Infinite Probabilistic and Nonprobabilistic Testing

K Narayan Kumar[1]* Rance Cleaveland[2] Scott A. Smolka[1]

[1] Dept. Comput. Sci., SUNY at Stony Brook, Stony Brook NY 11794
(Research supported in part by NSF Grant CCR-9505562 and
AFOSR Grants F49620-93-1-0250 and F49620-95-1-0508)
{kumar,sas}@cs.sunysb.edu
[2] Dept. Comput. Sci., N.C. State Univ., Raleigh, NC 27695
(Research supported in part by AFOSR Grant F49620-95-1-0508)
rance@csc.ncsu.edu

Abstract. We introduce several new notions of infinite testing, for both probabilistic and nonprobabilistic processes, and carefully examine their distinguishing power. We consider three kinds of infinite tests, simple, Büchi and fair. We show that Büchi tests do not change the distinguishing power of nondeterministic testing. In the probabilistic setting, we show that all three have the same distinguishing power as finite tests. Finally, we show that finite probabilistic tests are stronger than nondeterministic fair tests.

1 Introduction

Testing preorders [DNH83, Hen88] have proved to be a very versatile paradigm for relating specifications and their implementations. The basic idea is first to define when a process *may pass* a test (there exists a successful interaction between the process and test) and *must pass* a test (all process-test interactions are successful). The *may* preorder is then defined as follows: a process P is *may* less than a process Q if whenever P *may* pass a test T, then Q *may* passes T as well. The *must* preorder is defined similarly. The *testing* preorder is then the conjunction of the *may* and *must* preorders.

These preorders extend in a natural and useful way to *probabilistic processes*; see, for example, [CDS$^+$98, WSS97, YL92, Seg96, JY95]. A probabilistic process can be viewed as a *labeled Markov chain*, where the label attached to a Markovian transition indicates the communication action that takes place in conjunction with the transition. Testing preorders for probabilistic processes can be thought of as relating processes in terms of their relative reliability in different operating environments. Such environments can be modeled by *tests*, i.e. probabilistic processes equipped with a set of *success* states.

* On leave from SPIC Mathematical Institute, Chennai, India.

The majority of previous research on testing preorders has focused on finite testing. In the case of nondeterministic processes (which we also sometimes refer to as "nonprobabilistic processes") there are, however, two notable exceptions. Hennessey [Hen88] has shown that a certain simple extension of finite testing to infinite testing offers no additional distinguishing power in terms of the induced preorders. Natarajan and Cleaveland [NC95] and Brinksma et al. [BRV95] have shown that *fair testing* induces a finer testing preorder than finite testing, and has practical ramifications in applications such as protocol verification.

In this paper, we introduce several new notions of infinite testing, for both probabilistic and nonprobabilistic processes, and carefully examine their distinguishing power. In particular, we consider the testing preorders induced by the following success criteria for infinite tests: *simple infinite tests*: a finite execution ending in a success state is deemed successful while all infinite executions are treated as unsuccessful; *Büchi tests*: tests are equipped with a set of special states called Büchi states, and an infinite execution is considered successful if it visits the set of Büchi state infinitely often; and *fair testing*: a state of a test from which a success state can be reached is said to be "potentially successful," and an infinite execution is successful if it remains among the set of potentially successful states.

Simple tests and fair tests have been studied in literature and a motivation for their consideration may be found there. The motivation for Büchi tests is the following. A Büchi test corresponds to an environment that acts as a monitor, continously interacting with the system and periodically indicating whether the interaction up to that point is admissible. The entire computation is successful if the monitor reports success infinitely often.

As alluded to above, it is already known that, for nondeterministic processes, simple infinite tests offer no additional distinguishing power than finite tests [Hen88] (i.e. the induced preorders coincide), and that the fair testing preorder is finer than the one induced by finite tests [NC95, BRV95]. We complete this picture by showing that Büchi testing also coincides with finite testing.

In the probabilistic setting, only finite testing has been considered previously [CDS$^+$98]. We show that infinite tests offer no more distinguishing power than finite tests; i.e., the simple, Büchi, and fair testing preorders all coincide with the testing preorder induced by finite tests. Thus, fairness adds to the expressive power of nondeterministic testing but fails to do so in probabilistic testing.

Finally, we show that the probabilistic testing preorder refines the fair testing preorder for nondeterministic processes in the following sense: if two probabilistic processes are related by the probabilistic testing preorder, then their images under the forgetful homomorphism that erases probabilities are related by the fair testing preorder for nondeterministic processes.

2 Infinite Testing for Nondeterministic Processes

We begin by recalling the definitions of the testing preorders of [DNH83, Hen88]. Let Act denote a finite set of *actions* and let $\tau \notin \mathsf{Act}$. τ represents the distinguished unobservable internal action.

A *process* P over an alphabet Act is a labeled transition system $\langle K_P, \longrightarrow, p^I\rangle$ where K_P is a (countable) set of states, $p^I \in K_P$ is the initial state and $\longrightarrow \subseteq K_P \times \mathsf{Act} \cup \{\tau\} \times K_P$ is an image-finite *transition relation*. We write $p \xrightarrow{a} p'$ for $(p, a, p') \in \longrightarrow$ and $p \longrightarrow$ when there is an a and a p' such that $p \xrightarrow{a} p'$. A state p is *terminal* if $p \not\longrightarrow$.

A *finite execution fragment* of a process P is a sequence $p_0 \xrightarrow{a_1} p_1 \xrightarrow{a_2} \ldots \xrightarrow{a_n} p_n$ where $p_i \xrightarrow{a_i} p_{i+1}$. An *infinite execution fragment* is a sequence $p_0 \xrightarrow{a_1} p_1 \xrightarrow{a_2} \ldots \xrightarrow{a_n} p_n \xrightarrow{a_{n+1}} \ldots$ where $p_i \xrightarrow{a_i} p_{i+1}$. An *execution* is a finite or an infinite execution fragment with $p_0 = p^I$. The set of all executions of a process P is denoted by $Exec_P$. An execution is said to be *complete* if it is infinite or if it is finite and ends in a terminal state. A state p is said to be *divergent* if there is an infinite execution fragment of the form $p \xrightarrow{\tau} p_1 \xrightarrow{\tau} p_2 \ldots$ emanating from p. A process without divergent states is said to be divergence-free.

Given processes P and Q, the *composition* of P and Q is the process $P \,||\, Q = \langle K_P \times K_Q, \longrightarrow_{P\,||\,Q}, (p^I, q^I)\rangle$ where $\longrightarrow_{P\,||\,Q}$ is given by

$$\begin{aligned} p \xrightarrow{\tau} p' &\Rightarrow (p,q) \xrightarrow{\tau}_{P\,||\,Q} (p',q) \\ q \xrightarrow{\tau} q' &\Rightarrow (p,q) \xrightarrow{\tau}_{P\,||\,Q} (p,q') \\ a \neq \tau \;\&\; p \xrightarrow{a} p' \;\&\; q \xrightarrow{a} q' &\Rightarrow (p,q) \xrightarrow{a}_{P\,||\,Q} (p',q') \end{aligned}$$

A process P is said to be *finite* if it is finite-state and the relation $\longrightarrow$ is acyclic. Every execution of a finite process is finite. The composition of a finite process with any divergence-free process yields a finite process.

A *test* $T = \langle K_T, \longrightarrow_T, t^I, Suc_T\rangle$ is a process $\langle K_T, \longrightarrow_T, t^I\rangle$ with a distinguished set of *success* states $Suc_T \subseteq K_T$. (We often use T to refer its underlying process $\langle K_T, \longrightarrow_T, t^I\rangle$.) An execution of a test T is said to be *successful* if it ends in a state in Suc_T.

In this paper we require that processes be divergence-free and that every success state be a terminal state. The results in [CDS+98] that we build on assume that processes are divergence-free. The assumption on success states is for consistency with [NC95] and [CDS+98]. Traditional testing does not require that success states be terminal but any run that passes through a success state is considered successful. Thus, w.l.o.g., we may restrict success states to terminal states.

Given a process P and a test T, $P \,||\, T$ is a test with $Suc_{P\,||\,T} = \{(p,t) \mid t \in Suc_T \wedge (p,t) \not\longrightarrow\}$. We say that a process P may pass the test T, written

$P\ may\ T$, if the test $P \,||\, T$ has a successful execution. We say that a process P must pass the test T, written $P\ must\ T$, if every complete execution of the test $P \,||\, T$ is successful. (Usually success is indicated by a distinguished action w (see [Hen88, NC95]). It is easy to verify that our formulation using terminal success states is equivalent to the one presented in [NC95] for divergence-free processes).

A test T is said to be finite if its underlying process is finite. The *may-testing* preorder $\sqsubseteq_{may}$ is given by $P \sqsubseteq_{may} Q$ if for every finite test T, $P\ may\ T \Rightarrow Q\ may\ T$ and the *must-testing* preorder $\sqsubseteq_{must}$ is given by $P \sqsubseteq_{must} Q$ if for every finite test T, $P\ must\ T \Rightarrow Q\ must\ T$. Further, we say $P \sqsubseteq_{test} Q$ if $P \sqsubseteq_{may} Q$ and $P \sqsubseteq_{must} Q$.

If we remove the restriction that the tests be finite we get preorders $\sqsubseteq_{smay}$, $\sqsubseteq_{smust}$ and $\sqsubseteq_{s}$, respectively. We call such tests *simple infinite tests*. In [DNH83, Hen88], alternative characterizations of these preorders are derived and as a consequence it turns out that $\sqsubseteq_{may} = \sqsubseteq_{smay}$, $\sqsubseteq_{must} = \sqsubseteq_{smust}$ and $\sqsubseteq_{test} = \sqsubseteq_{s}$. This result relies heavily on the fact that every infinite execution is deemed unsuccessful.

Natarajan and Cleaveland [NC95] and Brinksma et al. [BRV95] present a different criterion for deciding which infinite executions are successful. A state t of a test is said to be *potentially successful* if there is an execution fragment starting at t that leads to a success state. Under fair testing, a complete finite execution is successful if it ends in a success state and an infinite execution is successful if every state it visits is potentially successful. When a process P may/must satisfy a test T is defined analogous to the finite case and let the resulting preorders be $\sqsubseteq_{fmay}$, $\sqsubseteq_{fmust}$ and $\sqsubseteq_{f}$. It turns out that $\sqsubseteq_{fmay}$ is the same as $\sqsubseteq_{may}$ while $\sqsubseteq_{fmust}$ is finer than $\sqsubseteq_{must}$.

We now propose a third kind of infinite test. *Büchi* tests, as the name suggests, are inspired by the theory of automata over infinite sequences. A Büchi test T is a tuple $\langle K_T, \longrightarrow_T, t^I, Suc_T, \mathrm{B}\rangle$ where $\langle K_T, \longrightarrow_T, t^I, Suc_T\rangle$ is a test and $\mathrm{B} \subseteq K_T$. A finite execution of a Büchi test is successful if it ends in a state in Suc_T. An infinite execution of a Büchi test is successful if it visits the set B infinitely often.

Given a process P and a Büchi test T we treat $P \,||\, T$ as a Büchi test with $\mathrm{B}_{P\,||\,T} = \{(p,t) \mid t \in \mathrm{B}\}$. We say $P\ may_b\ T$ if $P \,||\, T$ has a successful execution and $P\ must_b\ T$ if every complete execution of $P \,||\, T$ is successful. As in the case of finite tests, these relations induce three preorders $\sqsubseteq_{bmay}$, $\sqsubseteq_{bmust}$ and $\sqsubseteq_{b}$.

Let T be the test $\langle\{t,t'\}, \{(t,a,t),(t,b,t),(t,b,t'),(t',b,t')\}, t, \emptyset, \{t'\}\rangle$. Then, $P\ may_b\ T$ if and only if P has an infinite execution with finite number of a-transitions, a property seemingly beyond the reach of finite tests. Nevertheless, the following theorem shows that the preorders induced by Büchi test are no different from those induced by finite tests.

Theorem 1. *$\sqsubseteq_{bmay} = \sqsubseteq_{may}$, $\sqsubseteq_{bmust} = \sqsubseteq_{must}$ and $\sqsubseteq_{b} = \sqsubseteq_{test}$.*

Proof. (Sketch) For each test T let $Unf(T)_i$ be the "i-depth" approximation of T. We show that for any process P, P $must_b$ T if and only if P $must_b$ $Unf(T)_i$ for all i and P may_b T if and only if P may $Unf(T)_i$ for all i.

Theorem 1 critically depends on the image-finiteness assumption. Büchi tests do not coicide with simple tests for processes with infinite branching. Divergence can be used to simulate infinite branching and consequently Theorem 1 also does not hold for processes with divergence. It is possible to modify the definition of Büchi tests to generalize this theorem to arbitrary processes, but such modifications are technical and an operational justification seems hard.

3 Probabilistic Processes and Testing

In this section, we recall the basic ingredients of the probabilistic testing framework of [CDS$^+$98], namely, probabilistic processes and tests, the notion of a process passing a test with a certain probability, the probabilistic testing preorder, and the alternative characterization of the preorder. To simplify the presentation, we restrict ourselves to processes and tests without τ-moves. However, all our results extend to divergence-free processes and tests with τ-moves.

Let Act be a finite set of *atomic actions* and let a, b, ... range over Act. A *probabilistic process* P over Act is a tuple $\langle K_P, \longrightarrow, p^{\mathrm{I}}, \mu_P\rangle$ where $\langle K_P, \longrightarrow, p^{\mathrm{I}}\rangle$ is a process and $\mu_P : (K_P \times \mathsf{Act} \times K_P) \to [0, 1]$ is the *transition distribution function* such that for each $p \in K_P$,

$$\sum_{\substack{a \in \mathsf{Act},\\ p' \in K_P}} \mu_P(p, a, p') \in \{0, 1\}$$

Further, $p \xrightarrow{a} p'$ if and only if $\mu_P(p, a, p') > 0$. Intuitively, $\mu_P(p, a, p')$ is the probability that atomic action a is executed from state p to become state p'. We write $\mu_P(p, a)$, for $\sum_{p' \in P} \mu_P(p, a, p')$.

Notice that we are working in the generative model of probabilistic processes [GSS95]. Let $P = \langle K_P, \longrightarrow_P, p^{\mathrm{I}}, \mu_P\rangle$ and $Q = \langle K_Q, \longrightarrow_Q, q^{\mathrm{I}}, \mu_Q\rangle$ be probabilistic processes. Then their *composition*, denoted $P \,||\, Q$, is the tuple $\langle K_P \times K_Q, \longrightarrow_{P\,||\,Q}, (p^{\mathrm{I}}, q^{\mathrm{I}}), \mu_{P\,||\,Q}\rangle$ such that

$$\mu_{P\,||\,Q}((p, q), a, (p', q')) = \begin{array}{ll} \frac{\mu_P(p,a,p') \cdot \mu_Q(q,a,q')}{\nu(p,q)} & \text{if } \nu(p, q) > 0 \\ 0 & \text{otherwise} \end{array}$$

where $\nu(p, q) = \sum_{b \in \mathsf{Act}} \mu_P(p, b) \cdot \mu_Q(q, b)$.

Although P is a generative process it cannot autonomously decide which transitions to offer while interacting with an environment. The set of enabled transitions is determined by the environment and the probability assigned to an enabled transition is the conditional probability that this transition is taken from the set of all enabled transitions. In this sense, our processes are reactive and different from (labeled) Markov processes. As we shall see later, however, the interaction system consisting of a process and a test may be analysed as a Markov chain.

A probabilistic *test* is a probabilistic process with some of its states designated as success states. That is, a probabilistic test is a tuple of the form $\langle K_T, \longrightarrow_T, t^{\mathrm{I}}, \mu_T, Suc_T\rangle$ where $\langle K_T, \longrightarrow_T, t^{\mathrm{I}}, \mu_T\rangle$ is a probabilistic process and $Suc_T \subseteq K_T$ is the set of *success states*. Further we will assume that if $t \in Suc_T$ then $\sum_{a\in\mathsf{Act}} \mu_T(t,a) = 0$. That is, as in the nondeterministic case, there are no transitions out of success states. Once the test has reached a success state the run is deemed successful (and thus the rest of the run is of no interest). In this section we assume that tests are finite, i.e., they are finite-state and the transition relation $\longrightarrow$ of a test is acyclic.

For an execution $\sigma = p_1 \xrightarrow{a_1} p_2 \xrightarrow{a_2} \ldots \xrightarrow{a_{n-1}} p_n$, let

$$Prod(\sigma) = \prod_{1\leq i\leq n-1} \mu_P(p_i, a_i, p_{i+1})$$

For any finite process P, the *execution probability distribution function* $\Pr : Exec_P \to [0,1]$ is defined inductively as follows:

$$\begin{aligned}\Pr(p_1 \xrightarrow{a_1} p_2 \ldots \xrightarrow{a_{n-1}} p_n) &= Prod(p_1 \xrightarrow{a_1} p_2 \ldots \xrightarrow{a_{n-1}} p_n) && \text{if } p_n \text{ is terminal}\\ &= 0 && \text{otherwise}\end{aligned}$$

and is lifted to sets $\mathcal{E} \subseteq Exec_P$ of executions as follows:

$$\Pr(\mathcal{E}) = \textstyle\sum_{\sigma\in\mathcal{E}} \Pr(\sigma)$$

Intuitively, $\Pr(\sigma)$ is the probability that the process P engages in execution σ. We assume that an execution always proceeds to completion and thus the probability of incomplete executions is taken to be 0. The following proposition establishes that Pr is indeed a probability distribution on executions.

Proposition [CDS+98] *Let P be any finite process. Then* $\sum_{\sigma\in Exec_P} \Pr(\sigma) \in \{0,1\}$.

Since tests are finite processes the execution probability function is defined. The *success probability* of a test T is given by $Pr(\mathrm{SE}_T)$, where SE_T is the set of all executions of T that end in success states.

Every test can be treated as a process by disregarding its set of success states. Thus, given a probabilistic process $P = \langle K_P, \longrightarrow_P, p^I, \mu_P\rangle$ and a test $T = \langle K_T, \longrightarrow_T, t^I, \mu_T, Suc_T\rangle$, we can define the composition $P \,||\, T$. $P \,||\, T$ can also be treated as a test with $Suc_{P\,||\,T} = \{(p,t) \mid p \in K_P,\ t \in Suc_T\}$. By the definition of composition, there are no transitions out of states in $Suc_{P\,||\,T}$. Further, if T

is a finite process then so is $P \,||\, T$. The *probability that P passes T* is simply the success probability of $P \,||\, T$. We write $P\, pass_\pi\, T$ if $Prob(\mathrm{SE}_{P\,||\,T}) \geq \pi$.

The testing preorder for probabilistic processes [CDS+98] can now be defined as follows. Let P and Q be processes. Then $P \sqsubseteq^{pr}_{test} Q$ if for all tests T, $\pi \in [0,1]$, $P\, pass_\pi\, T$ implies $Q\, pass_\pi\, T$ (or equivalently $\Pr(\mathrm{SE}_{P\,||\,T}) \leq \Pr(\mathrm{SE}_{Q\,||\,T})$).

Let Erase be a mapping that transforms a probabilistic process or test into a nondeterministic process or test by erasing probabilities; i.e., if $P = \langle K_P, \longrightarrow_P, p^{\mathrm{I}}, \mu_P\rangle$, then $\mathrm{Erase}(P) = \langle K_P, \longrightarrow_P, p^{\mathrm{I}}\rangle$. [CDS+98] relates the probabilistic testing preorders to the $\sqsubseteq_{may}$ and $\sqsubseteq_{must}$ preorders. If $P \sqsubseteq^{pr}_{test} Q$ then $\mathrm{Erase}(P) \sqsubseteq_{must} \mathrm{Erase}(Q)$ and $\mathrm{Erase}(P) \sqsubseteq_{may} \mathrm{Erase}(Q)$. They also present an alternative characterization of the probabilistic testing preorder $\sqsubseteq^{pr}_{test}$ which we now describe.

A distribution over Act is a map $D : \mathsf{Act} \to [0,1]$ such that $\sum_{a \in \mathsf{Act}} D(a) \in \{0,1\}$. Let Dist denote the set of all distributions. A *probabilistic trace* $(a_1, D_1) \cdots (a_n, D_n)$ over Act is an element of $(\mathsf{Dist} \times \mathsf{Act})^*$. Tr_{Act} denotes the set of all probabilistic traces over Act, and $\varepsilon \in Tr_{\mathsf{Act}}$ denotes the empty probabilistic trace.

Intuitively, a probabilistic 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. The function M_P defined next computes the probability that a process P, from a given state, engages in such an interaction session.

Let $P = \langle K_P, \longrightarrow_P, p^{\mathrm{I}}, \mu_P\rangle$ be a probabilistic process. Function $M_P : K_P \times Tr_{\mathsf{Act}} \to [0,1]$ is defined as follows:

$$M_P(p, \varepsilon) = 1$$
$$M_P(p, (a, D)\ s') = \begin{cases} \sum_{p' \in K_P} \frac{\mu_P(p,a,p') \cdot D(a) \cdot M_P(p', s')}{\nu(p,D)} & \text{if } \nu(p,D) > 0 \\ 0 & \text{otherwise} \end{cases}$$

where $\nu(p, D) = \sum_{b \in \mathsf{Act}} \mu_P(p, b) \cdot D(b)$.

The preorder induced by probabilistic traces, $\ll$, is defined by $P \ll Q$ if for all $s \in Tr_{\mathsf{Act}}$, $M_P(p^{\mathrm{I}}, s) \leq M_Q(q^{\mathrm{I}}, s)$. The following characterization is shown in [CDS+98].

Theorem 2 [CDS+98].

$$P \sqsubseteq^{pr}_{test} Q \text{ if and only if } P \ll Q$$

4 A Probability Measure on Sets of Executions

For any finite execution σ, let $[\sigma]$ denote the set of finite and infinite executions of P that extend σ. A set of executions is said to be *simple basic measurable* if it is either of the form $[\sigma]$ or $\{\sigma\}$ where $\{\sigma\} \neq [\sigma]$.

We assign probabilities to simple basic measurable sets as follows:

$$Prob_P([\sigma]) = Prod(\sigma)$$

and when $\{\sigma\} \neq [\sigma]$

$$Prob_P(\{\sigma\}) = 0$$

We omit the subscript P when the identity of the process is clear from the context. A standard result tells us that this function extends to a probability measure over the σ-algebra generated by these sets. The reader is referred to [WSS97] for the details.

5 Simple Infinite Tests

As mentioned earlier, an execution of a simple infinite test is considered successful if it is finite and ends in a success state. Since there are no outgoing transitions from success states, $[\sigma] = \{\sigma\}$ for any successful execution σ. The probability of success is given by $Prob(\mathrm{SE}_T) = \sum_{\sigma \in \mathrm{SE}_T} Prob([\sigma])$.

Given a probabilistic process P and a simple infinite test T the probability with which P passes T is just the success probability of $P \,||\, T$. The class of simple infinite tests induce a preorder $\sqsubseteq_s^{pr}$ given by $P \sqsubseteq_s^{pr} Q$ if and only if for all probabilistic tests T, $Prob(\mathrm{SE}_{P \,||\, T}) \leq Prob(\mathrm{SE}_{Q \,||\, T})$.

The preorder induced by simple infinite tests is finer than the one induced by probabilistic traces as every finite test is also a simple infinite test. We now show that the converse is also true.

Theorem 3. $\sqsubseteq_{test}^{pr} \;=\; \ll \;=\; \sqsubseteq_s^{pr}$

Proof. It is easy to see that $\sqsubseteq_s^{pr} \subseteq \ll$. We now prove the converse. Let P and Q be probabilistic processes such that $P \ll Q$. Let us examine the set of successful executions of $R \,||\, T$ for any process R. Let σ be any execution of T. Let $\mathrm{E}_{R \,||\, T}(\sigma)$ denote the set of executions of $R \,||\, T$ whose projection onto the T component extends σ. In [CDS+98] it is shown that there is a probabilistic trace $tr(\sigma)$ and a constant v (that depends on T and σ) such that $Prob(\mathrm{E}_{R \,||\, T}(\sigma)) = \frac{M_R(r^I, tr(\sigma))}{v}$.

Let P and Q be probabilistic processes with $P \ll Q$. Then,

$$Prob(\mathrm{E}_{P \,||\, T}(\sigma)) \;=\; \frac{M_P(p^I, \mathrm{tr}(\sigma))}{v} \;\leq\; \frac{M_Q(q^I, \mathrm{tr}(\sigma))}{v} \;=\; Prob(\mathrm{E}_{Q \,||\, T}(\sigma))$$

But $\mathrm{SE}_{P \,||\, T} = \bigcup_{\sigma \in \mathrm{SE}_T} \mathrm{E}_{P \,||\, T}(\sigma)$. Moreover, $\mathrm{E}_{P \,||\, T}(\sigma) \cap \mathrm{E}_{P \,||\, T}(\sigma') = \emptyset$ whenever $\sigma \neq \sigma'$. Thus,

$$Prob(\mathrm{SE}_{P \,||\, T}) \;\leq\; Prob(\mathrm{SE}_{Q \,||\, T})$$

6 Büchi Tests

We next consider infinite probabilistic tests with a Büchi acceptance condition. Formally, a *Büchi test* is a tuple $\langle K_T, \longrightarrow_T, t^I, \mu^T, Suc_T, \mathrm{B}\rangle$ where $\langle K_T, \longrightarrow_T, t^I, \mu_T, Suc_T\rangle$ is a probabilistic test and B is a subset of K_T. An execution of a Büchi test is said to be successful if either it is finite and ends in a state in Suc_T or it is infinite and visits B infinitely often. We use SE_T to denote the set of successful executions of a Büchi test T. This set is measurable.

If P is a probabilistic process and T is a Büchi test then $P \,||\, T$ is a Büchi test with $\mathrm{B}_{P\,||\,T} = \{(r,t) \mid t \in \mathrm{B}\}$. The probability with which P passes T is given by $Prob(\mathrm{SE}_{P\,||\,T})$. The class of Büchi tests induces a preorder $\sqsubseteq_b^{pr}$ on the set of probabilistic processes such that $P \sqsubseteq_b^{pr} Q$ if and only if, for all Büchi tests T, $Prob(\mathrm{SE}_{P\,||\,T}) \leq Prob(\mathrm{SE}_{Q\,||\,T})$.

A Büchi test is said to be a *pure* Büchi test if its set of success states is empty.

Proposition 4. *For every Büchi test T there is a pure Büchi test T^b and a simple infinite test T^u such that, for any probabilistic process P, $Prob(\mathrm{SE}_{P\,||\,T}) = Prob_{P\,||\,T^b}(\mathrm{SE}_{P\,||\,T^b}) + Prob_{P\,||\,T^u}(\mathrm{SE}_{P\,||\,T^u})$.*

Theorem 5. $\sqsubseteq_{test}^{pr} \;=\; \ll \;=\; \sqsubseteq_b^{pr}$

Proof. (sketch) As every finite test may be treated as a Büchi test the preorder induced by Büchi tests is finer than $\ll$.

Let P and Q be probabilistic processes with $P \ll Q$. Let T be any pure Büchi test and R be any probabilistic process. It is easy to show that $Prob(\mathrm{SE}_{R\,||\,T}) = Prob(\bigcap_{i\in\omega} S_i^R)$, where S_i^R is the set of executions of $R \,||\, T$ that contain at least i states whose T component is from B.

Let T_i be the set of executions of T that end in a state from B and contain precisely i occurrences of states from B. Then, $S_i^R = \bigcup_{\sigma\in T_i} \mathrm{E}_{R\,||\,T}(\sigma)$. Using this characterization of S_i^R, an argument identical to that in the proof of Lemma 3 shows that $Prob(S_i^P) \leq Prob(S_i^Q)$. The proof also relies on the following property of probability measures.

Fact: Let $A_1, A_2, \ldots$ be a family of measurable sets with $A_i \supseteq A_{i+1}$. Then, $Prob(\bigcap_{i\in\omega} A_i) = \lim_i \{Prob(A_i)\}_{i\in\omega}$.

Thus, $Prob(P \,||\, T) = \lim_i \{Prob(S_i^P)\}_{i\in\omega} \leq \lim_i \{Prob(S_i^Q)\}_{i\in\omega} = Prob(Q \,||\, T)$.

If T is a general Büchi test, by Proposition 4, we can decompose T into a pure Büchi test and a simple infinite test. We then use the property of pure Büchi tests established above and Lemma 3 to derive the desired result.

7 Fair Tests

In this section we will restrict ourselves to finite-state processes and tests. A state t of a fair test T is said to be *potentially successful* if there is a finite execution fragment $t = t_1 \xrightarrow{a_1} \ldots \xrightarrow{a_{n-1}} t_n$ with $t_n \in Suc_T$. It is said to be *dead* otherwise. A potentially successful state that is not a success state is called a *live* state.

A finite execution of a test is said to be *successful* if it ends in a success state. An infinite execution is said to be successful if every state in the execution is potentially successful. It is not difficult to show that this set is measurable. Thus we can define the probability of a process passing a fair test and the associated preorder $\sqsubseteq_f^{pr}$.

A finite execution is said to be *potentially successful* if it ends in a potentially successful state and is said to be dead otherwise. An execution is said to be *unsuccessful* if it is not successful. An execution that ends in a live state is a *live execution*. We use SE_T, FSE_T, PSE_T, LE_T and DE_T to denote the sets of successful, finite successful, potentially successful, live and dead executions of T, respectively.

7.1 Turning Probabilistic Processes into Markov Chains

Given a probabilistic test T the process

$$Mv(T) = \langle K_T^M, \longrightarrow_T^M, t^I, \mu_T^M, Suc_{Mv(T)} \rangle$$

is defined as follows: :

$$\begin{aligned}
K_T^M &= \{t^I\} \cup \{(t,a) \mid t \in K_T \ \&\ a \in Act\} \\
\mu_T^M((t,a),b,(t',b)) &= \mu_T(t,b,t') \\
\mu_T^M(t^I,a,(t',a)) &= \mu_T(t^I,a,t') \\
\mu_T^M((t,a),b,(t',c)) &= 0 \qquad\qquad \{b \neq c\} \\
\mu_T^M(t^I,a,(t,b)) &= 0 \qquad\qquad \{a \neq b\}
\end{aligned}$$

The initial state of $Mv(T)$ is t^I. The set of success states of $Mv(T)$ is $Suc_{Mv(T)} = \{(t,a) \mid t \in Suc_T\}$.

We have made multiple copies of each state so that in the resulting process there is at most one transition from a given state to another. Furthermore, the transitions used in an execution can be determined merely from the sequence of states involved in the execution.

There is an obvious bijective correspondence between the set of executions of T beginning at t^I and that of $Mv(T)$ beginning at t^I: The finite execution $\sigma = t^I \xrightarrow{a_1} t_2 \xrightarrow{a_2} \ldots t_N$ corresponds to the execution $Mv(\sigma) = t^I \xrightarrow{a_1} (t_2, a_1) \xrightarrow{a_2} (t_3, a_2) \ldots (t_n, a_{n-1})$ and similarly for infinite executions. We use ρ to denote this bijection.

The process $Mv(T)$ is almost a Markov chain (see [Fel50]). In a Markov chain the sum of the probabilities of the transitions out of every state must be 1 while in $Mv(T)$ this sum is zero for some states (at least at every success state). We make one final modification to $Mv(T)$, without sacrificing the correspondence mentioned above, to turn it into a Markov chain. For each state x that has no outgoing transitions, add a self-loop labeled δ with $\mu_T^M(x, \delta, x) = 1$, where δ does not appear in Act. We continue to refer to this process as $Mv(T)$. (Formally we have extended μ_T^M so that it assigns a probability 0 for δ transitions at every other state and expanded the alphabet to $\mathsf{Act} \cup \{\delta\}$.)

Clearly ρ is still a bijective correspondence between executions of T beginning at t^I and those of $Mv(T)$ that begins at t^I which do not make use of transitions on δ. An execution of the Markov chain Mv(T) is said to be *successful* if it is infinite and all the states that it visits are potentially successful or if it is finite and it visits a success state.

An execution of $Mv(T)$ is said to be minimally good if it ends in a success state and none of its prefixes end in success states. We write $\mathrm{MGE}_{Mv(T)}$ for the set of minimally good executions of $Mv(T)$.

Lemma 6. *The map ρ relates minimally dead executions of T with minimally dead executions of $Mv(T)$ in a bijective manner. The map ρ (bijectively) relates the finite successful executions of T with the minimally good executions of $Mv(T)$. Further ρ preserves probability, that is $Prod(\sigma) = Prod(\rho(\sigma))$. The probability of success of T equals the probability of success of $Mv(T)$.*

Lemma 7. *The success probability of $Mv(T)$ is $\sum_{\gamma \in \mathrm{MGE}_{Mv(T)}} Prod(\gamma)$.*

Proof. Recall, that an execution of $Mv(T)$ is successful iff it is infinite and if every state it visits is potentially successful or it is finite and visits a success state. This set can be partitioned in two: the set of successful executions that visit a success state and those that don't.

Clearly, any successful execution that enters some success state necessarily belong to $[\gamma]$ for some $\gamma \in \mathrm{MGE}_{Mv(T)}$. On the other hand, the only transition out of a success state is a self-loop and so, every execution that is in $[\gamma]$, with $\gamma \in \mathrm{MGE}_{Mv(T)}$, is successful. Let $\mathrm{L}_{Mv(T)}$ be the set of infinite executions of $Mv(T)$ that are successful but never visit a success state.

$$\mathrm{SE}_{Mv(T)} \ = \ (\bigcup_{\gamma \in \mathrm{MGE}_{Mv(T)}} [\gamma]) \ \cup \ \mathrm{L}_{Mv(T)}$$

Thus

$$Prob(\mathrm{SE}_{Mv(T)}) = \sum_{\gamma \in \mathrm{MGE}_{Mv(T)}} Prob([\gamma]) \ + \ Prob(\mathrm{L}_{Mv(T)})$$

We now show that $Prob(\mathrm{L}_{Mv(T)}) = 0$. (Note that since $\mathrm{L}_{Mv(T)}$is the difference between two measurable sets it must indeed be measurable). Let

$$L_n = \bigcup_{\gamma \in \mathrm{LE}_{Mv(T)}, |\gamma| = n} [\gamma]$$

Then, $L_{n+1} \subseteq L_n$ and $\mathrm{L}_{Mv(T)} = \bigcap_{n \geq 0} L_n$. Thus,

$$Prob(\mathrm{L}_{Mv(T)}) \;=\; \mathsf{lim}_n \{Prob(L_n)\}_{n \in \omega}$$

Consider the underlying graph of the Markov chain $Mv(T)$ whose vertices are the states and whose edges are defined by transitions of nonzero probability (recall that there is at most one transition of nonzero probability from any state to another).

In the decomposition of this graph into its maximal strongly connected components, components containing live states will consist entirely of live states; each success state lies in a singleton maximal strongly connected component and live states are not reachable from dead states. Further the components that contain live states will necessarily have outgoing edges. In the terminology of Markov chains (see [Fel50, Shi80]) every live state is transient (or nonrecurrent). A basic result in the theory of Markov chains states that for any finite Markov chain and any real number ϵ there is an n such that, for any $m \geq n$ and state q, the probability that an m length execution beginning at q ends in a transient state is less than or equal to ϵ (see [Shi80]) But $Prob(L_n)$ is at most the probability that a n length run from q^I ends in a transient state. Thus the sequence $\{Prob(L_n)\}_{n \in \omega}$ must indeed converge to 0.

Thus $Prob(\mathrm{L}_{Mv(T)}) = 0$ and so $Prob(Mv(S)) = \sum_{\gamma \in \mathrm{MGE}_{Mv(T)}} Prod(\gamma)$.

Now, using Lemmas 6 and 7 it is easy to establish the following theorem.

Theorem 8. $\sqsubseteq^{pr}_{test} \;=\; \ll \;=\; \sqsubseteq^{pr}_{f}$

In contrast, nondeterministic fair tests induce a finer preorder than finite tests even when restricted to finite-state τ-free processes and tests.

7.2 Finite Probabilistic Testing Refines Fair Testing

We now show that finite probabilistic testing induces a preorder that is finer than that induced by nondeterministic fair testing.

Theorem 9. *For finite-state probabilistic processes P and Q,*

$$P \sqsubseteq^{pr}_{test} Q \Rightarrow \mathrm{Erase}(P) \sqsubseteq_f \mathrm{Erase}(Q).$$

8 Conclusions

We have carefully examined the distinguishing power of infinite testing. For nonprobabilistic testing, we showed that Büchi testing offers no additional distinguishing power over finite testing. In the probabilistic setting, we showed that nothing new is gained by considering infinite tests, as their induced preorders coincide with finite testing. Thus, fairness fails to add to the distinguishing power of probabilistic tests. Finally, we showed that the probabilistic testing preorder refines the fair testing preorder for nondeterministic processes.

The results of [CDS+98] can be extended to processes with divergence and consequently our results for simple tests and fair tests extend to processes with divergence. As future work, we would like to investigate the probabilistic preorder induced by infinite-state fair tests.

References

[BRV95] E. Brinksma, A. Rensink, and W. Vogler. Fair testing. In *CONCUR '95, LNCS 962*. Springer-Verlag, 1995.

[CDS+98] R. Cleaveland, Z. Dayar, S. A. Smolka, S. Yuen, and A. Zwarico. Testing preorders for probabilistic processes. *Information and Computation*, 1998. To Appear.

[DNH83] R. De Nicola and M. C. B. Hennessy. Testing equivalences for processes. *Theoretical Computer Science*, 34:83–133, 1983.

[Fel50] William Feller. *Probability Theory and its Applications.* John Wiley and Sons, 1950.

[GSS95] R. J. van Glabbeek, S. A. Smolka, and B. Steffen. Reactive, generative, and stratified models of probabilistic processes. *Information and Computation*, 121(1):59–80, August 1995.

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

[JY95] B. Jonsson and W. Yi. Compositional testing preorders for probabilistic processes. In *Proceedings of LICS '95*, pages 431–441, 1995.

[NC95] V. Natarajan and R. Cleaveland. Divergence and fair testing. In *Proceedings of ICALP '95,* Vol. 944 of *Lecture Notes in Computer Science*. Springer-Verlag, 1995.

[Seg96] R. Segala. Testing probabilistic automata. In *Proccedings of CONCUR '96, LNCS 1119*. Springer-Verlag, 1996.

[Shi80] A. N. Shiryayev. *Probability.* Graduate Texts in Mathematics, Springer-Verlag, 1980.

[WSS97] S.-H. Wu, S. A. Smolka, and E. W. Stark. Composition and behaviors of probabilistic I/O automata. *Theoretical Computer Science*, 176(1-2):1–38, 1997.

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

This article was processed using the LaTeX macro package with LLNCS style