Testing Finitary Probabilistic Processes
(Extended Abstract)

Yuxin Deng[1] ⋆, Rob van Glabbeek[2,4], Matthew Hennessy[3] ⋆⋆ & Carroll Morgan[4] ⋆⋆⋆

[1] Shanghai Jiao Tong University, China
[2] NICTA, Sydney, Australia
[3] Trinity College Dublin, Ireland
[4] University of New South Wales, Sydney, Australia

Abstract. We provide both modal- and relational characterisations of may- and must-testing preorders for recursive CSP processes with divergence, featuring probabilistic as well as nondeterministic choice. May testing is characterised in terms of simulation, and must testing in terms of failure simulation. To this end we develop weak transitions between probabilistic processes, elaborate their topological properties, and express divergence in terms of partial distributions.

1 Introduction

It has long been a challenge for theoretical computer scientists to provide a firm mathematical foundation for process-description languages that incorporate both nondeterministic and probabilistic behaviour in such a way that processes are semantically distinguished just when they can be told apart by some notion of testing.

In our earlier work [3, 1] a semantic theory was developed for one particular language with these characteristics, a finite process calculus called pCSP: nondeterminism is present in the form of the standard choice operators inherited from CSP [7], that is $P \sqcap Q$ and $P \,\Box\, Q$, while probabilistic behaviour is added via a new choice operator $P \;{}_p\!\oplus\; Q$ in which P is chosen with probability p and Q with probability $1{-}p$. The intensional behaviour of a pCSP process is given in terms of a probabilistic labelled transition system [14, 3], or pLTS, a generalisation of labelled transition systems [12]. In a pLTS the result of performing an action in a given state results in a *probability distribution* over states, rather than a single state; thus the relations $s \xrightarrow{\alpha} t$ in an LTS are replaced by relations $s \xrightarrow{\alpha} \Delta$, with Δ a distribution. Closed pCSP expressions P are interpreted as probability distributions $[\![P]\!]$ in the associated pLTS. Our semantic theory [3, 1] naturally generalises the two preorders of standard testing theory [5] to pCSP:

- $P \sqsubseteq_{\text{pmay}} Q$ indicates that Q is at least as good as P from the point of view of *possibly* passing probabilistic tests; and
- $P \sqsubseteq_{\text{pmust}} Q$ indicates instead that Q is at least as good as P from the point of view of *guaranteeing* the passing of probabilistic tests.

⋆ Deng was supported by the National Natural Science Foundation of China (60703033).
⋆⋆ Hennessy gratefully acknowledges the financial support of Science Foundation Ireland.
⋆⋆⋆ Morgan acknowledges the support of ARC Discovery Grant DP0879529.

The most significant result of [1] was an alternative characterisation of these preorders as particular forms of coinductively defined *simulation* relations, $\sqsubseteq_S$ and $\sqsubseteq_{FS}$, over the underlying pLTS. We also provided a characterisation in terms of a modal logic.

The object of the current paper is to extend the above results to a version of pCSP with recursive process descriptions: we add a construct $\mathsf{rec}\,x.\,P$ for recursion, and extend the intensional semantics of [1] in a straightforward manner. We restrict ourselves to *finitary* pCSP processes, those having finitely many states and displaying finite branching.

The simulation relations $\sqsubseteq_S$ and $\sqsubseteq_{FS}$ in [1] were defined in terms of weak transitions $\stackrel{\hat{\tau}}{\Longrightarrow}$ between distributions, obtained as the transitive closure of a relation $\stackrel{\hat{\tau}}{\longrightarrow}$ between distributions that allows one part of a distribution to make a τ-move with the other part remaining in place. This definition is however inadequate for processes that can do an unbounded number of τ-steps. The problem is highlighted by the process $Q_1 = \mathsf{rec}\,x.\,(\tau.x\ {}_{\frac{1}{2}}\!\oplus\ a.\,\mathbf{0})$ illustrated in Figure 1(a). Process Q_1 is indistinguishable, using tests, from the simple process $a.\,\mathbf{0}$: we have $Q_1 \simeq_{\text{pmay}} a.\,\mathbf{0}$ and $Q_1 \simeq_{\text{pmust}} a.\,\mathbf{0}$. This is because the process Q_1 will eventually perform the action a with probability 1. However, the action $[\![a.\,\mathbf{0}]\!] \xrightarrow{a} [\![\mathbf{0}]\!]$ can not be simulated by a corresponding move $[\![Q_1]\!] \stackrel{\hat{\tau}}{\Longrightarrow}\xrightarrow{a}$. No matter which distribution Δ we obtain from executing a finite sequence of internal moves $[\![Q_1]\!] \stackrel{\hat{\tau}}{\Longrightarrow} \Delta$, still part of it is unable to subsequently perform the action a.

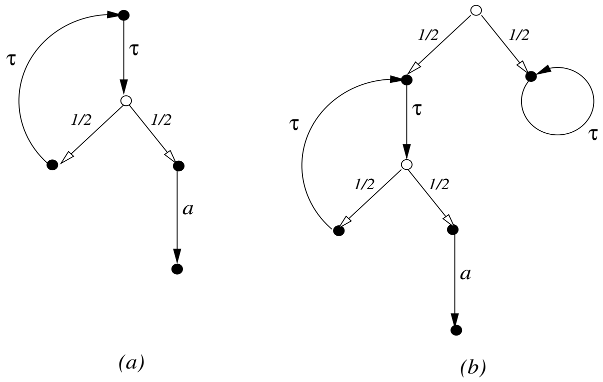

Fig. 1. The pLTSs of processes Q_1 and Q_2

To address this problem we propose a new relation $\Delta \Longrightarrow \Theta$, that indicates that Θ can be derived from Δ by performing an unbounded sequence of internal moves; we call Θ a *weak derivative* of Δ. For example $[\![a.\,\mathbf{0}]\!]$ will turn out to be a weak derivative of $[\![Q_1]\!]$, $[\![Q_1]\!] \Longrightarrow [\![a.\,\mathbf{0}]\!]$, via the infinite sequence of internal moves

$$[\![Q_1]\!] \xrightarrow{\tau} [\![Q_1\ {}_{\frac{1}{2}}\!\oplus\ a.\,\mathbf{0}]\!] \xrightarrow{\tau} [\![Q_1\ {}_{\frac{1}{2^2}}\!\oplus\ a.\,\mathbf{0}]\!] \xrightarrow{\tau} \ldots [\![Q_1\ {}_{\frac{1}{2^n}}\!\oplus\ a.\,\mathbf{0}]\!] \xrightarrow{\tau} \ldots .$$

One of our contributions here is the significant use of "sub distributions" that sum to *no more than* one [8, 11]. For example, the empty subdistribution ε elegantly represents the chaotic behaviour of processes that in CSP and in must-testing semantics is tantamount to divergence, because we have $\varepsilon \xrightarrow{\alpha} \varepsilon$ for any action α, and a process like $\mathsf{rec}\,x.\,x$ that diverges via an infinite τ path gives rise to the weak transition $\mathsf{rec}\,x.\,x \Longrightarrow \varepsilon$. So the process $Q_2 = Q_1\ {}_{\frac{1}{2}}\!\oplus\ \mathsf{rec}\,x.\,x$ illustrated in Figure 1(b) will enable the weak transition $[\![Q_2]\!] \Longrightarrow \frac{1}{2}[\![a.\,\mathbf{0}]\!]$, where intuitively the latter is a proper subdistribution mapping the state $a.\,\mathbf{0}$ to the probability $\frac{1}{2}$. Our weak transition relation $\Longrightarrow$ can be regarded as an extension of the *weak hyper-transition* from [10] to partial distributions; the latter, although defined in a very different way, can be represented in terms of ours by requiring weak derivatives to be total distributions.

We end this introduction with a brief glimpse at our proof strategy. In [1] the characterisations for finite pCSP processes were obtained using a probabilistic extension of the Hennessy-Milner logic [12]. Moving to recursive processes, we know that process behaviour can be captured by a finite modal logic only if the underlying LTS is finitely branching, or at least image-finite [12]. Thus to take advantage of a finite probabilistic HML we need a property of pLTSs corresponding to finite branching in LTSs: this is topological compactness, whose relevance we now sketch.

Subdistributions over (derivatives of) finitary pCSP processes inherit the standard (complete) Euclidean metric. One of our key results is that

Theorem 1. For every finitary pCSP process P, the set $\{ \Delta \mid [\![P]\!] \Longrightarrow \Delta \}$ is convex and compact.

Indeed, using techniques from Markov Decision Theory [13] we can show that the potentially uncountable set $\{ \Delta \mid [\![P]\!] \Longrightarrow \Delta \}$ is nevertheless the convex closure of a *finite* set of subdistributions, from which Theorem 1 follows.

This key result allows an *inductive* characterisation of the simulation preorders $\sqsubseteq_S$ and $\sqsubseteq_{FS}$, here defined using our novel weak derivation relation $\Longrightarrow$. We first construct a sequence of approximations $\sqsubseteq_S^k$ for $k \geq 0$ and, using Theorem 1, we prove

Theorem 2. For every finitary pCSP process P, and for every $k \geq 0$, the set $\{ \Delta \mid [\![P]\!] \sqsubseteq_S^k \Delta \}$ is convex and compact.

This in turn enables us to use the *Finite Intersection Property* of compact sets to prove

Theorem 3. For finitary pCSP processes we have $P \sqsubseteq_S Q$ iff $P \sqsubseteq_S^k Q$ for all $k \geq 0$.

Our main characterisation results can then be obtained by extending the probabilistic modal logic used in [1], so that for example

- it characterises $\sqsubseteq_S^k$ for every $k \geq 0$, and therefore it also characterises $\sqsubseteq_S$
- every probabilistic modal formula can be captured by a may-test.

Similar results accrue for must testing and the new failure simulation preorder $\sqsubseteq_{FS}$: details are given in Section 6.

Due to lack of space, we omit proofs, and most examples: they are reported in [2].

2 The Language pCSP

Let Act be a set of visible actions which a process can perform, and let Var be an infinite set of variables. The language pCSP of probabilistic CSP processes is given by the following two-sorted syntax, in which $p \in [0,1]$, $a \in \mathsf{Act}$ and $A \subseteq \mathsf{Act}$:

$$\begin{aligned} P &::= S \mid P\,{}_p\!\oplus P \\ S &::= \mathbf{0} \mid x \in \mathsf{Var} \mid a.P \mid P \sqcap P \mid S \,\Box\, S \mid S \,|_A\, S \mid \mathsf{rec}\, x.\, P\ . \end{aligned}$$

$$a.P \xrightarrow{a} [\![P]\!] \qquad \mathsf{rec}\,x.\,P \xrightarrow{\tau} [\![P[x \mapsto \mathsf{rec}\,x.\,P]]\!]$$

$$P \sqcap Q \xrightarrow{\tau} [\![P]\!] \qquad P \sqcap Q \xrightarrow{\tau} [\![Q]\!]$$

$$\frac{s_1 \xrightarrow{a} \Delta}{s_1 \,\Box\, s_2 \xrightarrow{a} \Delta} \qquad \frac{s_2 \xrightarrow{a} \Delta}{s_1 \,\Box\, s_2 \xrightarrow{a} \Delta}$$

$$\frac{s_1 \xrightarrow{\tau} \Delta}{s_1 \,\Box\, s_2 \xrightarrow{\tau} \Delta \,\Box\, s_2} \qquad \frac{s_2 \xrightarrow{\tau} \Delta}{s_1 \,\Box\, s_2 \xrightarrow{\tau} s_1 \,\Box\, \Delta}$$

$$\frac{s_1 \xrightarrow{\alpha} \Delta \quad \alpha \notin A}{s_1 \,|_A\, s_2 \xrightarrow{\alpha} \Delta \,|_A\, s_2} \qquad \frac{s_2 \xrightarrow{\alpha} \Delta \quad \alpha \notin A}{s_1 \,|_A\, s_2 \xrightarrow{\alpha} s_1 \,|_A\, \Delta}$$

$$\frac{s_1 \xrightarrow{a} \Delta_1,\; s_2 \xrightarrow{a} \Delta_2 \; a \in A}{s_1 \,|_A\, s_2 \xrightarrow{\tau} \Delta_1 \,|_A\, \Delta_2}$$

Fig. 2. Operational semantics of pCSP

This is essentially the finite language of $[1, 3]$ plus the recursive construct $\mathsf{rec}\,x.\,P$ in which x is a variable and P a term. The notions of free- and bound variables are standard; by $Q[x \mapsto P]$ we indicate substitution of term P for variable x in Q, with renaming if necessary. We write pCSP for the set of closed P-terms defined by this grammar, and sCSP for its *state-based* subset of closed S-terms.

Following $[3, 1]$, we interpret the language as a *probabilistic labelled transition system*. A (discrete) probability *subdistribution* over a set S is a function $\Delta : S \to [0, 1]$ with $\sum_{s \in S} \Delta(s) \leq 1$; the *support* of such a Δ is $\lceil\Delta\rceil := \{\, s \in S \mid \Delta(s) > 0 \,\}$, and its *mass* $|\Delta|$ is $\sum_{s \in \lceil\Delta\rceil} \Delta(s)$. A subdistribution is a (total, or full) *distribution* if $|\Delta| = 1$. The point distribution $\overline{s}$ assigns probability 1 to s and 0 to all other elements of S, so that $\lceil\overline{s}\rceil = \{s\}$. With $\mathcal{D}_{sub}(S)$ we denote the set of subdistributions over S, and with $\mathcal{D}(S)$ its subset of full distributions.

Let $\{\Delta_k \mid k \in K\}$ be a set of subdistributions, possibly infinite. Then $\sum_{k\in K} \Delta_k$ is the real-valued function in $S \to \mathbb{R}$ defined by $(\sum_{k\in K} \Delta_k)(s) := \sum_{k\in K} \Delta_k(s)$. This is a partial operation on subdistributions because for some state s the sum of $\Delta_k(s)$ might exceed 1. If the index set is finite, say $\{1..n\}$, we often write $\Delta_1 + \ldots + \Delta_n$. For p a real number from $[0, 1]$ we use $p\cdot\Delta$ to denote the subdistribution given by $(p\cdot\Delta)(s) := p\cdot\Delta(s)$. Finally we use ε to denote the everywhere-zero subdistribution that thus has empty support. These operations on subdistributions do not readily adapt themselves to distributions; yet if that $\sum_{k\in K} p_k = 1$ for some collection of $p_k \geq 0$, and the Δ_k are distributions, then so is $\sum_{k\in K} p_k\cdot\Delta_k$. In general when $0{\leq}p{\leq}1$ we write $x\,{}_p{\oplus}\,y$ for $p\cdot x + (1{-}p)\cdot y$ where that makes sense, so that for example $\Delta_1\,{}_p{\oplus}\,\Delta_2$ is always defined, and is full if Δ_1 and Δ_2 are.

The expected value $\sum_{s\in S} \Delta(s)\cdot f(s)$ over a distribution Δ of a bounded non-negative function f to the reals or tuples of them is written $\mathrm{Exp}_\Delta(f)$, and the image of a distribution Δ through a function f is written $\mathrm{Img}_f(\Delta)$ — the latter is the distribution over the range of f given by $\mathrm{Img}_f(\Delta)(t) := \sum_{f(s)=t} \Delta(s)$.

Definition 1. A *probabilistic labelled transition system* (pLTS) is a triple $\langle S, L, \rightarrow \rangle$, where

(i) S is a set of states,
(ii) L is a set of transition labels,
(iii) relation $\rightarrow$ is a subset of $S \times L \times \mathcal{D}(S)$.

A (non-probabilistic) labelled transition system (LTS) may be viewed as a degenerate pLTS — one in which only point distributions are used. As with LTSs, we write $s \xrightarrow{\alpha} \Delta$ for $(s, \alpha, \Delta) \in \rightarrow$, as well as $s \xrightarrow{\alpha}$ for $\exists \Delta : s \xrightarrow{\alpha} \Delta$ and $s \rightarrow$ for $\exists \alpha : s \xrightarrow{\alpha}$. A pLTS is *finitely branching* if the set $\{\Delta \mid s \xrightarrow{\alpha} \Delta,\ \alpha \in L\}$ is finite for all states s; if moreover S is finite, then the pLTS is *finitary*. A subdistribution Δ in an arbitrary pLTS is *finitary* if restricting the state set to the states reachable from Δ yields a finitary sub-pLTS.

The operational semantics of pCSP is defined by a particular pLTS $\langle \mathsf{sCSP}, \mathsf{Act}_\tau, \rightarrow \rangle$ in which sCSP is the set of states and $\mathsf{Act}_\tau := \mathsf{Act} \cup \{\tau\}$ is the set of transition labels; we let a range over Act and α over Act_τ. We interpret pCSP processes P as distributions $[\![P]\!] \in \mathcal{D}(\mathsf{sCSP})$ via the function $[\![\_]\!] : \mathsf{pCSP} \rightarrow \mathcal{D}(\mathsf{sCSP})$ defined by

$$[\![s]\!] \ := \ \overline{s} \quad \text{for } s \in \mathsf{sCSP}, \quad \text{and} \quad [\![P \,{}_p\!\oplus Q]\!] \ := \ [\![P]\!] \,{}_p\!\oplus [\![Q]\!] \,.$$

The relations $\xrightarrow{\alpha}$ are defined in Figure 2 which extends the rules used in [3, 1] for finite processes with a new rule for recursion. External choice and parallel composition use an abbreviation for distributing an operator over a distribution: for example $\Delta \,\Box\, s$ is the distribution given by $(\Delta \,\Box\, s)(t) := \Delta(s')$ if t is $s' \,\Box\, s$ and 0 otherwise. We sometimes write $\tau.P$ for $P \sqcap P$, thus giving $\tau.P \xrightarrow{\tau} [\![P]\!]$.

Note that this pLTS is finitely branching and for each $P \in \mathsf{pCSP}$ the distribution $[\![P]\!]$ has finite support. However, it is possible for there to be infinitely many states reachable from $[\![P]\!]$. If only finitely many states are reachable from $[\![P]\!]$, then P is called *finitary*.

3 Testing Probabilistic Processes

We follow the approach of [3, 1] to the testing of probabilistic processes. A *test* is simply a process from the language pCSP except that it may use extra visible actions $\omega_i \notin \mathsf{Act}_\tau$, which are assumed to be fresh, for reporting success. Given a set of test actions Ω, we write pCSP^Ω for the set of pCSP expressions using actions from $\Omega \cup \mathsf{Act}_\tau$, and sCSP^Ω for the set of state-based pCSP^Ω expressions. To apply test T to process P we form the process $T \mid_{\mathsf{Act}} P$ in which *all* visible actions except the ω_i must synchronise, leaving only actions τ and ω_i, and as in [3, 1] we extract testing outcomes from them. However, as processes $T \mid_{\mathsf{Act}} P$ are in general not of finite depth, we can no longer do this inductively. Below we outline two alternative methods that for finitary systems will turn out to be equivalent. The first one is slightly easier to explain, whereas the second one extends the work of [15, 14, 4] and is needed in establishing our results.

3.1 Extremal Testing

For the first method we assume that tests may use only a *single* success action ω. We view the unit interval $[0, 1]$ ordered in the standard manner as a complete lattice; this induces a complete lattice on the set of functions $\mathsf{sCSP}^\Omega \to [0,1]$. Now consider the function $\mathcal{R}_{\min} : (\mathsf{sCSP}^\Omega \to [0,1]) \to (\mathsf{sCSP}^\Omega \to [0,1])$ defined by

$$\mathcal{R}_{\min}(f)(s) := \begin{cases} 1 & \text{if } s \xrightarrow{\omega} \\ 0 & \text{if } s \not\rightarrow \\ \min\{ \operatorname{Exp}_\Delta(f) \mid s \xrightarrow{\alpha} \Delta \} & \text{otherwise.} \end{cases}$$

In a similar fashion we define the function $\mathcal{R}_{\max}$ which uses $\max$ in place of $\min$. Both these functions are monotonic, and therefore have least fixed points which we call $\mathbb{V}_{\min}$, $\mathbb{V}_{\max}$ respectively.

Now for a test T and a process P we have two ways of defining the outcome of the application of T to P:

$$\mathcal{A}^{\mathrm{e}}_{\min}(T,P) := \operatorname{Exp}_{[\![T|_{\mathsf{Act}}P]\!]}(\mathbb{V}_{\min})$$
$$\mathcal{A}^{\mathrm{e}}_{\max}(T,P) := \operatorname{Exp}_{[\![T|_{\mathsf{Act}}P]\!]}(\mathbb{V}_{\max})\ .$$

Here $\mathcal{A}^{\mathrm{e}}_{\min}(T,P)$ returns a single probability p, estimating the minimum probability of success; it is a pessimistic estimate. On the other hand $\mathcal{A}^{\mathrm{e}}_{\max}(T,P)$ is optimistic, in that it gives the maximum probability of success.

Definition 2. The *may*- and *must* preorders are given by

- $P \sqsubseteq^e_{\mathrm{pmay}} Q$ if for every test T we have $\mathcal{A}^{\mathrm{e}}_{\max}(T,P) \leq \mathcal{A}^{\mathrm{e}}_{\max}(T,Q)$
- $P \sqsubseteq^e_{\mathrm{pmust}} Q$ if for every test T we have $\mathcal{A}^{\mathrm{e}}_{\min}(T,P) \leq \mathcal{A}^{\mathrm{e}}_{\min}(T,Q)$.

3.2 Resolution-based Testing

In the second method we use Ω-*tests* for any given collection Ω of success actions disjoint from Act_τ; here ω will be a variable ranging over the individual success actions of Ω. We calculate the result of applying test T to process P in terms of the *resolutions* of the combined process $T\,|_{\mathsf{Act}}\,P$, where intuitively a resolution represents a *run* of a process and, as such, gives exactly one probability for each success action. So in general the application of T to P will yield a *set of vectors* of probabilities.

We define the resolutions of a process $T\,|_{\mathsf{Act}}\,P$ in terms of the distribution $[\![T\,|_{\mathsf{Act}}\,P]\!]$ in the pLTS $\langle \mathsf{sCSP}^\Omega\,|_{\mathsf{Act}}\,\mathsf{sCSP}, \Omega_\tau, \rightarrow\rangle$ obtained by restricting attention to states of the form $t\ |_{\mathsf{Act}}\ s$ with $t \in \mathsf{sCSP}^\Omega$ and $s \in \mathsf{sCSP}$. Note that all transitions in this pLTS have labels τ or $\omega \in \Omega$. Following [5, 15, 3, 1], and unlike [14, 4], this paper employs *state-based* testing [4, 1], meaning that transitions $s \xrightarrow{\omega} \Delta$ are merely expedients to mark the state s as an ω-success state — the target distribution Δ is wholly ignored. Hence the pLTS can also be regarded as having just τ-labels and moreover state markers $\omega \in \Omega$. Intuitively, a resolution of a distribution in such a pLTS is obtained by pruning away multiple τ-transitions from a state until only a single choice remains, possibly introducing some linear combinations in the process.

Definition 3. A pLTS $\langle R, L, \rightarrow\rangle$ is *deterministic* if for every $r \in R$ and every $\alpha \in L$ there is at most one $\Theta \in \mathcal{D}_{sub}(R)$ such that $r \xrightarrow{\alpha} \Theta$.

A *resolution* of a subdistribution $\Delta \in \mathcal{D}_{sub}(S)$ in a pLTS $\langle S, \Omega_\tau, \rightarrow\rangle$ is a triple $\langle R, \Theta, \rightarrow'\rangle$ where $\langle R, \Omega_\tau, \rightarrow'\rangle$ is a deterministic pLTS and $\Theta \in \mathcal{D}_{sub}(R)$, such that there exists a *resolving function* $f \in R \rightarrow S$ satisfying

1. $\mathrm{Img}_f(\Theta) = \Delta$
2. if $r \xrightarrow{\alpha}' \Theta'$ for $\alpha \in \Omega_\tau$ then $f(r) \xrightarrow{\alpha} \mathrm{Img}_f(\Theta')$
3. if $f(r) \xrightarrow{\alpha}$ for $\alpha \in \Omega_\tau$ then $r \xrightarrow{\alpha}'$.

By analogy with the functions $\mathcal{R}_{\text{min}}$ and $\mathcal{R}_{\text{max}}$ of Section 3.1, we define the function $\mathcal{R} : (R{\to}[0,1]^\Omega) \to (R{\to}[0,1]^\Omega)$ for a deterministic pLTS $\langle R, \Omega_\tau, \to\rangle$ as

$$\mathcal{R}(f)(r)(\omega) := \begin{cases} 1 & \text{if } r \xrightarrow{\omega} \\ 0 & \text{if } r \not\xrightarrow{\omega} \text{ and } r \not\xrightarrow{\tau} \\ \mathrm{Exp}_\Delta(f)(\omega) & \text{if } r \not\xrightarrow{\omega} \text{ and } r \xrightarrow{\tau} \Delta. \end{cases}$$

Once more this function has a least fixed point, which we denote by $\mathbb{V}_{\langle R,\Omega_\tau,\to\rangle}$.

Now let $\mathcal{A}^\Omega(T, P)$ denote the set of vectors

$$\{\, \mathrm{Exp}_\Theta(\mathbb{V}_{\langle R,\Omega_\tau,\to\rangle}) \mid \langle R, \Theta, \to\rangle \text{ is a resolution of } [\![T \mid_{\mathsf{Act}} P]\!] \,\} \ .$$

We compare two vectors of probabilities component-wise, and two sets of vectors of probabilities via the Hoare- and Smyth preorders:

$$\begin{array}{lll} X \leq_{\mathrm{Ho}} Y & \text{iff} & \forall x \in X : \exists y \in Y : x \leq y \\ X \leq_{\mathrm{Sm}} Y & \text{iff} & \forall y \in Y : \exists x \in X : x \leq y \ . \end{array}$$

Definition 4. Given two pCSP processes P and Q,
- $P \sqsubseteq^\Omega_{\text{pmay}} Q$ if for every Ω-test T, we have $\mathcal{A}^\Omega(T,P) \leq_{\mathrm{Ho}} \mathcal{A}^\Omega(T,Q)$
- $P \sqsubseteq^\Omega_{\text{pmust}} Q$ if for every Ω-test T, we have $\mathcal{A}^\Omega(T,P) \leq_{\mathrm{Sm}} \mathcal{A}^\Omega(T,Q)$.

These preorders are abbreviated to $P \sqsubseteq_{\text{pmay}} Q$ and $P \sqsubseteq_{\text{pmust}} Q$ when $|\Omega| = 1$.

3.3 Equivalence of Testing Methods

In this section we compare the two approaches of testing introduced in the previous two subsections. First of all, we recall the result from [4] which says that when testing finitary processes it suffices to use a single success action rather than multiple ones.[1]

Theorem 4. For finitary processes:

$$P \sqsubseteq^\Omega_{\text{pmay}} Q \ \text{ iff } \ P \sqsubseteq_{\text{pmay}} Q \qquad \text{and} \qquad P \sqsubseteq^\Omega_{\text{pmust}} Q \ \text{ iff } \ P \sqsubseteq_{\text{pmust}} Q.$$

The following theorem states that, for finitary processes, extremal testing yields the same preorders as resolution-based testing with a single success action.

Theorem 5. For finitary processes

$$P \sqsubseteq^e_{\text{pmay}} Q \ \text{ iff } \ P \sqsubseteq_{\text{pmay}} Q \qquad \text{and} \qquad P \sqsubseteq^e_{\text{pmust}} Q \ \text{ iff } \ P \sqsubseteq_{\text{pmust}} Q.$$

Neither result in Theorem 5 is true in the general (non-finitary) case, as counterexamples in [2, App. A] demonstrate. Although Theorem 4 suggests that we could have avoided multiple success actions in the resolution-based definition of testing, our completeness proof (Theorem 15) makes essential use of a countable set of them.

[1] The result in [4] is stated for *action-based* testing, meaning that it is the actual execution of a success action rather than reaching a success state that constitutes success, but, as mentioned in the conclusion of [4], it also holds in our current state-based setting.

4 A Novel Approach to Weak Derivations

In this section we develop a new definition of what it means for a recursive process to evolve by silent activity into another process; it allows the simulation and failure-simulation preorders of [1] to be adapted to characterise the testing preorders for at least finitary probabilistic processes. The key technical generalisation is the *subdistributions* that enable us to express divergence very conveniently. [2]

In a pLTS actions are only performed by states, in that actions are given by relations from states to distributions. But pCSP processes in general correspond to distributions over states, so in order to define what it means for a process to perform an action we need to *lift* these relations so that they also apply to (sub)distributions.

Definition 5. Let $(S, L, \rightarrow)$ be a pLTS and $\mathcal{R} \subseteq S \times \mathcal{D}_{sub}(S)$ be a relation from states to subdistributions. Then $\overline{\mathcal{R}} \subseteq \mathcal{D}_{sub}(S) \times \mathcal{D}_{sub}(S)$ is the smallest relation that satisfies

(1) $s \; \mathcal{R} \; \Theta$ implies $\overline{s} \; \overline{\mathcal{R}} \; \Theta$, and
(2) (Linearity) $\Delta_i \; \overline{\mathcal{R}} \; \Theta_i$ for $i \in I$ implies $(\sum_{i\in I} p_i \cdot \Delta_i) \; \overline{\mathcal{R}} \; (\sum_{i\in I} p_i \cdot \Theta_i)$ for any $p_i \in [0,1]$ with $\sum_{i\in I} p_i \leq 1$.

This applies when the relation is $\xrightarrow{\alpha}$ for $\alpha \in \mathsf{Act}_\tau$, where we also write $\xrightarrow{\alpha}$ for $\overline{\xrightarrow{\alpha}}$. Thus as source of a relation $\xrightarrow{\alpha}$ we now also allow distributions, and even subdistributions. A subtlety of this approach is that for any action α, we have $\varepsilon \xrightarrow{\alpha} \varepsilon$ simply by taking $I = \emptyset$ or $\sum_{i\in I} p_i = 0$ in Definition 5. That will turn out to make ε especially useful for modelling the "chaotic" aspects of divergence, in particular that in the must-case a divergent process can mimic any other.

We now formally define the notation of weak derivatives.

Definition 6. Suppose we have subdistributions $\Delta, \Delta_k^{\rightarrow}, \Delta_k^{\times}$, for $k \geq 0$, with the following properties:

$$\begin{aligned} \Delta \;&=\; \Delta_0^{\rightarrow} + \Delta_0^{\times} \\ \Delta_0^{\rightarrow} \;&\xrightarrow{\tau}\; \Delta_1^{\rightarrow} + \Delta_1^{\times} \\ &\vdots \\ \Delta_k^{\rightarrow} \;&\xrightarrow{\tau}\; \Delta_{k+1}^{\rightarrow} + \Delta_{k+1}^{\times} \; . \\ &\vdots \end{aligned}$$

Then we call $\Delta' := \sum_{k=0}^{\infty} \Delta_k^{\times}$ a *weak derivative* of Δ, and write $\Delta \Longrightarrow \Delta'$ to mean that Δ can make a *weak τ move* to its derivative Δ'.

It is easy to check that $\sum_{k=0}^{\infty} \Delta_k^{\times}$ is indeed a subdistribution, whereas in general it is not a full distribution: for instance we have $[\![\mathsf{rec}\, x.\, x]\!] \Longrightarrow \varepsilon$. By setting appropriate $\Delta_k^{\times}$'s to ε we see that $\Delta(\xrightarrow{\tau})^*\Phi$, where * denotes reflexive and transitive closure, implies $\Delta \Longrightarrow \Phi$. It is also easy to check that on recursion-free pCSP the relation $\Longrightarrow$ agrees with the one defined in [3, 1] by means of transitive closure. Moreover the standard notion of *divergence*, the ability of a subdistribution Δ to perform an infinite sequence of τ transitions, is neatly captured by the relation $\Delta \Longrightarrow \varepsilon$.

[2] Subdistributions' nice properties with respect to divergence are due to their being equivalent to the discrete probabilistic powerdomain over a flat domain [8].

Example 1. Consider the (infinite) collection of states s_k and probabilities p_k for $k \geq 2$ such that

$$s_k \xrightarrow{\tau} [\![a.\mathbf{0}]\!]_{p_k}\oplus \overline{s_{k+1}} \ ,$$

where we choose p_k so that starting from any s_k the probability of eventually taking a left-hand branch, and so reaching $[\![a.\mathbf{0}]\!]$ ultimately, is just $1/k$ in total. Thus p_k must satisfy $1/k = p_k + (1-p_k)/(k+1)$, whence by arithmetic we have that $p_k := 1/k^2$ will do. Therefore in particular $\overline{s_2} \Longrightarrow \frac{1}{2}[\![a.\mathbf{0}]\!]$, with the remaining $\frac{1}{2}$ lost in divergence.

Definition 7. Let Δ and its variants be subdistributions in a pLTS $\langle S, \mathsf{Act}_\tau, \rightarrow\rangle$.

- For $a \in \mathsf{Act}$ write $\Delta \stackrel{a}{\Longrightarrow} \Delta'$ whenever $\Delta \Longrightarrow \Delta^{\text{pre}} \xrightarrow{a} \Delta^{\text{post}} \Longrightarrow \Delta'$. Extend this to Act_τ by allowing as a special case that $\stackrel{\tau}{\Longrightarrow}$ is simply $\Longrightarrow$, i.e. including identity (rather than requiring at least one $\xrightarrow{\tau}$).
- For $A \subseteq \mathsf{Act}$ and $s \in S$ write $s \stackrel{A}{\not\longrightarrow}$ if $s \stackrel{\alpha}{\not\longrightarrow}$ for every $\alpha \in A \cup \{\tau\}$; write $\Delta \stackrel{A}{\not\longrightarrow}$ if $s \stackrel{A}{\not\longrightarrow}$ for every $s \in \lceil\Delta\rceil$.
- More generally write $\Delta \Longrightarrow\stackrel{A}{\not\longrightarrow}$ if $\Delta \Longrightarrow \Delta^{\text{pre}}$ for some Δ^{pre} such that $\Delta^{\text{pre}} \stackrel{A}{\not\longrightarrow}$.

For example, in Figure 1 we have $[\![Q_1]\!] \stackrel{a}{\Longrightarrow} [\![\mathbf{0}]\!]$, because $[\![Q_1]\!] \Longrightarrow [\![a.\mathbf{0}]\!] \xrightarrow{a} [\![\mathbf{0}]\!]$.

5 Some properties of weak derivations in finitary pLTSs

In this section we expose some less obvious properties of weak derivations from states in finitary pLTSs, relating to their behaviour at infinity; they underpin many results in the next section. One important property is that the set of weak derivations from a single starting point is *compact* in the sense (from analysis) of being bounded and containing all its limit points, where, in turn, limits depend on a Euclidean-style metric defining the distance between two distributions in a straightforward way. The other property is "distillation of divergence", allowing us to find in any weak derivation that partially diverges (by no matter how small an amount) a point at which the divergence is "distilled" into a state which wholly diverges.

Both properties depend on our working within *finitary* pLTSs — that is, ones in which the state space is finite and the (unlifted) transition relation is finite-branching.

5.1 Finite generability and closure

In a finitary pLTS, by definition the sets $\{\Delta \mid s \xrightarrow{\alpha} \Delta\}$ are finite, for every s and α. This of course is no longer true for the lifted relations $\xrightarrow{\alpha}$ over subdistributions; nevertheless, the sets $\{\Delta \mid \overline{s} \xrightarrow{\alpha} \Delta\}$ and their weak counterparts $\{\Delta \mid \overline{s} \stackrel{\alpha}{\Longrightarrow} \Delta\}$ can be finitely represented. Below, we focus on the set $\{\Delta \mid \overline{s} \Longrightarrow \Delta\}$.

Definition 8. A *static derivative policy* (SDP) for a pLTS $\langle S, \mathsf{Act}_\tau, \rightarrow\rangle$ is a partial function $\mathsf{pp} : S \rightharpoonup \mathcal{D}(S)$ such that if pp is defined at s then $s \xrightarrow{\tau} \mathsf{pp}(s)$.

Intuitively a policy pp decides for each state, once and for all, which of the available τ-choices to take, if any: since it either chooses a specific transition, or inaction (by being undefined), it does not interpolate via a convex combination of two different transitions; and since it is a function of the state, it makes the same choice on every visit.

The great importance for us of SDP's is that they give a particularly simple characterisation of weak derivatives, provided the state-space is finite and the pLTS is finitely branching. This is essentially a result of Markov Decision Processes [13], which we translate into our context. We first introduce a notion of SDP-derivatives by adapting Definition 6.

Definition 9 (SDP-derivatives). Let pp be a SDP. We write $\Delta \Longrightarrow_{\mathsf{pp}} \Delta'$ if $\Delta \Longrightarrow \Delta'$ and the following holds (using the notation of Def. 6 and writing Δ_k for $\Delta_k^{\rightarrow} + \Delta_k^{\times}$):

$$\begin{aligned} \Delta_k^{\times}(s) &= \begin{cases} 0 & \text{if pp defined at } s \\ \Delta_k(s) & \text{otherwise} \end{cases} \\ \Delta_{k+1} &= \sum\{\Delta_k(s)\cdot \mathsf{pp}(s) \mid s \in \lceil \Delta_k \rceil \text{ and pp defined at } s\}. \end{aligned}$$

Intuitively, $\Delta \Longrightarrow_{\mathsf{pp}} \Delta'$ means that Δ' is the single derivative of Δ that results from using policy pp to construct the weak transition $\Delta \Longrightarrow \Delta'$. Note that, for a given SDP pp, the relation $\Longrightarrow_{\mathsf{pp}}$ is actually a function; moreover in a finitary pLTS the set of all possible SDPs is finite, due to the constraints of Definition 8.

Theorem 6 (Finite generability). Let s be a state in a finitary pLTS $\langle S, \mathsf{Act}_\tau, \rightarrow\rangle$. Then $s \Longrightarrow \Delta$ for some $\Delta \in \mathcal{D}_{sub}(S)$ iff there is a finite index set I, probabilities p_i summing to 1 and static derivative policies pp_i with $s \Longrightarrow_{\mathsf{pp}_i} \Delta_i$ for each i, such that $\Delta = \sum_{i\in I} p_i \cdot \Delta_i$.

Since the convex closure of a finite set of points is always compact, we obtain

Corollary 1. For any state s in a finitary pLTS the set $\{\Delta \mid s \Longrightarrow \Delta\}$ is convex and compact.

A similar result is obtained by Desharnais, Gupta, Jagadeesan & Panagaden [6].

Although the pLTS $\langle \mathsf{sCSP}, \mathsf{Act}_\tau, \rightarrow\rangle$ is not finitary, the interpretation $[\![P]\!] \in \mathcal{D}(\mathsf{sCSP})$ of a finitary pCSP process P can also be understood to be a distribution in a finitary pLTS, namely the restriction of $\langle \mathsf{sCSP}, \mathsf{Act}_\tau, \rightarrow\rangle$ to the states reachable from $[\![P]\!]$. Using this, Corollary 1 leads to the essential Theorem 1, referred to in the introduction.

5.2 Distillation of divergence

Although it is possible to have processes that diverge with some probability strictly between zero and one, in a finitary pLTS we can *distill* divergence in the sense that for many purposes we can limit our analyses to processes that either wholly diverge (can do so with probability one) or wholly converge (can diverge only with probability zero). This property is based on the zero-one law for finite-state probabilistic systems, relevant aspects of which we present in this sub-section.

We first note that static derivative policies obey the following zero-one law.

Theorem 7 (Zero-one law). If for a static derivative policy pp over a finite-state pLTS there is for some s a derivation $s \Longrightarrow_{\mathsf{pp}} \Delta$ with $|\Delta| < 1$ then in fact for some (possibly different) state s_ε we have $s_\varepsilon \Longrightarrow_{\mathsf{pp}} \varepsilon$.

Based on Theorems 6 and 7, the following property of weak derivations can now be established.

Theorem 8 (Distillation of divergence). For any s, Δ in a finitary pLTS with $s \Longrightarrow \Delta$ there is a probability p and full distributions $\Delta_1, \Delta_\varepsilon$ such that $s \Longrightarrow (\Delta_1 \ {}_p\!\oplus\ \Delta_\varepsilon)$ and $\Delta = p\cdot\Delta_1$ and $\Delta_\varepsilon \Longrightarrow \varepsilon$.

6 Failure Simulation is Sound and Complete for Must Testing

In this section we define the failure-simulation preorder and show that it is sound and complete for the must-testing preorder. The following presentation is an enhancement of our earlier definition in [1].

Definition 10 (Failure-Simulation Preorder). Define $\sqsupseteq_{FS}$ to be the largest relation in $\mathcal{D}_{sub}(S) \times \mathcal{D}_{sub}(S)$ such that if $\Delta \sqsupseteq_{FS} \Theta$ then

1. whenever $\Delta \stackrel{\alpha}{\Longrightarrow} (\sum_i p_i \Delta'_i)$, for $\alpha \in \mathsf{Act}_\tau$ and certain p_i with $(\sum_i p_i) \leq 1$, then there are $\Theta'_i \in \mathcal{D}_{sub}(S)$ with $\Theta \stackrel{\alpha}{\Longrightarrow} (\sum_i p_i \Theta'_i)$ and $\Delta'_i \sqsupseteq_{FS} \Theta'_i$ for each i
2. and whenever $\Delta \Longrightarrow\stackrel{A}{\not\longrightarrow}$ then also $\Theta \Longrightarrow\stackrel{A}{\not\longrightarrow}$.

Naturally $\Theta \sqsubseteq_{FS} \Delta$ just means $\Delta \sqsupseteq_{FS} \Theta$. For pCSP processes P and Q and any preorder $\sqsubseteq\ \subseteq \mathcal{D}_{sub}(\mathsf{sCSP}) \times \mathcal{D}_{sub}(\mathsf{sCSP})$ we write $P \sqsubseteq Q$ for $[\![P]\!] \sqsubseteq [\![Q]\!]$.

Although the regularity of Definition 10 is appealing — for example it is trivial to see that $\sqsubseteq_{FS}$ is reflexive and transitive, as it should be — in practice, for specific processes, it is easier to work with a characterisation of the failure-simulation preorder in terms of a relation between *states* and subdistributions.

Definition 11 (Failure Similarity). Let $\lhd_{FS}$ be the largest relation in $S \times \mathcal{D}_{sub}(S)$ such that if $s \lhd_{FS} \Theta$ then

1. whenever $\overline{s} \Longrightarrow \varepsilon$ then also $\Theta \Longrightarrow \varepsilon$,
2. whenever $s \stackrel{\alpha}{\longrightarrow} \Delta'$, for $\alpha \in \mathsf{Act}_\tau$, then there is a Θ' with $\Theta \stackrel{\alpha}{\Longrightarrow} \Theta'$ and $\Delta' \ \overline{\lhd_{FS}}\ \Theta'$
3. and whenever $s \stackrel{A}{\not\longrightarrow}$ then $\Theta \Longrightarrow\stackrel{A}{\not\longrightarrow}$.

As an example, in Figure 1 it is straightforward to exhibit failure simulations to prove both $[\![Q_1]\!] \ \overline{\lhd_{FS}}\ [\![a.\mathbf{0}]\!]$ and the converse $[\![a.\mathbf{0}]\!] \ \overline{\lhd_{FS}}\ [\![Q_1]\!]$, the essential ingredient being the weak move $[\![Q_1]\!] \stackrel{a}{\Longrightarrow} [\![\mathbf{0}]\!]$. Likewise, we have $a.\mathbf{0} \lhd_{FS} [\![Q_1 \ {}_{\frac{1}{2}}\!\oplus\ \mathsf{rec}\, x.\, x]\!]$, the additional ingredient being $\mathbf{0} \lhd_{FS} \varepsilon$.

The next result shows how the failure-simulation preorder can alternatively be defined in terms of failure similarity. This is actually how we defined it in [1].

Theorem 9. For finitary $\Delta, \Theta \in \mathcal{D}_{sub}(S)$ we have $\Delta \sqsupseteq_{FS} \Theta$ just when there is a $\Theta^\natural$ with $\Theta \Longrightarrow \Theta^\natural$ and $\Delta \ \overline{\lhd_{FS}}\ \Theta^\natural$.

The proof of this theorem depends crucially on Theorems 1 and 8. The restriction to finitary subdistributions is essential, as in [2, App. A] we provide a counterexample to the general case. It is in terms of this characterisation that we establish soundness and completeness of the failure-simulation preorder with respect to the must-testing preorder; consequently we have these results for finitary processes only.

Theorem 10 (Precongruence). If P_1, P_2, Q_1 and Q_2 are finitary pCSP processes with $P_1 \sqsupseteq_{FS} Q_1$ and $P_2 \sqsupseteq_{FS} Q_2$ then we have $\alpha.P_1 \sqsupseteq_{FS} \alpha.Q_1$ for any $\alpha \in \mathsf{Act}_\tau$, as well as $P_1 \odot P_2 \sqsupseteq_{FS} Q_1 \odot Q_2$ for $\odot$ any of the operators $\sqcap$, $\Box$, ${}_p\oplus$ and $|_A$.

The proof of this precongruence property involves a significant complication: in order to relate two processes we have to demonstrate that if the first diverges then so does the second. This affects particularly the proof that $\sqsupseteq_{FS}$ is preserved by the parallel operator $|_A$. The approach we use involves first characterising divergence coinductively and then applying a novel coinductive proof technique.

Theorem 11 (Soundness and Completeness). For finitary pCSP processes P and Q we have $P \sqsubseteq_{FS} Q$ iff $P \sqsubseteq_{\text{pmust}} Q$.

Soundness, that $\sqsubseteq_{FS} \subseteq \sqsubseteq_{\text{pmust}}$, is a relatively easy consequence of $\sqsubseteq_{FS}$ being a precongruence (Theorem 10). The completeness proof (that $\sqsubseteq_{\text{pmust}} \subseteq \sqsubseteq_{FS}$) is much more complicated and proceeds in three steps, which we detail below. First we provide a characterisation of the preorder relation $\sqsubseteq_{FS}$ by finite approximations. Secondly, using this, we develop a modal logic which can be used to characterise the failure-simulation preorder on finitary processes. Finally, we adapt the results of [1] to show that the modal formulae can in turn be characterised by tests. From this, completeness follows.

6.1 Inductive Characterisation

The relation $\lhd_{FS}$ of Definition 11 is given coinductively: it is the largest fixpoint of an equation $\mathcal{R} = \mathcal{F}(\mathcal{R})$. An alternative approach is to use that $\mathcal{F}(-)$ to define $\lhd_{FS}$ as a limit of approximants:

Definition 12. For every $k \geq 0$ we define the relations $\lhd_{FS}^k \subseteq S \times \mathcal{D}_{sub}(S)$ as follows:

(i) $\lhd_{FS}^0 \;:=\; S \times \mathcal{D}_{sub}(S)$
(ii) $\lhd_{FS}^{k+1} \;:=\; \mathcal{F}(\lhd_{FS}^k)$

Finally let $\lhd_{FS}^\infty := \bigcap_{k=0}^\infty \lhd_{FS}^k$. Furthermore, for every $k \geq 0$ let $\Delta \sqsupseteq_{FS}^k \Theta$ if there exists a $\Theta \Longrightarrow \Theta^\natural$ with $\Delta \;\overline{\lhd_{FS}^k}\; \Theta^\natural$, and let $\sqsupseteq_{FS}^\infty$ denote $\bigcap_{k=0}^\infty \sqsupseteq_{FS}^k$.

Theorem 12. For finitary pCSP processes P and Q we have $P \sqsupseteq_{FS}^\infty Q$ iff $P \sqsupseteq_{FS} Q$.

To show this theorem, we need to use two key results, Propositions 1 and 2 below. We say a relation $\mathcal{R} \subseteq S \times \mathcal{D}(S)$ is convex (resp. compact) whenever the set $\{\Delta \mid s\;\mathcal{R}\;\Delta\}$ is convex (resp. compact) for every $s \in S$.

Proposition 1. In a finitary pLTS, the relation $\lhd_{FS}^k$ is convex and compact, for every $k \geq 0$.

The proof of this property heavily relies on Corollary 1.

Proposition 2. Suppose $\mathcal{R}^k \subseteq S \times \mathcal{D}_{sub}(S)$ is a sequence of convex and compact relations such that $\mathcal{R}^{k+1} \subseteq \mathcal{R}^k$. Then $\overline{(\bigcap_{k=0}^\infty \mathcal{R}^k)} \subseteq (\bigcap_{k=0}^\infty \overline{\mathcal{R}^k})$.

This proposition is proved using the Finite Intersection Property of compact sets [9].

6.2 A Modal Logic

Let $\mathcal{F}$ be the set of modal formulae defined inductively as follows:

- $\mathbf{div}, \top \in \mathcal{F}$
- $\mathbf{ref}(A) \in \mathcal{F}$ when $A \subseteq \mathsf{Act}$,
- $\langle a \rangle \varphi \in \mathcal{F}$ when $\varphi \in \mathcal{F}$ and $a \in \mathsf{Act}$,
- $\varphi_1 \wedge \varphi_2 \in \mathcal{F}$ when $\varphi_1, \varphi_2 \in \mathcal{F}$,
- $\varphi_1 \,{}_p\!\oplus \varphi_2 \in \mathcal{F}$ when $\varphi_1, \varphi_2 \in \mathcal{F}$ and $p \in [0,1]$.

This generalises the modal language used in [1] by the addition of the new constant $\mathbf{div}$, representing the ability of a process to diverge.

Relative to a given pLTS $\langle S, \mathsf{Act}_\tau, \rightarrow \rangle$ the *satisfaction relation* $\models \subseteq \mathcal{D}_{sub}(S) \times \mathcal{F}$ is given by:

- $\Delta \models \top$ for any $\Delta \in \mathcal{D}_{sub}(S)$,
- $\Delta \models \mathbf{div}$ iff $\Delta \Longrightarrow \varepsilon$,
- $\Delta \models \mathbf{ref}(A)$ iff $\Delta \Longrightarrow \xrightarrow{A}\!\!\!\!\!/\;$,
- $\Delta \models \langle a \rangle \varphi$ iff there is a Δ' with $\Delta \overset{a}{\Longrightarrow} \Delta'$ and $\Delta' \models \varphi$,
- $\Delta \models \varphi_1 \wedge \varphi_2$ iff $\Delta \models \varphi_1$ and $\Delta \models \varphi_2$,
- $\Delta \models \varphi_1 \,{}_p\!\oplus \varphi_2$ iff there are $\Delta_1, \Delta_2 \in \mathcal{D}_{sub}(S)$ with $\Delta_1 \models \varphi_1$ and $\Delta_2 \models \varphi_2$, such that $\Delta \Longrightarrow \Delta_1 \,{}_p\!\oplus \Delta_2$.

We write $\Delta \sqsupseteq^{\mathcal{F}} \Theta$ when $\Delta \models \varphi$ implies $\Theta \models \varphi$ for all $\varphi \in \mathcal{F}$, and can verify that $\sqsupseteq_{FS}$ is sound for $\sqsupseteq^{\mathcal{F}}$. In establishing the converse, we mimic the development in Section 7 of [1] by designing *characteristic formulae* which capture the behaviour of states in a pLTS. But here the behaviour is not characterised relative to $\lhd_{FS}$, but rather to the sequence of approximating relations $\lhd_{FS}^k$.

Definition 13. In a finitary pLTS $\langle S, \mathsf{Act}_\tau, \rightarrow \rangle$, the k^{th} *characteristic formulae* $\varphi_s^k, \varphi_\Delta^k$ of states $s \in S$ and subdistributions $\Delta \in \mathcal{D}_{sub}(S)$ are defined inductively as follows:

- $\varphi_s^0 = \top$ and $\varphi_\Delta^0 = \top$,
- $\varphi_s^{k+1} = \mathbf{div}$, provided $\overline{s} \Longrightarrow \varepsilon$,
- $\varphi_s^{k+1} = \mathbf{ref}(A) \wedge \bigwedge_{s \xrightarrow{a} \Delta} \langle a \rangle \varphi_\Delta^k$ where $A = \{a \in \mathsf{Act} \mid s \xrightarrow{a}\!\!\!\!\!/\;\}$, provided $s \xrightarrow{\tau}\!\!\!\!\!/\;$,
- $\varphi_s^{k+1} = \bigwedge_{s \xrightarrow{a} \Delta} \langle a \rangle \varphi_\Delta^k \wedge \bigwedge_{s \xrightarrow{\tau} \Delta} \varphi_\Delta^k$ otherwise,
- and $\varphi_\Delta^{k+1} = (\bigoplus_{s \in \lceil \Delta \rceil} \frac{\Delta(s)}{\lceil \Delta \rceil} \cdot \varphi_s^{k+1}) \;{}_{\lceil \Delta \rceil}\!\oplus\; (\mathbf{div})$.

The next result relates the k^{th} characteristic formulae to the k^{th} failure similarity.

Proposition 3. For $k \geq 0$ we have

(i) $\Theta \models \varphi_s^k$ implies $s \lhd_{FS}^k \Theta$,
(ii) $\Theta \models \varphi_\Delta^k$ implies $\Theta \sqsupseteq_{FS}^k \Delta$.

Using Proposition 3 we obtain a logical characterisation of $\sqsupseteq_{FS}^\infty$ (and hence of $\sqsupseteq_{FS}$):

Theorem 13. For finitary pCSP processes P and Q we have $P \sqsupseteq^{\mathcal{F}} Q$ iff $P \sqsupseteq_{FS}^\infty Q$.

6.3 Characteristic Tests for Formulae

The import of Theorems 12 and 13 is that we can obtain completeness of the failure-simulation preorder with respect to the must-testing preorder by designing for each formula φ a test which in some sense characterises the property that a process satisfies φ. This was achieved for the pLTS generated by the recursion-free fragment of pCSP in Section 8 of [1]. Here we have generalised this technique to the pLTS generated by the set of finitary pCSP terms. The crucial property is stated as follows.

Theorem 14. For every formula $\varphi \in \mathcal{F}$ there exists a pair (T_φ, v_φ) with T_φ an Ω-test and $v_\varphi \in [0,1]^\Omega$ such that $\Delta \models \varphi$ if and only if $\exists o \in \mathcal{A}^\Omega(T_\varphi, \Delta) : o \leq v_\varphi$. Test T_φ is called a *characteristic test* of φ and v_φ is its *target value*.

This property can be shown by exploiting several characteristics of the testing function $\mathcal{A}^\Omega(-,-)$; unlike in [1] these cannot be obtained inductively. The most complicated one is the following.

Proposition 4. If $o \in \mathcal{A}^\Omega(T_1 \sqcap T_2, \Delta)$ then there are a $q \in [0,1]$ and $\Delta_1, \Delta_2 \in \mathcal{D}_{sub}(\mathsf{sCSP})$ such that $\Delta \Longrightarrow q \cdot \Delta_1 + (1-q) \cdot \Delta_2$ and $o = q \cdot o_1 + (1-q) \cdot o_2$ for certain $o_i \in \mathcal{A}^\Omega(T_i, \Delta_i)$.

From Theorem 14 we obtain that the must-testing preorder is at least as discriminating as the logical preorder:

Theorem 15. Let P and Q be pCSP processes. If $P \sqsupseteq^\Omega_{\text{pmust}} Q$ then $P \sqsupseteq^{\mathcal{F}} Q$.

The completeness result in Theorem 11 follows by combining Theorems 15, 13 and 12.

7 Simulation is Sound and Complete for May Testing

We define a simulation preorder that can be shown sound and complete for may testing following the same strategy as for failure simulation and must testing, except that we restrict our treatment to full distributions, a simpler domain. This is possible because in may testing an infinite τ-path is not treated specially — it engages in no visible actions; in must testing, infinite τ-paths potentially can do anything (chaos).

Definition 14 (Simulation Preorder). Let $\sqsubseteq_S$ be the largest relation in $\mathcal{D}(S) \times \mathcal{D}(S)$ such that if $\Delta \sqsubseteq_S \Theta$ then

whenever $\Delta \overset{\alpha}{\Longrightarrow} (\sum_i p_i \Delta'_i)$, for $\alpha \in \mathsf{Act}_\tau$ and certain p_i with $(\sum_i p_i) \leq 1$,
then there are $\Theta'_i \in \mathcal{D}(S)$ with $\Theta \overset{\alpha}{\Longrightarrow} (\sum_i p_i \Theta'_i)$ and $\Delta'_i \sqsubseteq_S \Theta'_i$ for each i.

The technical development from this point on is similar to that given in Section 6. For the modal logic, we use the set of formulae obtained from $\mathcal{F}$ by skipping the $\mathbf{div}$ and $\mathbf{ref}(A)$ clauses. However the satisfaction relation used for this sub-logic is radically different from that given in Section 6.2, because here the interpretation is relative to full distributions. Nevertheless we still obtain the counterparts of Theorems 12, 13 and 15.

Theorem 16 (Soundness and Completeness). For finitary pCSP processes P and Q we have $P \sqsubseteq_{\text{pmay}} Q$ if and only if $P \sqsubseteq_S Q$.

8 Conclusion and Related Work

In this paper we continued our previous work [3, 4, 1] in our quest for a testing theory for processes which exhibit both nondeterministic and probabilistic behaviour. We have generalised our results in [1] of characterising the may preorder as a simulation relation and the must preorder as a failure-simulation relation, from finite processes to finitary processes. To do this it was necessary to investigate fundamental structural properties of derivation sets (finite generability) and similarities (infinite approximations), which are of independent interest. The use of Markov Decision Processes and Zero-One laws was essential in obtaining our results.

Segala [14] defined two preorders called trace distribution precongruence ($\sqsubseteq_{TD}$) and failure distribution precongruence ($\sqsubseteq_{FD}$). He proved that the former coincides with an action-based version of $\sqsubseteq^{\Omega}_{\text{pmay}}$ and that for "probabilistically convergent" systems the latter coincides with an action-based version of $\sqsubseteq^{\Omega}_{\text{pmust}}$. The condition of probabilistic convergence amounts in our framework to the requirement that for $\Delta \in \mathcal{D}(S)$ and $\Delta \Longrightarrow \Delta'$ we have $|\Delta'| = 1$. In [10] it has been shown that $\sqsubseteq_{TD}$ coincides with a notion of simulation akin to $\sqsubseteq_S$. Other probabilistic extensions of simulation occurring in the literature are reviewed in [3, 1].

References

1. Y. Deng, R.J. van Glabbeek, M. Hennessy & C.C. Morgan (2008): *Characterising testing preorders for finite probabilistic processes.* Logical Methods in Computer Science 4(4:4).
2. Y. Deng, R.J. van Glabbeek, M. Hennessy & C.C. Morgan (2009): *Testing finitary probabilistic processes.* Full version of this extended abstract. Available at `http://www.cse.unsw.edu.au/~rvg/pub/finitary.pdf`
3. Y. Deng, R.J. van Glabbeek, M. Hennessy, C.C. Morgan & C. Zhang (2007): *Remarks on testing probabilistic processes.* ENTCS 172, pp. 359–397.
4. Y. Deng, R.J. van Glabbeek, C.C. Morgan & C. Zhang (2007): *Scalar outcomes suffice for finitary probabilistic testing.* In Proc. ESOP'07, LNCS 4421, Springer, pp. 363–378.
5. R. De Nicola & M. Hennessy (1984): *Testing equivalences for processes.* Theoretical Computer Science 34, pp. 83–133.
6. J. Desharnais, V. Gupta, R. Jagadeesan & P. Panagaden (2002): *Weak bisimulation is sound and complete for PCTL.* In Proc. CONCUR'02, LNCS 2421, Springer, pp. 355-370.
7. C.A.R. Hoare (1985): *Communicating Sequential Processes*. Prentice-Hall.
8. C. Jones (1990) *Probabilistic Non-determinism.* Ph.D. Thesis, University of Edinburgh.
9. S. Lipschutz (1965): *Schaum's outline of theory and problems of general topology*. McGraw-Hill.
10. N. Lynch, R. Segala & F.W. Vaandrager (2007): *Observing branching structure through probabilistic contexts.* SIAM Journal on Computing 37(4), pp. 977–1013.
11. A.K. McIver & C.C. Morgan (2005): *Abstraction, Refinement and Proof for Probabilistic Systems*. Springer.
12. R. Milner (1989): *Communication and Concurrency*. Prentice-Hall.
13. M. Puterman (1994): *Markov Decision Processes*. Wiley.
14. R. Segala (1996): *Testing probabilistic automata.* In Proc. CONCUR'96, LNCS 1119, Springer, pp. 299–314.
15. Wang Yi & K.G. Larsen (1992): *Testing probabilistic and nondeterministic processes.* In Proc. PSTV'92, IFIP Transactions C-8, North-Holland, pp. 47–61.