$ is called a *probabilistic trace* of Act_τ where in each $\langle d, \alpha\rangle$ the set $\{\beta \mid d(\beta) > 0\}$ is finite. The set of all probabilistic traces of Act_τ is denoted by Tr_{Act_τ}.

Again we define a function to give the probability with which a process exhibits a given trace. Since probabilistic traces may include τ's, we must modify the functions defined in the previous section.

As before, the function $W_{\tau,\mathcal{P},\varepsilon}(p, p', d)$ is designed to compute the probability that process $\mathcal{P}$ may evolve from p to p' via some sequence of τ's in the context of an environment that enables actions as specified by d. The definition is given in terms of $W_{\tau,\mathcal{P}}(p, i, p', d)$ which records the probability of an evolution from p to p' using exactly i τ's.

$$W_{\tau,\mathcal{P}}(p, 0, p', d) = \begin{cases} 1 & \text{if } p = p' \\ 0 & \text{otherwise} \end{cases}$$

$$W_{\tau,\mathcal{P}}(p, i+1, p', d)$$

$$= \begin{cases} \frac{1}{\nu(p,d)} \cdot \sum_{p'' \in P} (1 - d(\tau)) \cdot \mu_P(p, \tau, p'') \cdot W_{\tau,\mathcal{P}}(p'', i, p', d) \\ \qquad \text{if } \nu(p, d) \neq 0 \\ 0 \qquad\quad \text{otherwise} \end{cases}$$

where $\nu(p, d) = (\sum_{\alpha \in Act} \mu_P(p, \alpha) \cdot d(\alpha)) + \mu_P(p, \tau) + d(\tau) - \mu_P(p, \tau) \cdot d(\tau)$.

$$W_{\tau,\mathcal{P},\varepsilon}(p, p', d) = \sum_{i=0}^{\infty} W_{\tau,\mathcal{P}}(p, i, p', d)$$

Note that this definition changes the normalization factor $\nu(p, d)$ to take into account the probability that the environment enables τ's and that d is not allowed to "execute" its τ when determining $W_{\tau,\mathcal{P}}(p, i, p', d)$. Also in this case, $W_{\tau,\mathcal{P}}(p, p', d)$ is well-defined since $\mathcal{P}$ is divergence-free.

We now define $N_{\tau,\mathcal{P}}(p, \alpha, p', d)$, which computes the probability that $\mathcal{P}$ evolves from p to p' using internal computation followed by α, given d. When $\alpha = \tau$, we must account for two possibilities, one in which the environment executes its τ and one in which it does not.

$$N_{\tau,\mathcal{P}}(p, \alpha, p', d)$$

$$= \begin{cases} \sum_{\{q \mid \nu(q,d) > 0\}} W_{\tau,\mathcal{P},\varepsilon}(p, q, d) \cdot \frac{d(\tau)}{\nu(q,d)} \cdot \mu_P(q, \tau, p') \\ \quad + W_{\tau,\mathcal{P},\varepsilon}(p, p', d) \cdot \frac{d(\tau)}{\nu(p',d)} \cdot (1 - \mu_P(p', \tau)) & \text{if } \alpha = \tau \text{ and } \nu(p', d) \neq 0 \\ \sum_{\{q \mid \nu(q,d) > 0\}} W_{\tau,\mathcal{P},\varepsilon}(p, q, d) \cdot \frac{d(\alpha)}{\nu(q,d)} \cdot \mu_P(q, \alpha, p') & \text{otherwise} \end{cases}$$

Function $M_{\tau,\mathcal{P}}$ computes the probability that from the given state process $\mathcal{P}$ completes the given trace.

Definition 13. Function $M_{\tau,\mathcal{P}} : P \times Tr_{Act_\tau} \to [0, 1]$ is defined as follows.

$$M_{\tau,\mathcal{P}}(p, \varepsilon) = 1$$
$$M_{\tau,\mathcal{P}}(p, \langle d, \alpha \rangle s') = \sum_{p' \in P} N_{\tau,\mathcal{P}}(p, \alpha, p', d) \cdot M_{\tau,\mathcal{P}}(p', s')$$

The alternative characterization for general tests is given in terms of $M_{\tau,\mathcal{P}}$.

Definition 14. $\mathcal{P} \ll \mathcal{Q}$ if for all $s \in Tr_{Act_\tau}$ $M_{\tau,\mathcal{P}}(p_0, s) \leq M_{\tau,\mathcal{Q}}(q_0, s)$

4.2 Full Abstraction

In this section we establish that $\ll$ is equivalent to $\sqsubseteq$. To do this, we first show that $\ll$ is an abstraction of $\sqsubseteq$, and then we show that $\ll$ is adequate for $\sqsubseteq$.

4.2.1 Abstraction

Theorem 4 (Abstraction). *$\mathcal{P} \sqsubseteq \mathcal{Q}$ implies $\mathcal{P} \ll \mathcal{Q}$.*

The construction of probabilistic tests from probabilistic traces is the same as in Section 3, the only difference being that the action alphabet is Act_τ. Namely, $\mathcal{T}(s) = (T, Act_\tau, \mu_s, t_0, \{t_{|s|}\})$ for $s \in Tr_{Act_\tau}$. The proof of Theorem 4 follows directly from the next lemma.

Lemma 6. *Let $\mathcal{P} = (P, Act, \mu_P, p_0)$. Then for all $s \in Tr_{Act_\tau}$, $Pr(\mathcal{S}_{\mathcal{P}\|\mathcal{T}(s)}) = M_{\tau,\mathcal{P}}(p_0, s)$.*

4.2.2 Adequacy

Theorem 5 (Adequacy). *$\mathcal{P} \ll \mathcal{Q}$ implies $\mathcal{P} \sqsubseteq \mathcal{Q}$.*

The proof of adequacy is essentially identical to the case of τ-free tests.

Lemma 7. *Let $\mathcal{P} = (P, Act, \mu_P, p_0)$ be a process and $\mathcal{T} = (T, Act, \mu_T, t_0, S)$ be a test, and let $F_\mathcal{T}$ be a transition sequence in $\mathcal{T}$ starting from t_0. Then $Pr(\mathcal{C}_{\mathcal{P}\|\mathcal{T}}\lceil_{F_\mathcal{T}}) = v \cdot M_{\tau,\mathcal{P}}(p_0, s)$ where $s = tr(F_\mathcal{T})$ and $v = \prod_{1 \le i \le |F_\mathcal{T}|} \frac{\mu_T(t_{i-1},\alpha_i,t_i)}{\mu_T(t_{i-1},\alpha_i)}$ such that $t_i = F_\mathcal{T}(i)$ with $\alpha_i = \lambda(s, i)$.*

As before, s and v are dependent only on $F_\mathcal{T}$. Thus the following lemma holds.

Lemma 8. *Let $F_\mathcal{T}$ be a non-empty transition sequence in $\mathcal{T}$ starting from t_0. Then $\mathcal{P} \ll \mathcal{Q}$ implies $Pr(\mathcal{C}_{\mathcal{P}\|\mathcal{T}}\lceil_{F_\mathcal{T}}) \le Pr(\mathcal{C}_{\mathcal{Q}\|\mathcal{T}}\lceil_{F_\mathcal{T}})$.*

Theorem 5 follows from the fact that $Pr(\mathcal{S}_{\mathcal{P}\|\mathcal{T}}) = \sum_{F_\mathcal{T} \in \mathcal{T}_S} Pr(\mathcal{C}_{\mathcal{P}\|\mathcal{T}}\lceil_{F_\mathcal{T}})$, where $\mathcal{T}_S$ is the set of successful transition sequences in $\mathcal{T}$.

4.3 Equivalence Property

Somewhat surprisingly, it turns out that $\sqsubseteq$ is in fact an equivalence relation. To prove this result, we first show that $\sqsubseteq$ is in fact characterized by *non-blocking* probabilistic traces.

Definition 15 (Non-blocking probabilistic trace). Probabilistic trace s is *non-blocking* if for all $i \le |s|$, $d_i(\tau) > 0$.

Definition 16. $\mathcal{P} \ll^{NB} \mathcal{Q}$ if $M_{\tau,\mathcal{P}}(p_0, s) \le M_{\tau,\mathcal{Q}}(q_0, s)$ for all non-blocking s.

We say that a probabilistic trace is "blocking" if it is not non-blocking. We also call an essential test "non-blocking" if it is generated from a non-blocking probabilistic trace, and "blocking" if it is generated from a blocking probabilistic trace.

We show that $\sqsubseteq$ is is actually characterized by $\ll^{NB}$. We first prove the following lemma.

Lemma 9. *$\mathcal{P} \ll^{NB} \mathcal{Q}$ implies $\mathcal{P} \ll \mathcal{Q}$*

Proofsketch. Suppose $M_{\tau,\mathcal{P}}(p_0, s) \leq M_{\tau,\mathcal{Q}}(q_0, s)$ for all non-blocking s and $\mathcal{P} \not\ll \mathcal{Q}$. Then, there exists a blocking s_0 such that $M_{\mathcal{P}}(p_0, s_0) > M_{\mathcal{Q}}(q_0, s_0)$. We can show that there exists a non-blocking s'_0 such that $M_{\tau,\mathcal{P}}(p_0, s'_0) > M_{\tau,\mathcal{Q}}(q_0, s'_0)$, deriving a contradiction.

To build s'_0, we modify any "blocking" distribution d in s_0 to d', where $d'(\tau) > 0$ in such a way that $M_{\tau,\mathcal{P}}(p_0, s'_0) > M_{\tau,\mathcal{Q}}(q_0, s'_0)$. Such d' always exist due to the fact that a process is finite-branching and has a lower bound on its transition probabilities.

That $\mathcal{P} \ll \mathcal{Q}$ implies $\mathcal{P} \ll^{NB} \mathcal{Q}$ is trivial; thus we can conclude that the non-blocking essential tests give the same preorder as do the essential tests.

Corollary 1. *$\mathcal{P} \ll \mathcal{Q}$ if and only if $\mathcal{P} \ll^{NB} \mathcal{Q}$*

An essential feature of non-blocking essential tests is that they can be "inverted"; from a given nonempty non-blocking essential test we may construct a set of non-blocking essential tests with the property that the cumulative probability with which any process passes the derived tests is 1 minus the probability that the process passes the original test. On the basis of the alternative characterization, then, we may conclude that if $\mathcal{P} \sqsubseteq \mathcal{Q}$ holds, then so does $\mathcal{Q} \sqsubseteq \mathcal{P}$. The reason that this inversion is possible is that unsuccessful computations can only end in configurations of the form $\langle p, t_{fail} \rangle$; that is, it is impossible to fail by getting "stuck" on the spine.

Definition 17 (Inverse Set of Non-blocking Probabilistic Traces). For a nonempty non-blocking s, a set of non-blocking probabilistic traces, denoted by s^{-1}, is inductively constructed as follows:

$$\langle d, \alpha \rangle^{-1} = \{\langle d, \beta \rangle \mid \beta \neq \alpha, d(\beta) > 0\}$$
$$(\langle d, \alpha \rangle s)^{-1} = \langle d, \alpha \rangle^{-1} \cup \{\langle d, \alpha \rangle s' \mid s' \in s^{-1}\}$$

The basic idea for the inverse set is to flip the failure state and the success state in an essential test. By simply flipping the states, the test has many successful paths. Therefore, we decompose the test again into a set of essential tests. The property is formally stated as follows.

Lemma 10. *If s is non-blocking, $M_{\tau,\mathcal{P}}(p_0, s) + \sum_{s' \in s^{-1}} M_{\tau,\mathcal{P}}(p_0, s') = 1$.*

Proof. We prove, by induction on the length of s, the slightly stronger fact: for a non-blocking probabilistic trace s, $M_{\tau,\mathcal{P}}(p, s) + \sum_{s' \in s^{-1}} M_{\tau,\mathcal{P}}(p, s') = 1$ for all $\mathcal{P}$ and for all $p \in P$.

Due to the existence of the inverse set, $\ll^{NB}$ is shown to be commutative. Namely, it forms an equivalence relation.

Lemma 11. *$\mathcal{P} \ll^{NB} \mathcal{Q}$ implies $\mathcal{Q} \ll^{NB} \mathcal{P}$*

Proof. By the definition of a non-blocking s, $M_{\tau,\mathcal{P}}(p_0, s) \leq M_{\tau,\mathcal{Q}}(q_0, s)$. And for all $u \in s^{-1}$, $M_{\tau,\mathcal{P}}(p_0, u) \leq M_{\tau,\mathcal{Q}}(q_0, u)$ since u is non-blocking. Thus

$$\begin{aligned} \textstyle\sum_{u \in s^{-1}} M_{\tau,\mathcal{P}}(p_0, u) &\leq \textstyle\sum_{u \in s^{-1}} M_{\tau,\mathcal{Q}}(q_0, u) \\ 1 - M_{\tau,\mathcal{P}}(p_0, s) &\leq 1 - M_{\tau,\mathcal{Q}}(q_0, s) \\ M_{\tau,\mathcal{Q}}(q_0, s) &\leq M_{\tau,\mathcal{P}}(p_0, s) \end{aligned}$$

The identity of $\ll$ and $\ll^{NB}$ and the full abstractness of $\ll$ with respect to $\sqsubseteq$ lead to the equivalence property of the preorder.

Theorem 6. *$\mathcal{P} \sqsubseteq \mathcal{Q}$ implies $\mathcal{Q} \sqsubseteq \mathcal{P}$*

4.4 A Proof Technique

We now provide a proof technique that is similar to the technique for τ-free tests. We first show that the traces of processes are identical if they are related by the testing preorder.

Lemma 12. *$\mathcal{P} \ll \mathcal{Q}$ implies $traces(\mathcal{P}) = traces(\mathcal{Q})$*

This follows from the fact that $\ll$ is an equivalence relation and $\ll$ implies $\ll_0$. It also turns out that if the "visible content" of a probabilistic trace is not a (nonprobabilistic) trace of a process, then the process is incapable of executing the probabilistic trace. Formally, if $\sigma \in Act_\tau^*$, then let $\widehat{\sigma}$ be σ with all τ's removed. We now have the following.

Lemma 13. *If $\widehat{\lambda(s)} \notin traces(\mathcal{P})$, then $M_{\tau,\mathcal{P}}(p_0, s) = 0$*

Theorem 7. *Let $S = \{s \mid \widehat{\lambda(s)} \in traces(\mathcal{P})\}$. Then, $\mathcal{P} \ll \mathcal{Q}$ if and only if $traces(\mathcal{P}) = traces(\mathcal{Q})$ and $\mathcal{P} \ll^S \mathcal{Q}$.*

This theorem and the full abstraction result imply the following proof technique for $\mathcal{P} \sqsubseteq \mathcal{Q}$.

- Check if $traces(\mathcal{P}) = traces(\mathcal{Q})$
- Check if $M_\mathcal{P}(p_0, s) = M_\mathcal{Q}(q_0, s)$ for all s such that $\widehat{\lambda(s)} \in traces(\mathcal{P})$.

Example 3. We use the proof technique to show that $\mathcal{P}_3 \sqsubseteq \mathcal{Q}_3$ where $\mathcal{P}_3$ and $\mathcal{Q}_3$ are given by:

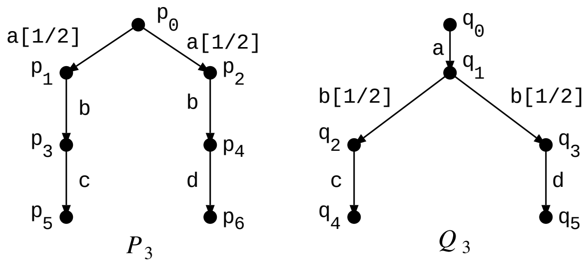

Since $traces(\mathcal{P}_3) = traces(\mathcal{Q}_3) = \{\varepsilon, a, ab, abc, abd\}$, we only need consider s such that $\widehat{\lambda(s)} \in \{\varepsilon, a, ab, abc, abd\}$. The details are left to the reader.

5 Concluding Remarks and Future Work

We have presented alternative characterizations of the testing preorders for probabilistic processes put forth in [CSZ92]. Specifically, we have investigated two preorders, distinguished on the basis of whether or not tests are allowed to make τ-transitions. In each case, we have proven the characterization to be fully abstract with respect to the corresponding testing preorder and given a proof technique that eases the task of establishing the preorder. Moreover, we have shown the second preorder to be an equivalence relation, a fact that went unnoticed in [CSZ92].

Our results are also applicable to *substochastic processes* [CSZ92], where the sum of the probabilities of outgoing transitions from a state may be strictly less than one. The deficit represents the process's capacity for *undefined behavior*.

Our present characterizations rely on succinctly encoding essential tests as probabilistic traces. As future work, we seek to establish more denotational characterizations based entirely on the structure of probabilistic processes. The acceptance tree model of [Hen88] for nondeterministic processes offers an appealing starting point for such an investigation.

We would also like to extend our testing preorder to handle time as well as probability. This would enable, for example, reasoning about "soft" real-time systems in which deadlines may be missed but not "too often."

Acknowledgement. The first author would like thank Prof. Yasuyoshi Inagaki of Nagoya University who provided support for his stay at SUNY Stony Brook.

References

[Chr90] I. Christoff. Testing equivalences and fully abstract models for probabilistic processes. In J. C. M. Baeten and J. W. Klop, editors, *Proceedings of CONCUR '90 – First Intl. Conf. on Concurrency Theory,* volume 458 of *Lecture Notes in Computer Science*, pages 126–140. Springer-Verlag, 1990.

[CSZ92] R. Cleaveland, S. A. Smolka, and A. E. Zwarico. Testing preorders for probabilistic processes. In *Proceedings of the 19th ICALP*, volume 623 of *Lecture Notes in Computer Science*, pages 708–719. Springer Verlag, July 1992.

[DeN87] R. DeNicola. Extensional equivalences for transition systems. *Acta Informatica*, 24:211–237, 1987.

[DH83] R. DeNicola and M.C.B. Hennessy. Testing equivalences for processes. *Theoretical Computer Science*, 34:83–133, 1983.

[Hen88] M.C.B. Hennessy. *Algebraic Theory of Processes.* MIT Press, 1988.

[HJ90] H. A. Hansson and B. Jonsson. A calculus for communicating systems with time and probabilities. In *Proceedings of the 11th IEEE Symposium on Real-Time Systems*, pages 278–287, 1990.

[LS92a] K. G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94(1):1–28, September 1992. Preliminary versions of this paper appeared as University of Aalborg technical reports R 88-18 and R 88-29, and in *Proceedings of the 16th Annual ACM Symposium on Principles of Programming Languages*, Austin, Texas, 1989.

[LS92b] K. G. Larsen and A. Skou. Compositional verification of probabilistic processes. In *Proceedings of CONCUR '92*, volume 630 of *Lecture Notes in Computer Science*, pages 456–471. Springer-Verlag, 1992.

[LT87] N.A. Lynch and M. Tuttle. Hierarchical correctness proofs for distributed algorithms. In *Proceeding the 6th Annual ACM symposium on Principles of Distributed Computing*, pages 137–151, 1987.

[Mil89] R. Milner. *Communication and Concurrency.* Prentice Hall, 1989.

[vGSST90] R. J. van Glabbeek, S. A. Smolka, B. Steffen, and C. M. N. Tofts. Reactive, generative, and stratified models of probabilistic processes. In *Proceedings of the 5th IEEE Symposium on Logic in Computer Science*, pages 130–141, Philadelphia, PA, 1990.

[WSS94] S. Wu, S.A. Smolka, and E.W. Stark. Composition and behaviors of probabilistic I/O automata. In *Proceedings of CONCUR '94*, Lecture Notes in Computer Science. Springer Verlag, 1994.

[YL92] W. Yi and Kim G. Larsen. Testing probabilistic and nondeterministic processes. In *Protocol Specification, Testing and Verification XII*, pages 47–61, 1992.