Testing Preorders for Probabilistic Processes

Rance Cleaveland
Computer Science Dept.
N.C. State Univ.
Raleigh, NC 27695
USA
rance@csc.ncsu.edu

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

Amy Zwarico
Computer Science Dept.
Johns Hopkins Univ.
Baltimore, MD 21218
USA
amy@cs.jhu.edu

Abstract

We present a testing preorder for probabilistic processes based on the natural notion of a process passing a test with a certain probability. The theory enjoys close connections with the classical testing theory of Hennessy and DeNicola in that whenever a process passes a test with probability 1 (respectively some non-0 probability) in our setting, then the process *must* (respectively *may*) pass the test in the classical theory. In addition, we develop an alternative characterization of the probabilistic testing preorder that is based on the "must sets" characterization of De Nicola. Finally, we extend our theory of testing to *substochastic processes*, in which the sum of the probabilities of a process's outgoing transitions may be strictly less than 1, with the deficit representing the process' capacity for *undefined behavior*. A simple example involving the construction of pipelines from faulty buffer cells is given to illustrate how substochastic processes, and the resulting preorder, can be used to model fault-tolerant systems and to reason about system reliability.

1 Introduction

Real-life processes often exhibit behavior of a statistical or probabilistic nature. For example, through repeated observations, a failure rate of 1.6×10^{-3} *failures/sec* can be attributed to an off-the-shelf memory chip, or a communication medium can be seen to be dropping two messages over a one-minute interval with probability 0.0945.

In this paper, we present a semantic framework for *probabilistic reactive processes* based on the *testing preorder* of Hennessy and De Nicola [DNH83, Hen88]. We begin by formalizing the notion of a process passing a test with a certain probability. This leads naturally to the definition of a *probabilistic testing preorder* where process $\mathcal{P}$ is less than process $\mathcal{P}'$ in the preorder when, for all tests $\mathcal{T}$, the probability that $\mathcal{P}$ passes $\mathcal{T}$ is no greater than the probability that $\mathcal{P}'$ passes $\mathcal{T}$. This testing preorder has an alternative characterization that extends the characterization of [DN87] for the nonprobabilistic preorder by taking into account information about the actions a process may *refuse* (i.e. execute with probability 0).

We establish close connections with the testing theory of Hennessy and De Nicola as follows: let $\mathcal{P}$ be a probabilistic process, $\mathcal{T}$ a probabilistic test, and $h(\mathcal{P})$, $h(\mathcal{T})$ the "deprobabilized" process and test obtained by erasing the probabilities in the underlying transition systems. Then we have that $h(\mathcal{P})$ *may* pass $h(\mathcal{T})$ if $\mathcal{P}$ passes $\mathcal{T}$ with some non-zero probability, and $h(\mathcal{P})$ *must* pass $h(\mathcal{T})$ if $\mathcal{P}$ passes $\mathcal{T}$ with probability 1. Using these results, we show that if two probabilistic processes are related by our probabilistic testing preorder, then their deprobabilized images are related by the testing preorder of De Nicola and Hennessy. Thus, the testing theory of probabilistic processes is a refinement of the testing theory of nonprobabilistic processes.

Finally, we extend the theory to handle *substochastic processes*. A substochastic process has the property that the sum of the probabilities of its outgoing transitions may be strictly less than 1. The "missing" part of the probability distribution can be interpreted as a measure of the *undefined behavior* of the process, and can be seen as a probabilistic analogue of the phenomenon of divergence in the testing theory of nonprobabilistic processes. Substochastic processes enable the modeling of faulty and fault-tolerant processes, while the testing preorder is useful in the analysis of system reliability. This is illustrated by a simple example involving the construction of pipelines of faulty buffer cells.

The structure of the rest of this paper is as follows. Section 2 defines probabilistic processes, tests, and interaction systems, as well as our testing preorder for probabilistic processes. Section 3 connects our theory of testing for probabilistic processes with the theory of Hennessy and De Nicola for nondeterministic processes. Section 4 proposes another characterization of probabilistic testing, and Section 5 extends the theory of probabilistic testing to substochastic processes. Section 6 contains our conclusions and directions for future work.

Related Work

Two approaches to testing probabilistic processes have previously been proposed. Larsen and Skou [LS89] focus on *probabilistic bisimulation*, a natural and elegant extension of Milner and Park's strong bisimulation to probabilistic processes. Based on a simple, yet powerful, language for tests (in which copying is permitted) they obtain a testing characterization of probabilistic bisimulation. They also present a probabilistic extension of Hennessy-Milner logic that induces an equivalence on probabilistic processes that coincides with probabilistic bisimulation.

Christoff [Chr90] introduces a testing framework based on the probability distributions over finite sequences of observable events induced by probabilistic processes in their interactions with deterministic tests. An alternative characterization of the resulting preorders is given. His approach differs from our own in three key ways: 1) his notion of testing is based on the probability of event sequences in interaction systems, where ours is based on the notion of a process passing a test with a certain probability; 2) he uses deterministic tests, where in our setup, processes and tests are essentially indistinguishable; 3) as a result of 1) and 2), our approach is more closely grounded in the testing theory of De Nicola and Hennessy.

Other work has focused on the development of models of probabilistic processes. In [vGSST90], van Glabbeek et al. introduce three such models—namely, *reactive*, *generative*, and *stratified*— and show that they form a hierarchy: the reactive model is an abstraction of the generative model which in turn is an abstraction of the stratified model. We use generative processes for the processes being tested, for tests and for the *interaction systems* resulting from the application of tests to processes. Christoff also uses generative processes but deterministic tests. Larsen and Skou use reactive processes and nonprobabilistic tests with copying.

In [JS90], probabilistic extensions of a number of other well-known process equivalences are examined, including probabilistic trace, completed trace, failure, and ready.

2 Probabilistic Processes, Tests and Interaction Systems

In this section we define the basic elements for our theory: probabilistic processes, tests, and interaction systems. This leads us to a natural notion of probabilistic testing and a probabilistic testing preorder.

2.1 Probabilistic Processes and Tests

Probabilistic processes are modeled using generative transition probabilities [vGSST90], meaning that from any given state, the sum of the probabilities on outgoing transitions is either 0 or 1; i.e. the process is either terminated or not. Even though our interest is in reactive systems, we have chosen this model rather than the reactive one of [vGSST90, LS89]. Reactive systems often make more than one possible action available to the environment, which may in turn attempt to execute any one of them. The generative model permits the resulting nondeterminism as to which action actually takes place to be resolved probabilistically. This expressiveness is not available in the probabilistic reactive model, in which the sum of the probabilities of all outgoing actions *labeled by the same action* must be either 0 or 1.

Definition 1 *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*, and* τ *is a distinguished "silent action" that is assumed not to be in* Act*;*
- $\mu : P \times (Act \cup \{\tau\}) \times P \to [0,1]$*, the* transition probability distribution function*, satisfies: for all* $p \in P$*,* $\sum_{p' \in P, a \in Act \cup \{\tau\}} \mu(p, a, p') \in \{0, 1\}$*;*
- $p_0 \in P$ *is the* start state*.*

The class of all probabilistic processes is denoted **PP**.

In the sequel we sometimes write $p \xrightarrow{a,\pi} p'$ when $\mu(p,a,p') = \pi > 0$ and $p \xrightarrow{a} p'$ if $\mu(p,a,p') > 0$. We also use $\mu(p,a)$ as shorthand for $\sum_{p' \in P} \mu(p,a,p')$, and if $A \subseteq Act \cup \{\tau\}$ then we use $\mu(p,A)$ for $\sum_{a \in A} \mu(p,a)$. In the remainder of the paper we restrict our attention to *image-finite, divergence-free* processes. A process $\mathcal{P} = (P, Act, \mu, p_0)$ is image-finite if for any $p \in P$ and $a \in Act \cup \{\tau\}$, the set $\{p' \mid p \xrightarrow{a} p'\}$ is finite; $\mathcal{P}$ is divergence-free if no $p \in P$ is divergent. Here a state p is divergent if there is an infinite seqence $\langle p_0, p_1, \ldots \rangle$ with $p = p_0$ and $p_i \xrightarrow{\tau} p_{i+1}$.

We now turn our attention to defining the *tests* that will be applied to processes. Probabilistic tests are simply probabilistic processes with a set of states designated as *successful*. The reader should note the slight departure from the standard formulation [Hen88], in which a state of a test is considered successful if it can perform a w-transition.

Definition 2 *A* probabilistic test *is a 5-tuple* (P, Act, μ, p_0, S) *where* (P, Act, μ, p_0) *is a probabilistic process and* $S \subseteq P$ *is a set of* success states*.*

In contrast to [LS89], we require that transitions in tests be assigned probabilities. This is because we permit τ-transitions to occur in tests, and since τ actions are different from other actions (because processes/tests execute them autonomously), the model must account for the probability of their occurrence. In the remainder of the paper, we only consider tests that are finite in the usual sense (i.e. finite-branching and incapable of an infinite execution sequence). For both probabilistic processes and tests, we often drop the adjective "probabilistic" when this is clear from the context.

For notational conciseness, we make use of the following simple algebra of probabilistic processes/tests in the remainder of the paper.

Definition 3 *The* algebra of finite serial probabilistic processes and tests *is defined inductively as follows.*

- nil *is the finite serial process given by* $(\{nil\}, \emptyset, \mu_{nil}, nil)$ *with* $\mu_{nil}(nil, \tau) = 0$.
- good *is the finite serial test obtained by designating the start state of* nil *to be successful, and* bad *is the finite serial test obtained by designating the start state of* nil *to be unsuccessful. Formally,* good *is the finite serial test given by* $(\{nil\}, \emptyset, \mu_{nil}, nil, \{nil\})$, *while* bad *is the finite serial test given by* $(\{nil\}, \emptyset, \mu_{nil}, nil, \emptyset)$.
- *Let* $\mathcal{P}_i = (P_i, Act_i, \mu_i, p_i)$ *be a finite set of finite serial processes with the property that for any* $p \in P_i \cap P_j$, $\mu_i(p, a, p') = \mu_j(p, a, p')$ *for all* $a \in Act_i \cup Act_j$ *and* $p' \in P_i \cup P_j$. *Then* $a_1[\pi_1] \,.\, \mathcal{P}_1 + \cdots + a_n[\pi_n] \,.\, \mathcal{P}_n$, *with* $\sum \pi_i = 1$, *is a finite serial process. It is given by* (P, Act, μ, p_0), *where* $P = \bigcup_{i=1}^{n} P_i \cup \{p_0\}$ *(with* $p_0 \notin \bigcup_{i=1}^{n} P_i$*),* $Act = (\bigcup_{i=1}^{n} A_i \cup \{a_1, \ldots, a_n\}) - \{\tau\}$, *and* μ *given as follows.*
$$\mu(p, a, p') = \begin{cases} \pi_i & \text{if } p = p_0, a = a_i \text{ and } p' = p_i \\ \mu_i(p, a, p') & \text{if } p, p' \in P_i \\ 0 & \text{otherwise} \end{cases}$$
Similarly, if $\mathcal{T}_i = (T_i, Act_i, \mu_i, t_i, S_i)$ *is a finite set of finite serial tests with the appropriate property, then* $a_1[\pi_1] \,.\, \mathcal{T}_1 + \cdots + a_n[\pi_n] \,.\, \mathcal{T}_n$, *with* $\sum \pi_i = 1$, *is a finite serial test given by* (T, Act, μ, t_0, S), *where* (T, Act, μ, t_0) *is as specified above and* $S = \bigcup_{i=1}^{n} S_i$.

So $a[1] \,.\, nil$, $a[\frac{1}{3}] \,.\, nil + b[\frac{2}{3}] \,.\, c[1] \,.\, nil$, and $a[1] \,.\, (b[\frac{1}{2}] \,.\, nil + \tau[\frac{1}{2}] \,.\, nil)$ all represent finite serial processes. In what follows we sometimes leave off trailing $.\, nil$'s and probabilities of 1; thus, e.g., $b \,.\, c$ is shorthand for $b[1] \,.\, c[1] \,.\, nil$. Readers should note the (intended) similarity with Milner's CCS [Mil80, Mil89].

2.2 Interaction Systems

We use *interaction systems* to formalize the notion of applying a test $\mathcal{T}$ to a process $\mathcal{P}$. Intuitively our definition of an interaction system can be interpreted in terms of a *coin-flipping protocol*. Consider the interaction between a process $\mathcal{P}$ and test $\mathcal{T}$ where initially $\mathcal{P}$ has n transitions of non-zero probability and $\mathcal{T}$ has m transitions of non-zero probability. To commence interacting, $\mathcal{P}$ flips an n-sided coin (weighted in accordance with the probability distribution on its outgoing transitions) and, simultaneously, $\mathcal{T}$ flips an m-sided coin (similarly weighted). Should $\mathcal{T}$'s coin flip indicate a transition a (where a is not τ) and $\mathcal{P}$'s coin flip indicate an a-transition as well, then the interaction system will perform an a-transition as a joint transition between the process and the test. On the other hand, if either $\mathcal{P}$'s or $\mathcal{T}$'s coin, but not both, indicates a τ-transition, then it will do its τ independently of the other. If $\mathcal{P}$ and $\mathcal{T}$'s coins both indicate a τ, a coin biased by ρ to performing $\mathcal{P}$'s τ is flipped to decided whether $\mathcal{P}$ or $\mathcal{T}$ should perform its τ. Otherwise, the coin flips are repeated. The protocol continues until the process and test reach a configuration where no further transitions are possible. Due to the finite nature of tests, the protocol is guaranteed to terminate in the sense that the probability of not terminating is 0. The reader should note the similarity with traditional testing in which exactly one of the process or test can perform a τ at any time.

Definition 4 *Let $\mathcal{P} = (P, Act, \mu_P, p_0)$ be a probabilistic process, $\mathcal{T} = (T, Act, \mu_T, t_0, S)$ be a test and $\rho \in (0,1)$. Then the* interaction system *for $\mathcal{P}$ and $\mathcal{T}$ biased by ρ, $\mathcal{P}||_\rho\mathcal{T}$, has the form $(P \times T, Act, \mu_\rho, \langle\, p_0, t_0 \,\rangle)$ where μ_ρ is defined with respect to a normalization function $\nu : P \times T \to [0,1]$, with*

$$\nu(p,t) = (\sum_{a \in Act} \mu_P(p,a) * \mu_T(t,a)) + \mu_T(t,\tau) + \mu_P(p,\tau) - \mu_T(t,\tau) * \mu_P(p,\tau)$$

as follows:

$$\mu_\rho(\langle\, p,t \,\rangle, a, \langle\, p',t' \,\rangle) = \begin{cases} 0 & \text{if } \nu(p,t) = 0 \\ \frac{\mu_P(p,a,p') * \mu_T(t,a,t')}{\nu(p,t)} & \text{if } a \neq \tau \ \wedge \ \nu(p,t) \neq 0 \\ \frac{\mu_P(p,\tau,p') * (1-(1-\rho) * \mu_T(t,\tau))}{\nu(p,t)} & \text{if } a = \tau \ \wedge \ t = t' \ \wedge \ \nu(p,t) \neq 0 \\ \frac{\mu_T(t,\tau,t') * (1-\rho * \mu_P(p,\tau))}{\nu(p,t)} & \text{if } a = \tau \ \wedge \ p = p' \ \wedge \ \nu(p,t) \neq 0 \end{cases}$$

The normalization factor $\nu(p,t)$ records the sum of the probabilities of the transitions possible from state $\langle\, p,t \,\rangle$, and is used to compute the conditional probabilities of transitions from this interaction system state. It is interesting to note that the normalization factor is independent of the bias factor, ρ. In the definition of μ_ρ, the first clause pertains to a terminated interaction system state; the second clause computes the probability that the process and test synchronize on a non-τ action; the third clause computes the conditional probability that the process performs an autonomous τ action; and the fourth clause is similar to the third with the roles of p and t reversed.

Since tests are finite and processes are image-finite and divergence-free, interaction systems turn out to be finite. The normalization that is performed on the transition probabilities reflects our desire to model reactive processes. In a reactive system in which several actions (including, possibly, a τ) are enabled, if a test is capable of activating at least one of these actions (or a τ) with a certain probability then some interaction between the system and test is guaranteed to take place. Viewed in this light, the normalized probabilities represent the conditional probability of a particular interaction taking place, *given that some interaction takes place.* A by-product of the normalization is that interaction systems themselves are stochastic.

To clarify, consider the process $a[1]$ and the test $a[\frac{1}{3}]\,.\,good + b[\frac{2}{3}]\,.\,bad$ interacting with a bias factor of $\frac{1}{2}$. By our definition we end up with the interaction system that is isomorphic to $a[1]\,.\,good$; in other words, the process and test synchronize on a with probability 1. Without normalization, we would obtain a substochastic interaction system whose start state only performs an a-transition with probability $\frac{1}{3}$. As an illustration of the use of the bias factor in deciding whether the process or test does its τ action, consider the process $\tau[\frac{1}{2}].b + a[\frac{1}{2}]$ and the test $\tau[1].(a[\frac{1}{2}].good + b[\frac{1}{2}].good)$ interacting with a bias factor of $\frac{2}{3}$. The resulting interaction system is $\tau[\frac{1}{3}]\,.\,\tau\,.\,b\,.\,good + \tau[\frac{2}{3}]\,.\,(\tau[\frac{2}{3}]\,.\,b\,.\,good + a[\frac{1}{3}]\,.\,good)$. Even though the bias is towards the process doing its τ action if both want to do τ actions simultaneously, the probability that the test will do its τ action is still higher as represented by the right-hand term of the resulting interaction system. This is because the only action the test ever does initially is a τ.

As borne out by the definitions, processes, tests, and interaction systems are all objects of essentially the same type. This homogeneity, which is not present in the approaches to probabilistic testing found in [LS89, Chr90], mirrors that found in the original testing theory of De Nicola and Hennessy. One desirable consequence of this is that there is a strong connection between applying a test to a process and running a process in a network with other processes.

2.3 Testing Probabilistic Processes

We now formalize the notion of a process passing a test with a certain probability, and we use this idea to define a probabilistic testing preorder on processes.

Definition 5 *Let $\mathcal{P} = (P, Act, \mu_P, p_0)$ be a process, $\mathcal{T} = (T, Act, \mu_T, t_0, S)$ a test, and $\rho \in [0,1]$ with $\mathcal{P}||_\rho\mathcal{T}$ the associated interaction system.*

- *Each pair $\langle p, t \rangle \in P \times T$ is called 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*

$$\langle p, t \rangle = \langle p_0, t_0 \rangle \xrightarrow{a_1, \pi_1} \langle p_1, t_1 \rangle \xrightarrow{a_2, \pi_2} \ldots \xrightarrow{a_n, \pi_n} \langle p_n, t_n \rangle$$

  where

  - *for all $1 \leq i \leq n$, $\mu_\rho(\langle p_{i-1}, t_{i-1} \rangle, a_i, \langle p_i, t_i \rangle) = \pi_i$ (recall $\pi_i > 0$);*
  - *for all $0 \leq i \leq n-1$, $\langle p_i, t_i \rangle$ is unsuccessful,* i.e., *$t_i \notin S$; and*
  - *either $\langle p_n, t_n \rangle$ is successful or $\sum_{a \in Act \cup \{\tau\}} \mu_\rho(\langle p_n, t_n \rangle, a) = 0$.*

- *$\mathcal{C}_{p||_\rho t}$ is the set of all computations of the form $\langle p, t \rangle \xrightarrow{a,\pi} C$, $Comp = \bigcup_{p,t,\rho} \mathcal{C}_{p||_\rho t}$ is the set of all computations, and $\mathcal{S}_{p||_\rho t} \subseteq \mathcal{C}_{p||_\rho t}$ is the set of all successful computations from $\langle p, t \rangle$. $Comp$ is ranged over by C.*
- *The* computation probability density function *$Pr : Comp \to [0,1]$ is defined inductively as follows:*

$$Pr(\langle p, t \rangle) = 1$$
$$Pr(\langle p, t \rangle \xrightarrow{a,\pi} C) = \pi * Pr(C)$$

  If $\mathcal{C} \subseteq Comp$ then we define $Pr(\mathcal{C}) = \sum_{C \in Comp} Pr(C)$.

We now define the notion of passing a test with a particular probability and from this derive a testing preorder for probabilistic processes.

Definition 6 *Let $\mathcal{P} = (P, Act, \mu_P, p_0)$ and $\mathcal{P}' = (P', Act, \mu_{P'}, p_0')$ be processes, $\mathcal{T} = (T, Act, \mu_T, t_0, S)$ be a test, $\pi \in [0,1]$, and $\rho \in (0,1)$.*

- *$Succ(p, t, \rho) = Pr(\mathcal{S}_{p||_\rho t})$.*
- *$Succ(\mathcal{P}, \mathcal{T}, \rho) = Succ(p_0, t_0, \rho)$.*
- *p $pass_\pi$ t if $Succ(p, t, \rho) \geq \pi$.*
- *$\mathcal{P}$ $pass_\pi$ $\mathcal{T}$ if p_0 $pass_\pi$ t_0.*
- *$\mathcal{P} \sqsubseteq \mathcal{P}'$ if for all $\mathcal{T}$ and π, $\mathcal{P}$ $pass_\pi$ $\mathcal{T}$ implies $\mathcal{P}'$ $pass_\pi$ $\mathcal{T}$.*

Intuitively, $Succ(p, t, \rho)$ represents the probability that the process in state p will successfully complete the test in state t; p $pass_\pi$ t then holds if the probability of success is at least π, regardless of the bias factor. $\mathcal{P} \sqsubseteq \mathcal{P}'$ holds if, whenever $\mathcal{P}$ passes a test with a certain probability, then $\mathcal{P}'$ does as well. We also note the following.

Proposition 1 *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. If either $\sum_{p \in P} \mu_P(p, \tau) = 0$ or $\sum_{t \in T} \mu_T(t, \tau) = 0$, then $Succ(p, t, \rho) = Succ(p, t, \rho')$ for any $p \in P, t \in T$, and $\rho, \rho' \in (0, 1)$.*

Example 1

1. Let $\mathcal{P}_1 = a[\frac{1}{3}] \,.\, b + a[\frac{2}{3}] \,.\, c$ and $\mathcal{P}_1' = a[\frac{1}{2}] \,.\, b + a[\frac{1}{2}] \,.\, c$. Then $\mathcal{P}_1 \not\sqsubseteq \mathcal{P}_1'$ and $\mathcal{P}_1' \not\sqsubseteq \mathcal{P}_1$. To see this, consider the considering the tests $\mathcal{T}_1 = a \,.\, c \,.\, good$ and $\mathcal{T}_1' = a \,.\, b \,.\, good$. On the basis of the definition of $Succ$ and the previous proposition, it follows that $Succ(\mathcal{P}_1, \mathcal{T}_1, \rho) = \frac{2}{3}$, while $Succ(\mathcal{P}_1', \mathcal{T}_1, \rho) = \frac{1}{2}$, for any $\rho \in (0, 1)$. Thus $\mathcal{P}_1 \, pass_{\frac{2}{3}} \, \mathcal{T}_1$, while $\mathcal{P}_1' \, \not{p}ass_{\frac{2}{3}} \, \mathcal{T}_1$. Similarly, it follows that $\mathcal{P}_1' \, pass_{\frac{1}{2}} \, T_1'$, but $\mathcal{P}_1 \, \not{p}ass_{\frac{1}{2}} \, \mathcal{T}_1'$.

2. Let $\mathcal{P}_2 = a[\frac{1}{2}] \,.\, b \,.\, c + a[\frac{1}{2}] \,.\, b \,.\, d$ and $\mathcal{P}_2' = a \,.\, (b[\frac{1}{2}] \,.\, c + b[\frac{1}{2}] \,.\, d)$. Then $\mathcal{P}_2 = \mathcal{P}_2'$. It should be noted that these processes are distinguished by probabilistic bisimulation [LS89].

3. Let $\mathcal{P}_3 = a[\frac{1}{2}].b + a[\frac{1}{2}].c$ and $\mathcal{P}_3' = a.(b[\frac{1}{2}] + c[\frac{1}{2}])$. Then $\mathcal{P}_3 \not\sqsubseteq \mathcal{P}_3'$ and $\mathcal{P}_3' \not\sqsubseteq \mathcal{P}_3$. To see the former, consider the test $\mathcal{T}_3 = a \,.\, (b[\frac{1}{4}] \,.\, good + c[\frac{3}{4}] \,.\, bad)$. $Succ(\mathcal{P}_3, \mathcal{T}_3, \rho) = \frac{1}{2}$ for any rho, while $Succ(\mathcal{P}_3', \mathcal{T}_3', \rho) = \frac{1}{4}$, and thus $\mathcal{P}_3 \, pass_{\frac{1}{2}} \, \mathcal{T}_3$ while $\mathcal{P}_3' \, \not{p}ass_{\frac{1}{4}} \, \mathcal{T}_3$. For the latter, let $\mathcal{T}_3' = a \,.\, b \,.\, good$. Then $\mathcal{P}_3' \, pass_1 \, \mathcal{T}_3'$, but $\mathcal{P}_3' \, \not{p}ass_1 \, \mathcal{T}_3'$.

4. Let $\mathcal{P}_4 = \tau[\frac{1}{2}].b + a[\frac{1}{2}]$ and $\mathcal{P}_4' = a[\frac{1}{2}] + b[\frac{1}{2}]$. Then $\mathcal{P}_4 \not\sqsubseteq \mathcal{P}_4'$ and $\mathcal{P}_4' \not\sqsubseteq \mathcal{P}_4$. For the former, consider the test $\mathcal{T}_4 = a[\frac{1}{2}] \,.\, bad + b[\frac{1}{2}] \,.\, good$. It follows that $Succ(\mathcal{P}_4, \mathcal{T}_4, \rho) = \frac{2}{3}$ while $Succ(\mathcal{P}_4', \mathcal{T}_4, \rho) = \frac{1}{2}$ for any ρ; thus, $\mathcal{P}_4 \, pass_{\frac{2}{3}} \, \mathcal{T}_4$, while $\mathcal{P}_4' \, \not{p}ass_{\frac{2}{3}} \, \mathcal{T}_4$. For the latter, consider the test $\mathcal{T}_4' = a[\frac{1}{3}] \,.\, good + b[\frac{1}{3}] \,.\, bad + \tau[\frac{1}{3}] \,.\, good$. Since $Succ(\mathcal{P}_4', \mathcal{T}_4', \rho) = \frac{3}{4}$, it follows that $\mathcal{P}_4' \, pass_{\frac{3}{4}} \, \mathcal{T}_4'$. However, $Succ(\mathcal{P}_4, \mathcal{T}_4, \rho) = \frac{6-\rho}{10}$, which implies that if $\pi > \frac{1}{2}$, then $\mathcal{P}_4 \, \not{p}ass_\pi \, \mathcal{T}_4$.

5. Let $\mathcal{P}_5 = a \,.\, (\tau[\frac{1}{2}] \,.\, b + \tau[\frac{1}{2}] \,.\, c)$ and $\mathcal{P}_5' = a[\frac{1}{2}] \,.\, b + a[\frac{1}{2}] \,.\, c$. Then $\mathcal{P}_5 \sqsubseteq \mathcal{P}_5'$ but $\mathcal{P}_5' \not\sqsubseteq \mathcal{P}_5$. To see the latter, consider the test $\mathcal{T}_5' = a[1] \,.\, (\tau[\frac{1}{2}] \,.\, bad + b[\frac{1}{2}] \,.\, good)$. As $Succ(\mathcal{P}_5, \mathcal{T}_5', \rho) = \frac{1}{4}$ and $Succ(\mathcal{P}_5', \mathcal{T}_5', \rho) = \frac{1}{4} - \frac{1-\rho}{8}$, it follows that $\mathcal{P}_5' \, pass_{\frac{1}{4}} \, \mathcal{T}_5'$ and $\mathcal{P}_5 \, \not{p}ass_{\frac{1}{4}} \, \mathcal{T}_5'$. [DISCUSSION TO FOLLOW]

6. Let $\mathcal{P}_6 = a$ and $\mathcal{P}_6' = a[\frac{1}{2}] + \tau[\frac{1}{2}]$. Then $\mathcal{P}_6 \not\sqsubseteq \mathcal{P}_6'$ and $\mathcal{P}_6' \not\sqsubseteq \mathcal{P}_6$. For the former, consider the test $\mathcal{T}_6 = a[\frac{1}{2}] \,.\, bad + \tau[\frac{1}{2}] \,.\, good$. For all ρ it follows that $Succ(\mathcal{P}_6, \mathcal{T}_6, \rho) = \frac{3}{4}$ and $Succ(\mathcal{P}_6', \mathcal{T}_6, \rho) = \frac{1}{2}$; thus $\mathcal{P}_6 \, pass_{\frac{3}{4}} \, \mathcal{T}_6$, while $\mathcal{P}_6' \, \not{p}ass_{\frac{3}{4}} \, \mathcal{T}_6$. To see the latter, consider the test $\mathcal{T}_6' = a[\frac{1}{2}] \,.\, good + \tau[\frac{1}{2}] \,.\, bad$. In this case, $Succ(\mathcal{P}_6, \mathcal{T}_6', \rho) = \frac{1}{4}$, while $Succ(\mathcal{P}_6', \mathcal{T}_6', \rho) = \frac{1}{2}$, and it then follows that $\mathcal{P}_6' \, pass_{\frac{1}{2}} \, \mathcal{T}_6'$ while $\mathcal{P}_6 \, \not{p}ass_{\frac{1}{2}} \, \mathcal{T}_6'$.

3 Recovering Nonprobabilistic from Probabilistic Testing

We now establish that the probabilistic testing preorder $\sqsubseteq$ is a strict refinement of the classical testing preorder $\sqsubseteq_{test}$ of [DNH83, Hen88]. The proof relies on a correspondence between a probabilistic process passing a test with probability 1 (respectively a non-zero probability) and the associated nonprobabilistic process *must-passing* (respectively *may-passing*) the associated nonprobabilistic test.

We begin by introducing our renderings of the definitions of classical (nonprobabilistic) testing [DNH83, Hen88]. Although our framework differs slightly from theirs, the differences are shown to be inconsequential.

Definition 7 *A (nonprobabilistic) process is a quadruple $(P, Act, \rightarrow, p_0)$, where P is a set of states, $Act \cup \{\tau\}$ a set of actions, $\rightarrow \subseteq P \times Act \times P$ the* transition relation, *and $p_0 \in P$ the* start state. *We use* $\mathbf{P}$ *to denote the class of nonprobabilistic processes. A (nonprobabilistic) test is a quintuple $(T, Act, \rightarrow, t_0, S)$, where $(T, Act, \rightarrow, t_0) \in \mathbf{P}$ and $S \subseteq T$ is the set of* sucessful states.

As in our definition of probabilistic tests (Definition 2), our definition of nonprobabilistic tests differs slightly from that of [DNH83, Hen88]; there, a test is a process, and the successful states are defined to be those capable of performing a special "success action," w. This difference, however, is minor. In what follows, we write $p \xrightarrow{a}$ (resp. $p \not\xrightarrow{a}$) if there exists (resp. does not exist) a p' with $p \xrightarrow{a} p'$. We also write $p \not\rightarrow$ if $p \not\xrightarrow{a}$ holds for all $a \in Act \cup \{\tau\}$.

One result of DeNicola and Hennessy is that it suffices to consider only finite tests when defining the classical testing preorders. Accordingly, we only consider such tests in what follows. We now define what it means for a nonprobabilistic process to pass a test. We do so by introducing (nonprobabilistic) interaction systems, which record the computations in which a process may engage in response to a test.

Definition 8 *Let $\mathcal{P} = (P, Act, \rightarrow_P, p_0)$ be a nonprobabilistic process and $\mathcal{T} = (T, Act, \rightarrow_T, t_0, S)$ be a nonprobabilistic test. Then the nonprobabilistic interaction system $\mathcal{P}||\mathcal{T}$ is a quadruple $(P \times T, Act, \rightarrow, \langle p_0, t_0 \rangle)$, where $\rightarrow \subseteq (\mathcal{P} \times \mathcal{T}, Act \cup \{\tau\}, \mathcal{P} \times \mathcal{T})$ is defined as follows.*

- $t \xrightarrow{\tau}_T t' \;\Rightarrow\; \langle p, t \rangle \xrightarrow{\tau} \langle p, t' \rangle$.
- $p \xrightarrow{\tau}_P p' \;\Rightarrow\; \langle p, t \rangle \xrightarrow{\tau} \langle p', t \rangle$.
- $p \xrightarrow{a}_P p', t \xrightarrow{a}_T t', a \neq \tau \;\Rightarrow\; \langle p, t \rangle \xrightarrow{a} \langle p', t' \rangle$.

Definition 9 *Let $\mathcal{P} = (P, Act, \rightarrow_P, p_0)$ be a nonprobabilistic process and $\mathcal{T} = (T, Act, \rightarrow_T, t_0)$ be a nonprobabilistic test, with $\mathcal{P}||\mathcal{T}$ the associated interaction system.*

- Configurations *and* successful configurations *are defined as in Definition 5.*
- *A* computation *from a configuration $\langle p, t \rangle$ is a (maximal) sequence of the form*

$$\langle p, t \rangle = \langle p_0, t_0 \rangle \xrightarrow{a_1} \langle p_1, t_1 \rangle \xrightarrow{a_2} \ldots \xrightarrow{a_n} \langle p_n, t_n \rangle$$

  where

  - *for all $0 \leq i \leq n-1$, $\langle p_i, t_i \rangle$ is unsuccessful; and*
  - *either $\langle p_n, t_n \rangle$ is successful or $\langle p_n, t_n \rangle \not\rightarrow$.*

- *$\mathcal{C}_{p||t}$, $\mathcal{C}$, and $\mathcal{S}_{p||t}$ are defined as in Definition 5 except that the bias factor is no longer necessary.*
- *$\mathcal{P}$ may $\mathcal{T}$ if $S_{p_0||t_0} \neq \emptyset$.*
- *$\mathcal{P}$ must $\mathcal{T}$ if $S_{p_0||t_0} = C_{p_0||t_0}$.*

Definition 10 *Let $\mathcal{P}, \mathcal{P}' \in \mathbf{P}$.*

- *$\mathcal{P} \sqsubseteq_{may} \mathcal{P}'$ if for all $\mathcal{T}, \mathcal{P}$ may $\mathcal{T}$ implies $\mathcal{P}'$ may $\mathcal{T}$.*

- $\mathcal{P} \sqsubseteq_{must} \mathcal{P}'$ *if for all* $\mathcal{T}, \mathcal{P}$ *must* $\mathcal{T}$ *implies* $\mathcal{P}'$ *must* $\mathcal{T}$.
- $\mathcal{P} \sqsubseteq_{test} \mathcal{P}'$ *if* $\mathcal{P} \sqsubseteq_{may} \mathcal{P}'$ *and* $\mathcal{P} \sqsubseteq_{must} \mathcal{P}'$.

In the remainder of this section, we show that $\sqsubseteq$ refines $\sqsubseteq_{test}$. We begin by establishing that Pr (Definition 5) is indeed a probability distribution function for all maximal computations in a probabilistic interaction system.

Lemma 1 *Let* $\mathcal{P} = (P, Act, \mu_P, p_0)$, $\mathcal{T} = (T, Act, \mu_T, t_0)$ *and* $\rho \in (0,1)$. *Then* $\sum_{C \in \mathcal{C}_{p_0 ||_\rho t_0}} Pr(C) = 1$.

Proof: We prove something stronger: for all $\langle\, p, t\,\rangle \in P \times T$ and all $\rho \in (0,1)$, $\sum_{C \in \mathcal{C}_{p ||_\rho t}} Pr(C) = 1$. The proof is by induction on the length of the longest computation in $\mathcal{C}_{p_0 ||_\rho t_0}$. $\square$

We next define the function h that maps probabilistic processes to nonprobabilistic ones essentially by erasing the probabilities on transitions.

Definition 11 *The mapping* $h : \mathbf{PP} \rightarrow \mathbf{P}$ *is defined by* $h((P, Act, \mu_P, p_0)) = (P, Act, \rightarrow_P, p_0)$, *where* $(p, a, p') \in \rightarrow_P$ *if* $\mu_P(p, a, p') > 0$.

Note that we abuse notation slightly by applying h to tests in the obvious way: $h((T, Act, \mu_T, t_0, S)) = (T, Act, \rightarrow_T, t_0, S)$. The next lemma establishes that h is a homomorphism with respect to the interaction system constructors $||_\rho, ||$.

Lemma 2 *Let* $\mathcal{P}$ *be a probabilistic process,* $\mathcal{T}$ *a test, and* $\rho \in (0,1)$. *Then* $h(\mathcal{P} ||_\rho \mathcal{T}) = h(\mathcal{P}) \,||\, h(\mathcal{T})$.

Proof: Let $\langle\, p, t\,\rangle$ be a state in $h(\mathcal{P}||_\rho\mathcal{T})$ and $a \neq \tau$. $\langle\, p, t\,\rangle \xrightarrow{a} \langle\, p', t'\,\rangle$ iff $\mu_\rho(\langle\, p, t\,\rangle, a, \langle\, p', t'\,\rangle) > 0$ iff $\mu_P(p, a, p') > 0$ and $\mu_T(t, a, t') > 0$ iff $p \xrightarrow{a} p'$ and $t \xrightarrow{a} t'$ iff $\langle\, p, t\,\rangle \xrightarrow{a} \langle\, p', t'\,\rangle$ in $h(\mathcal{P})||h(\mathcal{T})$. The proof is similar for τ actions. $\square$

The proof that $\sqsubseteq$ implies $\sqsubseteq_{test}$ is broken into several steps. We first establish a relationship between $pass_\pi$ and may and $must$. We then show that $\sqsubseteq$ implies both $\sqsubseteq_{may}$ and $\sqsubseteq_{must}$. From this we obtain the result that the probabilistic testing preorder is a refinement of the nonprobabilistic testing preorder. In the following lemmas, let $\mathcal{P}$,$\mathcal{P}'$ be probabilistic processes and $\mathcal{T}$ a probabilistic test.

Lemma 3 *Let* $\mathcal{P}$ *and* $\mathcal{T}$ *be a probabilistic process and test, respectively.*

- $\mathcal{P}\ pass_\pi\ \mathcal{T}$ *for some* $\pi > 0$ *if and only if* $h(\mathcal{P})\ may\ h(\mathcal{T})$.
- $\mathcal{P}\ pass_1\ \mathcal{T}$ *if and only if* $h(\mathcal{P})\ must\ h(\mathcal{T})$.

This lemma follows directly from Lemmas 1 and 2. Using it, we may now prove the following.

Lemma 4 *Let* $\mathcal{P}$ *and* $\mathcal{P}'$ *be probabilistic processes.*

- $\mathcal{P} \sqsubseteq \mathcal{P}'$ *implies* $h(\mathcal{P}) \sqsubseteq_{may} h(\mathcal{P}')$.
- $\mathcal{P} \sqsubseteq \mathcal{P}'$ *implies* $h(\mathcal{P}) \sqsubseteq_{must} h(\mathcal{P}')$.

The next theorem is a trivial consequence of this lemma.

Theorem 1 *$\mathcal{P} \sqsubseteq \mathcal{P}'$ implies $h(\mathcal{P}) \sqsubseteq_{test} h(\mathcal{P}')$.*

Note that the converse of Theorem 1 does not hold. For example, consider once again processes $\mathcal{P}_1 = a[\frac{1}{3}] \,.\, b + a[\frac{2}{3}] \,.\, c$ and $\mathcal{P}_1' = a[\frac{1}{2}] \,.\, b + a[\frac{1}{2}] \,.\, c$ of Example 1: $h(\mathcal{P}_1) =_{test} h(\mathcal{P}_1')$ but $\mathcal{P}_1 \not\sqsubseteq \mathcal{P}_1'$ and $\mathcal{P}_1' \not\sqsubseteq \mathcal{P}_1$. Thus, the testing theory of nonprobabilistic processes is strictly more abstract than the testing theory of probabilistic processes.

4 Toward an Alternative Characterization

While intuitively appealing, the operational definition of our preorder 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, in this section we propose an alternative characterization of the preorder that is designed to ease the task of establishing relationships between systems. The characterization also has a denotational flavor that may be useful in future developments of denotational models.

The characterization is reminiscent of De Nicola's account of the testing preorder [DN87] in that it relates processes on the basis of the sets of actions they can execute after engaging in a sequence of visible actions (where arbitrary amounts of internal computation may be performed in between). In addition, we also take account of the actions that a process *refuses* after a sequence.

The alternative characterization now relies on the two functions $\mathcal{A}_\mathcal{P}$ and $\mathcal{R}_\mathcal{P}$.

Definition 12 *Let $\mathcal{P} = (P, Act, \mu_P, p_0)$ be a probabilistic process. Then $\mathcal{A}_\mathcal{P}, \mathcal{R}_\mathcal{P} : P \times Act^* \times 2^{Act} \times 2^{Act} \to [0,1]$ may be given inductively as follows.*

$$
\begin{aligned}
\mathcal{A}_\mathcal{P}(p,\epsilon,S) &= \begin{cases} 1 & \textit{if } \exists a \in S, \pi > 0. p \xrightarrow{a,\pi} \\ 0 & \textit{otherwise} \end{cases} \\
\mathcal{A}_\mathcal{P}(p,as,S) &= \begin{cases} 1 & \textit{if } \not\exists \pi > 0. p \xrightarrow{a,\pi} \\ \frac{1}{\mu_P(p,a)} * \sum_{p' \in P} \mu_P(p,a,p') * \mathcal{A}_\mathcal{P}(p',a,S) & \textit{otherwise} \end{cases} \\
\mathcal{R}_\mathcal{P}(p,\epsilon,S) &= 1 - \mathcal{A}_\mathcal{P}(p,\epsilon,S) \\
\mathcal{R}_\mathcal{P}(p,as,S) &= \begin{cases} 1 & \textit{if } \not\exists \pi > 0. p \xrightarrow{a,\pi} \\ \frac{1}{\mu_P(p,a)} * \sum_{p' \in P} \mu_P(p,a,p') * \mathcal{R}_\mathcal{P}(p',s,S) & \textit{otherwise} \end{cases}
\end{aligned}
$$

Intuitively, $\mathcal{A}_\mathcal{P}(p,s,S)$ and $\mathcal{R}_\mathcal{P}(p,s,S)$ record the probability that, after engaging in sequence s from state p, the process $\mathcal{P}$ can be in states in which some action in S is possible and every action in S is refuSed. We think of these two functions as the *acceptances* and *refusals* respectively. We now proceed to define the alternative characterization.

Definition 13 *Let $\mathcal{P}, \mathcal{Q} \in$ **PP**, with start states p_0 and q_0, respectively. Then $\mathcal{P} \ll \mathcal{Q}$ if for all $s \in Act \cup \{\tau\}^*$ and finite $S \subseteq Act$, $\mathcal{A}_\mathcal{P}(p_0,s,S) \le \mathcal{A}_\mathcal{Q}(q_0,s,S)$ and $\mathcal{R}_\mathcal{P}(p_0,s,S) \le \mathcal{R}_\mathcal{Q}(q_0,s,S)$.*

We now conjecture the following.

Conjecture 1 *Let $\mathcal{P}, \mathcal{Q} \in$ **PP**. Then $\mathcal{P} \ll \mathcal{Q}$ if and only if $\mathcal{P} \sqsubseteq \mathcal{Q}$.*

We have proven the $\Leftarrow$ direction of the conjecture by giving a technique for building tests from s and S in such a way that for some monotonic function f and state p, $\mathcal{A}_{\mathcal{P}}(p, s, S) = \pi$ if and only if $p\, pass_{f(\pi)}\, t$ and . Such a test can be constructed using a technique similar to De Nicola's for nonprobabilistic processes [DN87]. We have not yet been able to prove (or disprove) the $\Rightarrow$ direction. Our proposed proof technique involves analyzing the computations of $\mathcal{P}||\mathcal{T}$ for a test $\mathcal{T}$ to show that the associated probabilities correspond to certain values of $\mathrm{AC}_{\mathcal{P}}$. The negative information in S_2 is important for showing that normalization factors at various stages in the computations of $\mathcal{P}||\mathcal{T}$ and $\mathcal{Q}||\mathcal{T}$ are related appropriately.

We end this section with an example.

Example 2 *Let* $\mathcal{P} = a[\frac{1}{2}]\,.\,b + a[\frac{1}{2}]\,.\,c$ *and* $\mathcal{Q} = a\,.\,(b[\frac{1}{2}] + c[\frac{1}{2}])$*. The pair* $\langle\, a, \{b\}\,\rangle$ *is sufficient for establishing* $\mathcal{P} \not\ll \mathcal{Q}$ *and* $\mathcal{Q} \not\ll \mathcal{P}$ *since* $\mathcal{A}_{\mathcal{P}}(\mathcal{P}, a, \{b\}) = \frac{1}{2}$ *and* $\mathcal{R}_{\mathcal{P}}(\mathcal{P}, a, \{b\}) = \frac{1}{2}$ *but* $\mathcal{A}_{\mathcal{Q}}(\mathcal{Q}, a, \{b\}) = 1$ *while* $\mathcal{R}_{\mathcal{Q}}(\mathcal{Q}, a, \{b\}) = 0$*.*

5 Substochastic Processes

In this section we investigate the family of *substochastic* probabilistic processes. A substochastic process has the property that the sum of the probabilities of its outgoing transitions may be less than 1. The "missing" part of the probability distribution can be viewed as a measure of the undefined behavior of the process. For example, the behavior of process $\mathcal{P} = a[\frac{1}{4}]\,.\,b + a[\frac{1}{2}]\,.\,c$ is undefined with probability $\frac{1}{4}$ in its initial state.

Substochastic processes prove useful in the modeling of fault-tolerant systems and in the analysis of system reliability. System faults and unreliability can be captured by a process exhibiting undefined behavior, and the probabilistic testing preorder we define in this section will relate substochastic processes $\mathcal{P}$ and $\mathcal{Q}$ when $\mathcal{P}$ is less reliable than $\mathcal{Q}$.

In order to introduce substochasticness into our framework, we augment the definition of a probabilistic process with a function δ that represents the probability that a process will fail in a particular state. So, $\delta(p) = \pi$ means that, in state p, the behavior of the process will be undefined with probability π. The formal definition of substochastic processes is as follows.

Definition 14 *A* substochastic process *is a 5-tuple* $(P, Act, \delta, \mu, p_0)$ *where*

- P*,* Act *and* p_0 *are as in Definition 1;*
- $\delta : P \to [0, 1]$ *is the* undefinedness distribution*;*
- $\mu : P \times (Act \cup \{\,\tau\,\}) \times P \to [0, 1]$ *is the* transition probability distribution function *such that, for all* $p \in P$*,* $(\sum_{a \in Act} \mu(p, a)) + \delta(p) \in \{\,0, 1\,\}$

Substochastic processes will be subjected to *stochastic* tests; that is, tests are considered reliable. In general, this restriction seems reasonable, since if a process fails an unreliable test, it would not in general be possible to identify the source of unreliability in the interaction. Nonetheless, our theory could be modified in a straightforward way to accommodate substochastic tests.

The definition of an interaction system for substochastic processes is the same as Definition 4, except for two points: (1) The normalization factor $\nu(p, t)$, used to compute the probabilities of transitions emanating from configuration $\langle\, p, t\,\rangle$, contains the additional term $\delta(p)$; i.e.,

$$\nu(p, t) = (\sum_{a \in Act} \mu_P(p, a) * \mu_T(t, a)) + \mu_P(p, \tau) + \mu_T(t, \tau) - \mu_P(p, \tau) * \mu_T(t, \tau) + \delta_P(p)$$

As such, the normalized transition probabilities will reflect the measure of substochasticness present in the original process, and thus the interaction system itself will be substochastic. (2) Interaction systems for substochastic processes are equipped with an undefinedness distribution function δ_I, which is the δ function of the process normalized by $\nu(p,t)$ defined above.

Given these notions, the definitions of $pass_\pi$ and the probabilistic testing preorder do not change from those given earlier for stochastic processes (Definition 6). The following examples illustrate the relationships induced by the probabilistic testing preorder on substochastic processes.

Example 3

1. Let $\mathcal{P}_1 = a[\frac{3}{4}]$ and $\mathcal{P}'_1 = a$. Then $\mathcal{P}_1 \sqsubset_{\frac{1}{2}} \mathcal{P}'_1$, i.e., $\mathcal{P}_1 \sqsubseteq \mathcal{P}'_1$ and $\mathcal{P}'_1 \not\sqsubseteq \mathcal{P}_1$, as $\mathcal{P}_1$ is more substochastic than $\mathcal{P}'_1$. To see this one need only consider the test $a\,.\,good$, which $\mathcal{P}_1$ passes with probability $\frac{3}{4}$ and $\mathcal{P}'_1$ passes with probability 1, when both are biased with a factor of $\frac{1}{2}$
2. For a more complex example, let $\mathcal{P}_2 = a[\frac{1}{4}]\,.\,b[\frac{1}{2}] + a[\frac{1}{4}]\,.\,c[\frac{1}{2}]$ and $\mathcal{P}'_2 = a[\frac{1}{4}]\,.\,b + a[\frac{1}{4}]\,.\,c$. Then $\mathcal{P}_2 \sqsubset \mathcal{P}'_2$
3. Let $\mathcal{P}_3 = a[\frac{1}{4}]\,.\,b[\frac{1}{2}] + a[\frac{1}{4}]\,.\,c[\frac{1}{2}]$ and $\mathcal{P}'_3 = a\,.\,(b[\frac{1}{2}] + c[\frac{1}{2}])$. Then $\mathcal{P}_3$ and $\mathcal{P}'_3$ are unrelated by probabilistic testing. The test and bias factor which prevents $\mathcal{P}_3 \sqsubseteq \mathcal{P}'_3$ are $\mathcal{T} = a\,.\,(b[\frac{1}{16}]\,.\,good + c[\frac{15}{16}]\,.\,bad)$ and $\frac{1}{2}$ because $\mathcal{P}_3\, pass_{\frac{1}{8}}\, \mathcal{T}$ but $\mathcal{P}'_3\, pass_{\frac{1}{32}}\, \mathcal{T}$. Thus, even for substochastic processes, a sufficiently biased test can be found which prevents the more deterministic process from being greater than the less deterministic one.

Substochasticness and Reliability

To illustrate the utility of our testing theory for substochastic processes, we present a simple example of how to model unreliable system components. The example also illustrates what it means for one system to be more reliable (or unreliable) than another in our testing preorder.

We first define a probabilistic *chaining* operator that permits probabilistic processes to be linked together into networks. This requires some notation. If $S \subseteq Act_P \times Act_Q$ is a set of pairs, then the first and second projection operators, $(\cdot)_1$ and $(\cdot)_2$, respectively, are defined in the usual way:

$$\begin{aligned}(S)_1 &= \{\, a \in Act_P \mid \exists b \in Act_Q . \langle\, a, b\,\rangle \in S\,\}\\ (S)_2 &= \{\, b \in Act_Q \mid \exists a \in Act_P . \langle\, a, b\,\rangle \in S\,\}.\end{aligned}$$

The ordered pairs in S will be used to specify which pairs of ports are to be linked in the chaining operation defined next.

Definition 15 *Let $\mathcal{P} = (P, Act_P, \delta_P, \mu_P, p_0)$ and $\mathcal{Q} = (Q, Act_Q, \delta_Q, \mu_Q, q_0)$ be substochastic processes. Also let $S \subseteq Act_P \times Act_Q$ and $0 < \pi < 1$. Then $\mathcal{P}|S|_\pi\mathcal{Q}$ is the substochastic process $(P \times Q, Act, \mu_{||}, \langle\, p_0, q_0\,\rangle, \delta_{||})$, where $Act = (Act_P - (S)_1) \cup (Act_Q - (S)_2)$ and where $\delta_{||}$ and $\mu_{||}$ are defined with respect to a normalization factor*

$$\begin{aligned}\nu_{||}(p,q) = & \sum_{\langle\, p',q'\,\rangle \in P\times Q} Syn(\langle\, p,q\,\rangle, \langle\, p',q'\,\rangle) \;+\\ & \sum_{a\in Act\cup\{\tau\}, \langle\, p',q'\,\rangle \in P\times Q} (ASyn_1(\langle\, p,q\,\rangle, a, \langle\, p',q'\,\rangle) + ASyn_2(\langle\, p,q\,\rangle, a, \langle\, p',q'\,\rangle)) \;+\\ & \delta_P(p) + \delta_Q(q) - \delta_P(p) * \delta_Q(q)\end{aligned}$$

where

$$Syn(\langle\, p, q\,\rangle, \langle\, p', q'\,\rangle) = \sum_{\langle\, a,b\,\rangle \in S} \mu_P(p, a, p') * \mu_Q(q, b, q')$$

and

$$ASyn_1(\langle\, p, q\,\rangle, a, \langle\, p', q'\,\rangle) = \begin{cases} \mu_P(p, a, p') * \mu_Q(q, (Act_Q - (S)_2) \cup \{\tau\}) * \pi + & \\ \quad \mu_P(p, a, p') * \mu_Q(q, (S)_2) & \textit{if } q = q' \textit{ and } a \notin (S)_1 \\ 0 & \textit{otherwise} \end{cases}$$

$$ASyn_2(\langle\, p, q\,\rangle, a, \langle\, p', q'\,\rangle) = \begin{cases} \mu_Q(q, a, q') * \mu_P(p, (Act_P - (S)_1) \cup \{\tau\}) * (1 - \pi) + & \\ \quad \mu_Q(q, a, q') * \mu_P(p, (S)_1) & \textit{if } p = p' \textit{ and } a \notin (S)_2 \\ 0 & \textit{otherwise} \end{cases}$$

as follows:

$$\delta_{||}(\langle\, p, q\,\rangle) = \frac{\delta_P(p) + \delta_Q(q) - \delta_P(p) * \delta_Q(q)}{\nu_{||}(p, q)}$$

$$\mu(p|S|_\pi q, a, p'|S|_\pi q') = \begin{cases} 0 & \textit{if } \nu(p|S|_\pi q) = 0 \\ \frac{\mu(p,a,p') * \mu(q,a,q')}{\nu_{||}(p|S|_\pi q)} & \textit{if } a \in S \\ \frac{(p,a,p') * (1-(1-\pi) * \mu(q,(Act \cup \{\tau\}) - S))}{\nu_{||}(p|S|_\pi q)} & \textit{if } a \notin S \textit{ and } q = q' \\ \frac{\mu(q,a,q') * (1-\pi * \mu(p,(Act \cup \{\tau\}) - S))}{\nu_{||}(p|S|_\pi q)} & \textit{if } a \notin S \textit{ and } p = p' \end{cases}$$

Intuitively, in $\mathcal{P}|S|_\pi\mathcal{Q}$, $\mathcal{P}$ and $\mathcal{Q}$ execute "in parallel" in an interleaved fashion, with synchronization required on the pairs of ports mentioned in S. A synchronization results in a τ. Note that Syn calculates the unnormalized probabilities of $\mathcal{P}$ and $\mathcal{Q}$ synchronizing on a pair in S, while $ASyn_1$ and $ASyn_2$ compute the unnormalized probabilities of the indicated asynchronous transitions for the first and second component, respectively. The normalization factor follows from the fact that we want the operator to have the property that, if any interaction is possible, then some interaction will take place (modulo undefinedness).

The formal complexity of the definition belies the relatively simple "coin flipping" intuition underlying the semantics of this chaining operator, which is given as follows. Assume that $\mathcal{P}$ initially has m transitions and let $l_P = 1$ if $\delta(p_0) > 0$ and 0 otherwise. Likewise, assume that $\mathcal{Q}$ initially has n transitions and let l_Q be defined similarly to l_P. To determine the transitions for $\mathcal{P}|S|_\pi\mathcal{Q}$, we flip an $(m + l_P)$-sided coin (weighted according to the relevant probabilities) for $\mathcal{P}$ and a similarly weighted $(n + l_Q)$-sided coin for $\mathcal{Q}$. If either coin indicates that $\mathcal{P}$ or $\mathcal{Q}$ is undefined, then so is $\mathcal{P}|S|_\pi\mathcal{Q}$. If the labels of the indicated transitions (with the label for $\mathcal{P}$ being ordered before the label for $\mathcal{Q}$) constitute a pair in S, then $\mathcal{P}|S|_\pi\mathcal{Q}$ executes a τ, and both $\mathcal{P}$ and $\mathcal{Q}$ change state. Otherwise, if one of $\mathcal{P}$'s and $\mathcal{Q}$'s coins (but not both) indicates a transition whose label is not mentioned in S, then the associated component ($\mathcal{P}$ or $\mathcal{Q}$) engages asynchronously in this transition while the other component remains idle. Finally, if both $\mathcal{P}$'s and $\mathcal{Q}$'s coins indicate that each process wishes to engage in an asynchronous transition, then a two-sided coin, with faces weighted by π (for $\mathcal{P}$) and $1 - \pi$ (for $\mathcal{Q}$) is flipped to determine which process executes its transition. If none of these cases holds, the coins for $\mathcal{P}$ and $\mathcal{Q}$ are flipped again. If $\mathcal{P}$ and $\mathcal{Q}$ are capable of any interaction (synchronous or asynchronous), then this procedure is guaranteed, with probability 1, to terminate with some action being performed.

We note here that $|S|_\pi$ could be defined in terms of more primitive operators (such as probabilized counterparts of hiding, restriction, renaming, and CSP parallel composition). To illustrate the utility of our testing theory for substochastic processes, we consider, in Figure 1a, two unreliable 1-place buffers connected by the chaining operator $|\{\,\langle\, out, in\,\rangle\,\}|_{\frac{1}{2}}$. Each process may fail to

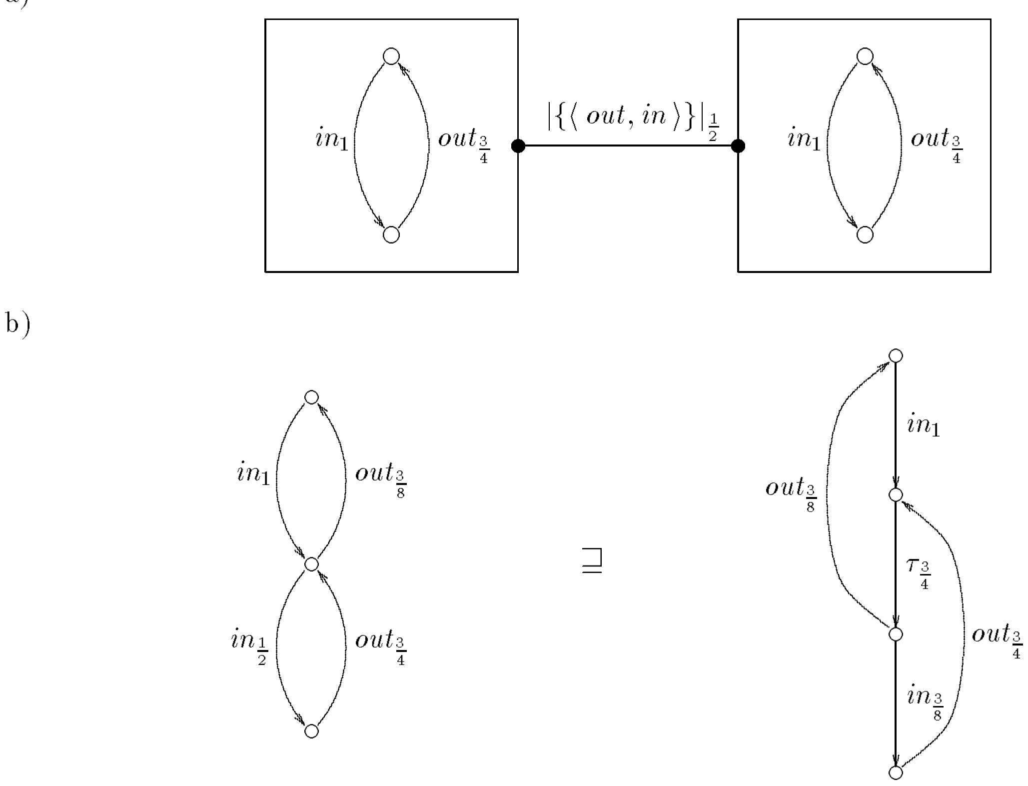

Figure 1: Example of substochastic processes

output the contents of its buffer with probability $\frac{1}{4}$. In the left-hand side of Figure 1b we have an unreliable 2-place buffer which, like the 1-place buffers, one quarter of the time fails to produce output when an *out*-transition is attempted. (Technically, $\delta(p) = \frac{1}{8}$ in the middle state as, ideally, the *out*-transition should occur with probability equal to that of the *in*-transition.) Finally, the right-hand side of Figure 1b depicts the composite process resulting from the chaining of the two 1-place buffers in Figure 1a. The τ-transition is a result of the synchronization resulting when the left-hand process passes along its buffer contents to the right-hand process. Note that the probabilities associated with the actions *in* and *out* in third state from the top are both $\frac{3}{8}$. This results from the undefinedness associated with the state from which *out* can occur.

According to our probabilistic testing preorder we have that the 2-place buffer is more reliable than the chained 1-place buffers. For example, the 2-place buffer passes the test $\mathcal{T} = in\,.\,out\,.\,good$ with probability $\frac{3}{4}$ while the chained buffer passes $\mathcal{T}$ with probability only $\frac{9}{20}$. This difference in behavior can be attributed to the extra communication (and concomitant possibility for failure) inherent to the chained architecture.

6 Conclusions

In this paper we have developed a testing theory for probabilistic processes based on the notion of a process passing a test with a certain probability. Close connections with the testing theory of De Nicola and Hennessy have been demonstrated, providing a well-established foundation for our approach. By extending our theory to substochastic processes, we have obtained techniques for modeling faulty systems and for distinguishing processes based on their degree of reliability.

Several topics for future study can be identified. Modifying the theory to accommodate processes with arbitrary τ transitions would likely necessitate the use of measure theory to handle infinite computations in interaction systems appropriately. The same can also be said of allowing infinite (as opposed to finite) tests. Other interesting topics to pursue are the integration of real-time into our probabilistic testing preorder (the work reported in [CZ91] is a natural candidate); and axiomatizations of our preorders with respect to a probabilistic process algebra.

References

[BM89] B. Bloom and A. R. Meyer. A remark on bisimulation between probabilistic processes. In Meyer and Tsailin, editors, *Logik at Botik*, Springer-Verlag, 1989.

[Chr90] I. Christoff. *Testing Equivalences for Probabilistic Processes.* Technical Report DoCS 90/22, Ph.D. Thesis, Department of Computer Science, Uppsala University, Uppsala, Sweden, 1990.

[CZ91] R. Cleaveland and A. Zwarico. A theory of testing for real-time. In *Proceedings of the Sixth IEEE Symposium on Logic in Computer Science*, July 1991.

[DN87] R. De Nicola. Extensional equivalences for transition systems. *Acta Informatica*, 24:211–237, 1987.

[DNH83] R. De Nicola 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, Boston, Mass., 1988.

[JS90] C.-C. Jou and S. A. Smolka. Equivalences, congruences, and complete axiomatizations for probabilistic processes. In J. C. M. Baeten and J. W. Klop, editors, *Proceedings of CONCUR '90*, pages 367–383, Springer-Verlag, Berlin, 1990.

[LS89] K. G. Larsen and A. Skou. Bisimulation through probabilistic testing. In *Proceedings of 16th Annual ACM Symposium on Principles of Programming Languages*, 1989.

[Mil80] R. Milner. *A Calculus of Communicating Systems.* Volume 92 of *Lecture Notes in Computer Science*, Springer-Verlag, Berlin, 1980.

[Mil89] R. Milner. *Communication and Concurrency. International Series in Computer Science*, 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.