

# Compositional Reasoning for Markov Decision Processes (Extended Abstract)

Yuxin Deng<sup>1,2</sup> \* and Matthew Hennessy<sup>3</sup> \*\*

<sup>1</sup> Dept. Comp. Sci. & Eng. and MOE-Microsoft Key Lab for Intell. Comp. & Syst.,  
Shanghai Jiao Tong University, China

<sup>2</sup> State Key Lab of Comp. Sci., Inst. of Software, Chinese Academy of Sciences

<sup>3</sup> Trinity College Dublin, Ireland

**Abstract.** Markov decision processes (MDPs) have long been used to model qualitative aspects of systems in the presence of uncertainty. However, much of the literature on MDPs takes a monolithic approach, by modelling a system as a particular MDP; properties of the system are then inferred by analysis of that particular MDP. In this paper we develop compositional methods for reasoning about the qualitative behaviour of MDPs. We consider a class of labelled MDPs called weighted MDPs from a process algebraic point of view. For these we define a coinductive simulation-based behavioural preorder which is compositional in the sense that it is preserved by structural operators for constructing MDPs from components.

For finitary convergent processes, which are finite-state and finitely branching systems without divergence, we provide two characterisations of the behavioural preorder. The first uses a novel qualitative probabilistic logic, while the second is in terms of a novel form of testing, in which benefits are accrued during the execution of tests.

## 1 Introduction

Markov decision processes (MDPs) have long been used to model qualitative aspects of systems in the presence of uncertainty [12, 13, 1]. A comprehensive account of analysis techniques may be found in [12], while [13] provides a good account of *model-checking*.

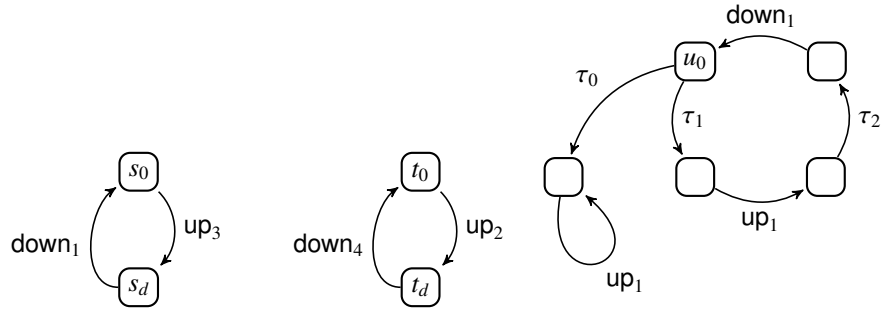
However, much of the literature on MDPs takes a monolithic view of systems; essentially a system is modelled using a particular MDP, and properties of the system are then inferred by analysis of that MDP. In this paper, instead, we would like to develop compositional methods for reasoning about qualitative behaviour of Markov decision processes. This involves defining an appropriate method for comparing the behaviour MDPs which is susceptible to compositional analysis; the behaviour of a composite system should be determined by that of its components.

Our starting point is the idea of one system being able to *simulate* another. For example consider the following three systems:

---

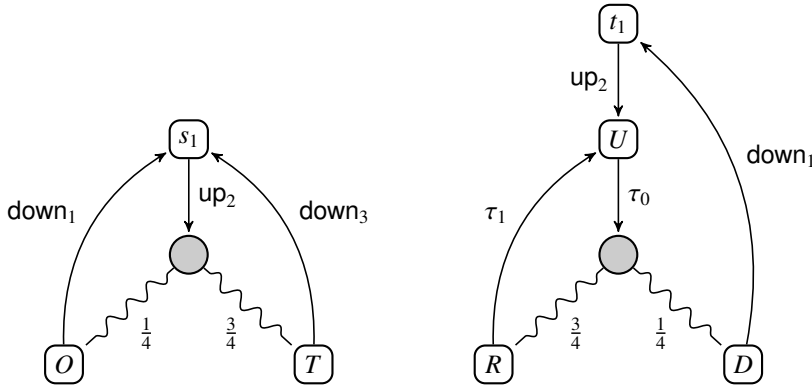
\* Supported by the National Natural Science Foundation of China (61033002, 61011140074).

\*\* Supported by SFI project SFI 06 IN.1 1898.



The first, a two-state machine, continually performs an **up** action, which accrues a benefit of 3 units, followed by a **down** action, which accrues a benefit of 1. The second machine performs the same actions but with benefits 2 and 4 respectively. In some sense  $t_0$  is an improvement on  $s_0$ ; intuitively  $t_0$  can simulate the behaviour of  $s_0$  but in so doing accrue more benefits; this is true even if one of its actions **up** is less beneficial than the corresponding action of  $s_0$ . The same is true for the machine  $u_0$ ; it can also simulate the behaviour of  $s_0$ , with more benefit, although in this case some internal weighted actions, denoted by  $\tau$ , participate in the simulation and add to the accumulation of benefits. In our terminology we will write  $s_0 \sqsubseteq_{sim} t_0$ ,  $s_0 \sqsubseteq_{sim} u_0$ . However, we will have  $t_0 \not\sqsubseteq_{sim} u_0$  because although  $u_0$  can simulate the behaviour of  $t_0$  it accumulates less benefit.

Similar informal reasoning can also be applied to probabilistic systems. Consider the following systems:




The first, from state  $s_1$ , can perform the **up** action with benefit 2 and a quarter of the time it ends up in a state in which **down** can be performed with benefit only 1. But for the remaining three-quarters it ends up in a state in which **down** can be performed for the larger benefit 3. The circular darkened node represents a distribution of states, with its outgoing edges describing the associated probabilities. Again intuitively we can see that  $s_1$  is an improvement on  $s_0$  because it can simulate  $s_0$  and on average accrue slightly more benefits; in our theory we will have  $s_0 \sqsubseteq_{sim} s_1$ .

The mixture of probabilistic behaviour and internal actions introduces complications. Consider the system  $t_1$  above which after performing an **up** probabilistically decides internally whether to perform a **down** action for benefit 1, or branch back to make

another probabilistic choice. However, each time it reverts back it accumulates a non-zero benefit via the internal weighted action  $\tau_1$ , albeit with diminishing probability. Nevertheless, it will turn out that  $s_0 \sqsubseteq_{sim} t_1$  and indeed  $s_1 \sqsubseteq_{sim} t_1$ .

Systems exhibiting both probabilistic and nondeterministic behaviour require more complicated analysis. Consider the following system:



After performing the action `up` it finds itself either in a state in which the action `down` will accrue the benefit 2, or 25% of the time there will be a nondeterministic choice between it accruing either 1 or 6. In the literature there are numerous mechanisms, such as policies, schedulers, adversaries, etc. [12, 14, 13] for resolving such choices. Here one can see if this choice systematically leads to the lower benefit 1 then  $s_2$  will not simulate  $s_0$  as it does not accrue sufficient benefits. This is a pessimistic outlook; an optimistic outlook means that the best choices are systematically made. If this is assumed then we will have  $s_0 \sqsubseteq_{sim} s_2$ ; in  $s_2$  one execution of `up` followed by `down` will yield on average the benefit  $1 + (\frac{3}{4} \cdot 2 + \frac{1}{4} \cdot 6) = 4$ .

The main contribution of the paper is a coinductively defined behavioural preorder  $\sqsubseteq_{sim}$  between MDPs based on simulations which validate the examples discussed informally above. We confine our attention to the optimistic approach to the resolution of nondeterministic choices, although in a later paper we hope to investigate the pessimistic approach. We also show that this preorder is compositional in the sense that it is preserved by structural operators for constructing MDPs from components. The main operator is one for composing two MDPs in parallel. In  $P \mid Q$  the two MDPs  $P$  and  $Q$  remain independent, execute in parallel and may communicate by synchronising on complementary actions; these internal synchronisations accrue the combined benefits of the associated complementary actions.

For finitary convergent MDPs, which are finite-state and finitely branching systems without divergence, we also provide two characterisations for the behavioural preorder  $\sqsubseteq_{sim}$ . The first is in terms of a *qualitative probabilistic logic*  $\mathcal{L}_Q$ . In addition to the standard logical connectives  $\wedge$ ,  $\vee$  and both maximal and minimal fixpoints this contains a novel qualitative *possibility* modality  $\langle \alpha \rangle_w (\phi_1 \oplus_p \phi_2)$ , where  $p$  is some probability between 0 and 1. Intuitively this is satisfied by an MDP which can accrue at least the benefit  $w$  by performing the action  $\alpha$ , and subsequently satisfy the probabilistic assertion  $\phi_1 \oplus_p \phi_2$ . It turns out that the simulation preorder is completely determined by the logic  $\mathcal{L}_Q$ . Further evidence of the compatibility between the logic and the simulation relation is the fact that every system  $P$  has a *characteristic formula*  $\phi(P)$  in the logic which captures its behaviour; informally system  $Q$  can simulate  $P$  if and only if it satisfies the characteristic formula  $\phi(P)$ .

Our second characterisation is in terms of a novel form of testing called *benefits* testing. Intuitively a system  $P$  can be tested by running it in parallel with another testing system  $T$ , and seeing the possible accrued benefits. In the presence of nondeterminism the execution of the combined system  $(T \mid P)$  will result in a non-empty set of benefits,  $\mathbf{Benefits}(T \mid P)$ . Then systems  $P$  and  $Q$  can be compared by comparing the associated benefit sets  $\mathbf{Benefits}(T \mid P)$  and  $\mathbf{Benefits}(T \mid Q)$  where  $T$  ranges over some collection of possible tests. We show that the simulation preorder  $\sqsubseteq_{sim}$  is also determined in this manner by a suitable collection of tests  $T$ .

The rest of this paper is organised as follows. In Section 2 we introduce the model of weighted MDPs, the notation of hyper-derivations and some important properties. Then we define a behavioural preorder based on amortised weighted simulations, which turns out to be a precongruence in a CCS-like process calculus for MDPs. Next, we provide logical and testing characterisations of the behavioural preorder over finitary convergent processes. In Section 3 we present a qualitative probabilistic logic whose formulae completely determine the behavioural preorder. We also show that characteristic formulae can be constructed for any state in such an MDP. In Section 4 we propose a testing framework where our behavioural preorder is sound and complete for may testing preorder. Finally, we conclude in Section 5.

Due to lack of space, we omit all detailed proofs: they are reported in [4].

## 2 Simulations for weighted Markov decision processes

There is considerable variation in the literature in the formal definition of a (labelled) Markov decision process [13, 12]. For the purpose of this paper we use Definition 1.

We first fix some notation. A (discrete) probability subdistribution over a countable set  $S$  is a function  $\Delta : S \rightarrow [0, 1]$  with  $\sum_{s \in S} \Delta(s) \leq 1$ ; the support of such an  $\Delta$  is the set  $[\Delta] = \{s \in S \mid \Delta(s) > 0\}$ . A subdistribution is a (total, or full) distribution if  $\sum_{s \in S} \Delta(s) = 1$ . The point distribution  $\bar{s}$  assigns probability 1 to  $s$  and 0 to all other elements of  $S$ , so that  $[\bar{s}] = s$ . We use  $\mathcal{D}_{sub}(S)$  to denote the set of subdistributions over  $S$ , and  $\mathcal{D}(S)$  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  $S \rightarrow \mathbb{R}$  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  $\Delta_1 + \dots + \Delta_n$ . 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  $\varepsilon$  to denote the everywhere-zero subdistribution that thus has empty support. These operations on subdistributions do not readily adapt themselves to distributions; yet if  $\sum_{k \in K} p_k = 1$  for some collection of  $p_k \geq 0$ , and the  $\Delta_k$  are distributions, then so is  $\sum_{k \in K} p_k \cdot \Delta_k$ .

**Definition 1 (Weighted Markov decision process).** A weighted Markov decision process or *wMDP* is a 4-tuple  $\langle S, \mathbf{A}, W, \longrightarrow \rangle$  where  $S$  is a set of states,  $\mathbf{A}$  a set of actions,  $W$  a set of weights, and  $\longrightarrow \subseteq S \times \mathbf{A} \times W \times \mathcal{D}(S)$ . We normally write  $s \xrightarrow{\alpha}_w \Delta$  to mean  $(s, \alpha, w, \Delta) \in \longrightarrow$ . In this paper we set  $W$  to be  $\mathbb{R}_{\geq 0}$ , the set of non-negative real numbers, and we assume  $\mathbf{A}$  has the structure  $\mathbf{Act}_\tau = \mathbf{Act} \cup \{\tau\}$  where, for the purpose of communication, each  $a$  in  $\mathbf{Act}$  has an inverse  $\bar{a}$  satisfying  $\bar{\bar{a}} = a$ .  $\square$

A wMDP is

- *finite-state* if  $S$  is a finite set;
- *finitely branching* if for each state  $s$ , the set  $\{(\alpha, w, \Delta) \mid s \xrightarrow{\alpha}_w \Delta\}$  is finite;
- *finitary* if it is both finite-state and finitely branching.

In the Introduction we have used a straightforward graphical representation for wMDPs; a state  $s$  is represented by a node  $\boxed{s}$  while darkened circular nodes are used for distributions, and arrows between nodes and distributions are annotated with their weights. Often a point distribution is represented by the unique state in its support; see the first series of examples with initial states  $s_0$ ,  $t_0$  and  $u_0$ .

## 2.1 Hyper-derivations

As we have seen in the Introduction, when reasoning informally that  $t_1$  can simulate  $s_0$ , the limiting behaviour of internal computations must be taken into account. We formalise this by extending the approach originally given in [5].

In a wMDP actions are only performed by states, in that actions are given by relations from states to distributions. But formally, systems or 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 distributions. In fact we will find it convenient to lift them to subdistributions.

**Definition 2.** Let  $\mathcal{R} \subseteq S \times (\mathbb{R}_{\geq 0} \times \mathcal{D}_{\text{sub}}(S))$  be a relation from states to pairs of weights and subdistributions. Then  $\overline{\mathcal{R}} \subseteq \mathcal{D}_{\text{sub}}(S) \times (\mathbb{R}_{\geq 0} \times \mathcal{D}_{\text{sub}}(S))$  is the smallest relation that satisfies:

- (i)  $s \mathcal{R} \langle r, \Theta \rangle$  implies  $\overline{s} \overline{\mathcal{R}} \langle r, \Theta \rangle$ , and
- (ii) (Linearity)  $\Delta_i \overline{\mathcal{R}} \langle r_i, \Theta_i \rangle$  for  $i \in I$  implies  $(\sum_{i \in I} p_i \cdot \Delta_i) \overline{\mathcal{R}} (\sum_{i \in I} p_i \cdot \langle r_i, \Theta_i \rangle)$  for any  $p_i \in [0, 1]$  ( $i \in I$ ) with  $\sum_{i \in I} p_i \leq 1$ .  $\square$

An application of Definition 2 to the arrow relation  $\xrightarrow{\alpha}$  in a wMDP gives a relation  $(\xrightarrow{\alpha}) \subseteq \mathcal{D}(S) \times (W \times \mathcal{D}(S))$ ; for convenience we also denote elements of this relation as  $\Delta \xrightarrow{\alpha}_w \Theta$ . Thus, as source of a relation  $\xrightarrow{\alpha}$  we now also allow distributions, and even subdistributions.

**Definition 3 (Hyper-derivations).** A hyper-derivation consists of a collection of subdistributions  $\Delta, \Delta_k^{\rightarrow}, \Delta_k^{\times}$ , for  $k \geq 0$ , with the following properties:

$$\begin{aligned}
 \Delta &= \Delta_0^{\rightarrow} + \Delta_0^{\times} \\
 \Delta_0^{\rightarrow} &\xrightarrow{\tau}_{w_0} \Delta_1^{\rightarrow} + \Delta_1^{\times} \\
 &\vdots \\
 \Delta_k^{\rightarrow} &\xrightarrow{\tau}_{w_k} \Delta_{k+1}^{\rightarrow} + \Delta_{k+1}^{\times} \\
 &\vdots \\
 \Delta' &= \sum_{k=0}^{\infty} \Delta_k^{\times}
 \end{aligned} \tag{1}$$

Then we call  $\Delta' = \sum_{k=0}^{\infty} \Delta_k^\times$  a hyper-derivative of  $\Delta$ , and write  $\Delta \xrightarrow{\tau}_w \Delta'$ , where  $w = \sum_{k \geq 0} w_k$ . Note that in general  $w \in \mathbb{R}_{\geq 0} \cup \{\infty\}$ ; that is there is no guarantee that the sum  $\sum_{k \geq 0} w_k$  has a finite limit.  $\square$

The reader is referred to [5] for a detailed discussion of the concept of hyper-derivation.

**Example 1.** Consider the wMDP with initial state  $t_1$  discussed in the Introduction. Then we have the following hyper-derivation:

$$\begin{aligned}
\bar{U} &= \bar{U} + \varepsilon \\
\bar{U} &\xrightarrow{\tau}_0 \frac{3}{4} \cdot \bar{R} + \frac{1}{4} \cdot \bar{D} \\
\frac{3}{4} \cdot \bar{R} &\xrightarrow{\tau}_{\frac{3}{4}} \frac{3}{4} \cdot \bar{U} + \varepsilon \\
\frac{3}{4} \cdot \bar{U} &\xrightarrow{\tau}_0 \left(\frac{3}{4}\right)^2 \cdot \bar{R} + \left(\frac{3}{4}\right) \frac{1}{4} \cdot \bar{D} \\
\left(\frac{3}{4}\right)^2 \cdot \bar{R} &\xrightarrow{\tau}_{\left(\frac{3}{4}\right)^2} \left(\frac{3}{4}\right)^2 \cdot \bar{U} + \varepsilon \\
&\vdots \\
\left(\frac{3}{4}\right)^k \cdot \bar{U} &\xrightarrow{\tau}_0 \left(\frac{3}{4}\right)^{(k+1)} \cdot \bar{R} + \left(\frac{3}{4}\right)^k \frac{1}{4} \cdot \bar{D} \\
\left(\frac{3}{4}\right)^{(k+1)} \cdot \bar{R} &\xrightarrow{\tau}_{\left(\frac{3}{4}\right)^{(k+1)}} \left(\frac{3}{4}\right)^{(k+1)} \cdot \bar{U} + \varepsilon \\
&\vdots
\end{aligned}$$

That is,  $\bar{U} \xrightarrow{\tau}_w \sum_{k \geq 0} \left(\frac{3}{4}\right)^k \left(\frac{1}{4} \cdot \bar{D}\right)$  where  $w = \sum_{k \geq 1} \left(\frac{3}{4}\right)^k$ . However this weight evaluates to 3, while the sum of the sub-distributions is the full point distribution  $\bar{D}$ . In other words  $\bar{U} \xrightarrow{\tau}_3 \bar{D}$ .  $\square$

Hyper-derivations satisfy the transitivity property: if  $\Delta \xrightarrow{\tau}_{w_1} \Delta_1$  and  $\Delta_1 \xrightarrow{\tau}_{w_2} \Delta_2$  then  $\Delta \xrightarrow{\tau}_{w_1+w_2} \Delta_2$ . The generation of a hyper-derivative is in general highly non-deterministic. In (1) of Definition 3 the calculation of  $\Delta_{k+1}^\rightarrow$  and  $\Delta_{k+1}^\times$  from  $\Delta_k^\rightarrow$  involves making nondeterministic choices. But these choices can be governed by *policies*.

**Definition 4 (Static policies).** A static policy for a wMDP is a partial function  $\text{pp} : S \rightarrow \mathbb{R}_{\geq 0} \times \mathcal{D}(S)$  such that if  $\text{pp}(s) = \langle w, \Delta \rangle$  then  $s \xrightarrow{\tau}_w \Delta$ .  $\square$

Let us write  $\Delta \xrightarrow{\tau}_{\text{pp}, w} \Delta'$  to mean that the hyper-derivative  $\Delta'$  is generated with weight  $w$  from  $\Delta$  using the policy  $\text{pp}$ . Formally this means that in (1) of Definition 3 the weights  $w_k$  and subdistributions  $\Delta_k^\rightarrow$ ,  $\Delta_k^\times$  are calculated as follows:

- $s \in [\Delta_k^\times]$  if and only if  $\text{pp}(s)$  is undefined
- $s \in [\Delta_k^\rightarrow]$  if and only if  $\text{pp}(s)$  is defined
- $\langle w_{k+1}, \Delta_{k+1}^\rightarrow + \Delta_{k+1}^\times \rangle = \sum_{s \in [\Delta_k^\rightarrow]} \Delta_k^\rightarrow(s) \cdot \text{pp}(s)$  for all  $k \geq 0$ .

**Theorem 1 (Finite generability).** Let  $\text{pp}_1, \dots, \text{pp}_n$  ( $n \geq 1$ ) be all the static policies in a finitary wMDP. Suppose  $\Delta \xrightarrow{\tau}_{\text{pp}_i, w_i} \Delta'_i$  and  $w_i < \infty$  for all  $1 \leq i \leq n$ . If  $\Delta \xrightarrow{\tau}_w \Delta'$  then there are probabilities  $p_i$  for all  $1 \leq i \leq n$  with  $\sum_{i=1}^n p_i = 1$  such that  $\langle w, \Delta' \rangle = \sum_{i=1}^n p_i \cdot \langle w_i, \Delta'_i \rangle$ .  $\square$

In later developments it will be important to rule out the possibility of hyper-derivatives generating an infinite weight.

**Definition 5.** A *wMDP* is convergent if  $\bar{s} \xrightarrow{w} \varepsilon$  for no state  $s$  and weight  $w$ ; in other words there is no divergent internal computation from any state. A *wMDP* is bounded if it is finitary and whenever  $\Delta \xrightarrow{w} \Delta'$  then  $w \in \mathbb{R}_{\geq 0}$ .  $\square$

A simple source of unboundedness is divergence. Consider the trivial *wMDP* consisting of one state  $s$  and one arrow  $s \xrightarrow{1} \bar{s}$ . This is not a bounded *wMDP* because  $\bar{s} \xrightarrow{\infty} \bar{s}$ . In fact for finitary *wMDPs*, this is the only source of unboundedness:

**Theorem 2.** Every finitary convergent *wMDP* is also bounded.  $\square$

The proof of the above theorem relies on Theorem 1. Another important consequence of Theorem 1 is the following.

**Corollary 1.** In a bounded *wMDP*, for every  $\Delta$  the set  $\{\langle w, \Delta' \rangle \mid \Delta \xrightarrow{w} \Delta'\}$  is compact, in the standard Euclidean topology.  $\square$

## 2.2 (Amortised weighted) simulations

Here we assume some *wMDP*  $\langle S, \text{Act}_\tau, \mathbb{R}_{\geq 0}, \longrightarrow \rangle$ . Our simulation relation is parametrised on an initial investment  $r \in \mathbb{R}_{\geq 0}$  and relates states to distributions, rather than states to states. It also uses *weak* arrow relations, defined using hyper-derivations: we write  $\Delta \xrightarrow{w} \Theta$  whenever  $\Delta \xrightarrow{w_1} \Delta' \xrightarrow{w_2} \Theta' \xrightarrow{w_3} \Theta$ , where  $w = w_1 + w_2 + w_3$ .

**Definition 6.** Given a relation  $\mathcal{R} \subseteq S \times (\mathbb{R}_{\geq 0} \times \mathcal{D}(S))$ , let  $\mathcal{S}(\mathcal{R}) \subseteq S \times (\mathbb{R}_{\geq 0} \times \mathcal{D}(S))$  be the relation defined by letting  $s \mathcal{S}(\mathcal{R}) \langle r, \Theta \rangle$  whenever

$s \xrightarrow{v} \Delta$  implies the existence of some  $w$  and  $\Theta'$  with  $\Theta \xrightarrow{w} \Theta'$  and  $\Delta \bar{\mathcal{R}} \langle r+w-v, \Theta' \rangle$ .

We say  $\mathcal{R}$  is an (amortised weighted) simulation if  $\mathcal{R} \subseteq \mathcal{S}(\mathcal{R})$ . The operator  $\mathcal{S}(-)$  is (pointwise) monotonic and so it has a maximal fixpoint, which is also a simulation, and which we denote by  $\triangleleft$ . We often write  $s \triangleleft_r \Theta$  for  $s \triangleleft \langle r, \Theta \rangle$  and use  $\Delta \sqsubseteq_{sim} \Theta$  to mean that there is some initial investment  $r$  such that  $\Delta \bar{\triangleleft}_r \Theta$ .  $\square$

The basic idea here is that  $s \triangleleft_r \Theta$  intuitively means that  $\Theta$  can simulate the actions of  $s$  but with *more benefit*, or at least not less benefit. The parameter  $r$  should be viewed as compensation which  $\Theta$  has accumulated and can be used in local comparisons between the benefits of individual actions. Thus when we simulate  $s \xrightarrow{v} \Delta$  with  $\Theta \xrightarrow{w} \Theta'$  there are two possibilities:

- (i)  $w > v$ ; here the accumulated compensation is increased from  $r$  to  $r + (w - v)$ . In subsequent rounds this extra compensation may be used to successfully simulate a heavier action with a lighter one.
- (ii)  $w \leq v$ ; here the compensation is decreased from  $r$  to  $r - (v - w)$ .

Finally it is important that  $r \geq 0$ , and remains greater than or equal to zero, or otherwise the presence of weights would have no effect. Thus in case (ii) if  $(v - w) > r$  then the attempted simulation is not successful.

We now show that with this formal definition of the relation  $\sqsubseteq_{sim}$  the various statements asserted in the Introduction are true:

$$\begin{array}{l}
\text{(L-ACT)} \\
\alpha_w \cdot (\oplus_{i \in I} P_i \cdot P_i) \xrightarrow{\alpha}_w \text{Dist}(\{(P_i, P_i) \mid i \in I\}) \\
\text{(L-COMM)} \\
\frac{P_1 \xrightarrow{a}_{w_1} \Delta_1, P_2 \xrightarrow{\bar{a}}_{w_2} \Delta_2}{P_1 \mid P_2 \xrightarrow{\tau}_w \Delta_1 \mid \Delta_2} \quad w = w_1 + w_2 \\
\text{(L-HIDE)} \\
\frac{P \xrightarrow{\alpha}_w \Delta}{P \setminus a \xrightarrow{\alpha}_w \Delta \setminus a} \quad \alpha \neq a, \bar{a} \\
\text{(L-ALT)} \\
\frac{P_1 \xrightarrow{\alpha}_w \Delta}{P_1 + P_2 \xrightarrow{\alpha}_w \Delta} \\
\text{(L-PAR)} \\
\frac{P_1 \xrightarrow{\alpha}_w \Delta}{P_1 \mid P_2 \xrightarrow{\alpha}_w \Delta \mid \bar{P}_2} \\
\text{(L-DEF)} \\
\frac{P_A \xrightarrow{\alpha}_w \Delta}{A \xrightarrow{\alpha}_w \Delta} \quad A \Leftarrow P_A
\end{array}$$

Fig. 1. Weighted actions for CCMDP

**Example 2.** Consider the first two systems,  $s_0$  and  $t_0$ , viewed as wMDPs. Then the relation  $\mathcal{R}$  given by  $\mathcal{R} = \{(s_0, \langle r, \bar{t}_0 \rangle) \mid r \geq 1\} \cup \{(s_d, \langle r, \bar{t}_d \rangle) \mid r \geq 0\}$  is a simulation. Thus  $s_0 \triangleleft_r \bar{t}_0$  for any  $r \geq 1$ . As pointed out in [11] this example shows the need for the parametrisation with respect to initial investments  $r$ ; Because of the weights associated with the action up an initial investment of at least one is required in order for  $\bar{t}_0$  to be able to match  $s_0$ .

We also have  $s_0 \triangleleft_r \bar{s}_1$  for any  $r \geq 1$  because of the following simulation:

$$\mathcal{R} = \{(s_0, \langle r, \bar{s}_1 \rangle) \mid r \geq 1\} \cup \{(s_d, \langle r, \bar{\Delta} \rangle) \mid r \geq 0\}$$

where  $\bar{\Delta}$  is the distribution  $\frac{1}{4} \cdot \bar{O} + \frac{3}{4} \cdot \bar{T}$ . Note that this is indeed a simulation because  $\bar{\Delta} \xrightarrow{\text{down}}_{2.5} \bar{s}_1$ . Incidentally this example shows the necessity of relating states to distributions, rather than states; no individual state accessible from  $s_1$  can simulate  $s_d$ .

Similarly  $s_1 \triangleleft_r \bar{t}_1$  for every  $r \geq 0$  because of the simulation:

$$\mathcal{R} = \{(s_1, \langle r, \bar{t}_1 \rangle) \mid r \geq 0\} \cup \{(O, \langle r, \bar{U} \rangle) \mid r \geq 0\} \cup \{(T, \langle r, \bar{U} \rangle) \mid r \geq 0\}$$

This relies on the fact that  $\bar{U} \xrightarrow{\text{down}}_4 \bar{t}_1$ , which follows by transitivity, since we have already seen in Example 1 that  $\bar{U} \xrightarrow{\tau}_3 \bar{D}$ .

Finally  $s_0 \triangleleft_2 s_2$  because of the following simulation:

$$\mathcal{R} = \{(s_0, \langle r, \bar{s}_2 \rangle) \mid r \geq 2\} \cup \{(s_d, \langle r, \bar{\Delta} \rangle) \mid r \geq 0\}$$

where  $\bar{\Delta}$  is the distribution  $\frac{1}{4} \cdot \bar{S} + \frac{3}{4} \cdot \bar{T}$ . Note that  $\bar{\Delta} \xrightarrow{\text{down}}_3 \bar{s}_2$  although it is also possible for it to do the **down** action for much less benefit.  $\square$

The simulation relations  $\triangleleft_r$  are defined coinductively. But in bounded wMDPs they can also be characterised inductively.

**Definition 7.** For every  $k \geq 0$  we define the relation  $\triangleleft^k \subseteq S \times (\mathbb{R}_{\geq 0} \times \mathcal{D}(S))$  as follows:

- (i)  $\triangleleft^0 = S \times (\mathbb{R}_{\geq 0} \times \mathcal{D}(S))$
- (ii)  $\triangleleft^{k+1} = \mathcal{S}(\triangleleft^k)$  for every  $k \geq 0$ .

Finally we let  $\triangleleft^\infty$  be  $\bigcap_{k=0}^\infty \triangleleft^k$ .  $\square$

**Theorem 3.** In a bounded wMDP, the two relations  $\triangleleft$  and  $\triangleleft^\infty$  coincide.  $\square$

Corollary 1 plays a crucial role in proving the above theorem.

The simplest approach to discussing compositionality is, as in [8], to introduce a process calculus-like syntax for wMDPs. Our calculus, CCMDP, is based on CCS:

$$P ::= \alpha_w \cdot (\oplus_{i \in I} P_i \cdot P_i) \mid P \mid P \mid P + P \mid \mathbf{0} \mid P \setminus a \mid A \quad (2)$$



The main operator is prefixing,  $\alpha_w.(\oplus_{i \in I} p_i \cdot P_i)$ . Here  $\alpha$  is taken from  $\text{Act}_\tau$ ,  $w$  from  $\mathbb{R}_{\geq 0}$ ,  $I$  is a non-empty finite index set and  $p_i$  are probabilities satisfying  $\sum_{i \in I} p_i = 1$ . We also assume a set of definitional constants, ranged over by  $A$ , and each such  $A$  has a definition associated with it, a process term  $P_A$ . We often write these definitions as  $A \Leftarrow P_A$ .

Intuitively, we view each process term as describing a wMDP. Formally we describe one overarching wMDP where the states are all terms  $P$  in the grammar (2) and the weighted actions  $P \xrightarrow{\alpha}_w \Delta$  are those which can be derived by the rules in Figure 1; obvious symmetric counterparts to the rules (L-ALT) (L-PAR) are omitted. In rule (L-ACT) we use the obvious notation  $\mathcal{D}ist(\{(p_i, P_i) \mid i \in I\})$  for constructing a distribution from the formal term  $\oplus_{i \in I} p_i \cdot P_i$ . In rules (L-COMM) and (L-PAR) we use an abbreviation for distributing parallel composition over a distribution: e.g.  $\Delta_1 \mid \Delta_2$  is the distribution given by  $(\Delta_1 \mid \Delta_2)(R) = \Delta_1(P_1) \cdot \Delta_2(P_2)$  if  $R = P_1 \mid P_2$  and 0 otherwise. Similar is the hiding operator in (L-HIDE). Note that all of the wMDPs described graphically in the Introduction can be described in CCMDP. In the sequel we will not distinguish between the syntactic term  $P$ , its interpretation as a state in the wMDP defined in Figure 1, and the wMDP it induces by considering only those states accessible from it.

**Theorem 4 (Compositionality).** *The preorders  $\bar{\prec}_r$ , for each  $r \in \mathbb{R}_{\geq 0}$ , are preserved by each of the operators in the language CCMDP.  $\square$*

**Example 3.** Let  $P, Q$  be two processes with  $P <_0 \bar{Q}$ . Consider the following processes:

$$\begin{aligned} U &\Leftarrow \tau_0.(\tau_1.U_{\frac{3}{4}} \oplus \text{down}_1.Q) \\ P' &\equiv \text{up}_2.(\text{down}_1.P_{\frac{1}{4}} \oplus \text{down}_3.P) \\ Q' &\equiv \text{up}_2.U \end{aligned}$$

By the analysis in Example 1 we know that  $\bar{U} \xrightarrow{\tau}_3 \overline{\text{down}_1.Q}$ , thus  $\bar{U} \xrightarrow{\text{down}}_4 \bar{Q}$ . Then it is easy to see that  $\text{down}_1.P <_0 \bar{U}$  and  $\text{down}_3.P <_0 \bar{U}$ . It follows from the compositionality of  $\bar{\prec}_0$  that  $(\text{down}_1.P_{\frac{1}{4}} \oplus \text{down}_3.P) \bar{\prec}_0 \bar{U}$  and furthermore  $\bar{P}' \bar{\prec}_0 \bar{Q}'$ .  $\square$

### 3 A qualitative probabilistic logic

Let  $\text{Var}$  be a set of variables, ranged over by  $X$ . We define the set of formulae as follows:

$$\begin{aligned} \phi ::= & \text{tt} \mid \text{ff} \mid \langle \alpha \rangle_w (\phi_1 \oplus \phi_2), \alpha \in \text{Act}_\tau, w \in \mathbb{R}_{\geq 0}, p \in [0, 1] \\ & \mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2 \mid X \mid \min X. \phi \mid \max X. \phi \end{aligned}$$

The two fixpoint operators  $\min X. -$  and  $\max X. -$  act as binders in the standard manner; we use  $\mathcal{L}_Q$  to denote the set of *closed* formulae, that is containing no free variables. As a shorthand, we write  $\langle \alpha \rangle_w \phi$  for  $\langle \alpha \rangle_w (\phi_1 \oplus \phi')$  for any  $\phi'$ .

Let  $\text{Con}$  denote the set of *configurations*, pairs  $\langle r, \Delta \rangle$  where  $r \in \mathbb{R}_{\geq 0}$  and  $\Delta \in \mathcal{D}(S)$ , with  $S$  denoting the state space of some wMDP. Intuitively this represents a system which has accumulated compensation  $r$  which it can use to satisfy formulae in the future. A formula from  $\mathcal{L}_Q$  determines a set of configurations, those which satisfy it; their calculation is standard, apart from the novel qualitative possibility operator. An environment  $\rho$  is a function that maps each variable in  $\text{Var}$  to a subset of  $\text{Con}$ . For a set  $V \subseteq \mathbb{R}_{\geq 0} \times \mathcal{D}(S)$  and a variable  $X \in \text{Var}$ , we write  $\rho[X \mapsto V]$  for the environment that maps  $X$  to  $V$  and  $Y$  to  $\rho(Y)$  for all  $Y \neq X$ . The semantics of a formula  $\phi$  is given by the set of configurations  $\llbracket \phi \rrbracket_\rho$  defined as follows:

- $\llbracket \text{tt} \rrbracket_\rho = \text{Con}, \quad \llbracket \text{ff} \rrbracket_\rho = \emptyset$
- $\llbracket \phi_1 \wedge \phi_2 \rrbracket_\rho = \llbracket \phi_1 \rrbracket_\rho \cap \llbracket \phi_2 \rrbracket_\rho, \quad \llbracket \phi_1 \vee \phi_2 \rrbracket_\rho = \llbracket \phi_1 \rrbracket_\rho \cup \llbracket \phi_2 \rrbracket_\rho$
- $\llbracket \langle \alpha \rangle_v(\phi_1 \oplus \phi_2) \rrbracket_\rho = \{ \langle r, \Delta \rangle \mid \Delta \xrightarrow{\alpha}_w \Theta \text{ where } \langle r+w-v, \Theta \rangle = \langle r_1, \Theta_1 \rangle \oplus \langle r_2, \Theta_2 \rangle \text{ and } \langle r_i, \Theta_i \rangle \in \llbracket \phi_i \rrbracket_\rho \}$
- $\llbracket X \rrbracket_\rho = \rho(X)$
- $\llbracket \min X. \phi \rrbracket_\rho = \bigcap \{ V \mid \llbracket \phi \rrbracket_{\rho[X \mapsto V]} \subseteq V \}$
- $\llbracket \max X. \phi \rrbracket_\rho = \bigcup \{ V \mid V \subseteq \llbracket \phi \rrbracket_{\rho[X \mapsto V]} \}$

When  $\phi$  is closed the set  $\llbracket \phi \rrbracket_\rho$  is independent of the environment  $\rho$ , and in this case we use the standard notation  $C \models \phi$  in place of  $C \in \llbracket \phi \rrbracket$ .

The novel qualitative formula  $\langle \alpha \rangle_v(\phi_1 \oplus \phi_2)$  represents the ability to do an  $\alpha$  action with benefit at least  $v$  and then probabilistically satisfy the property  $\phi_1 \oplus \phi_2$ ; we have  $\langle r, \Delta \rangle \models \langle \alpha \rangle_v(\phi_1 \oplus \phi_2)$  whenever  $\Delta \xrightarrow{\alpha}_w \Theta_1 \oplus \Theta_2$  and  $\langle r_i, \Theta_i \rangle \models \phi_i$  for some  $r_i$  satisfying  $r + w - v = p \cdot r_1 + (1 - p) \cdot r_2$ . Here there are two possibilities:

- (i)  $v > w$ : here the compensation comes into play. The action may be accepted despite being too heavy but the compensation for future use is reduced from  $r$  to  $r - (v - w)$ ; this is split into  $r_1, r_2$  via the probability  $p$ . Note this possibility will only exist if  $r - (v - w) \geq 0$ .
- (ii)  $v \leq w$ : The action is accepted and then the compensation is increased from  $r$  to  $r + (w - v)$ , which again is split proportionally into  $r_1, r_2$ , to satisfy  $\phi_1$  and  $\phi_2$  respectively.

**Example 4.** Both liveness and safety properties can be expressed in  $\mathcal{L}_Q$ . For example, suppose  $AB$  denote the formula  $\langle a \rangle_0(\langle b \rangle_{10} \text{tt} \oplus_{\frac{9}{10}} \text{tt})$  and  $C$  is a configuration. Then  $C \models AB$  means that  $C$  can perform an  $a$  action such that at least 90% of the time it can subsequently perform a  $b$  action with a benefit of at least 10. So the formula

$$\min X. \langle \text{up} \rangle_0(\langle \text{down} \rangle_{10} X \oplus_{\frac{9}{10}} \text{tt}) \vee \langle \text{up} \rangle_0 X$$

expresses the liveness property of being able to perform a sequence of  $\text{up}$  actions to arrive at a state where at least 90% of the time a  $\text{down}$  action for benefit at least 10 can be performed. On the other hand, the formula

$$\max X. \langle \text{up} \rangle_0(\langle \text{down} \rangle_{10} X \oplus_{\frac{9}{10}} \text{tt}) \wedge \langle \text{stay} \rangle_0 X$$

expresses the safety property of always being able to perform a  $\text{stay}$  action and at the same time to perform an  $\text{up}$  action to arrive at a state where at least 90% of the time a  $\text{down}$  action for benefit at least 10 can be performed.  $\square$

Let  $\mathcal{L}_Q(r, \Delta) = \{ \phi \in \mathcal{L}_Q \mid \langle r, \Delta \rangle \models \phi \}$ . We have the following logical characterisation of simulations.

**Theorem 5 (Logical characterisation).** *In a bounded wMDP,  $s \triangleleft_r \Theta$  if and only if  $\mathcal{L}_Q(0, \bar{s}) \subseteq \mathcal{L}_Q(r, \Theta)$ .*  $\square$

The proof of Theorem 5 exploits Theorem 3, which says that  $\triangleleft$  can be approximated by a family of stratified relations  $\triangleleft^k$  for  $k \in \mathbb{N}$ . So it suffices to prove that each approximant  $\triangleleft^k$  is completely determined by the finite fragment of the qualitative probabilistic logic.

The import of Theorem 5 is that if  $s \triangleleft_r \Theta$  does not hold then there is a formula  $\phi$  from  $\mathcal{L}_Q$  which  $\bar{s}$  satisfies but  $\Theta$  does not. Furthermore, it turns out that in a bounded wMDP this distinguishing formula will always be finite; that is contains no occurrence

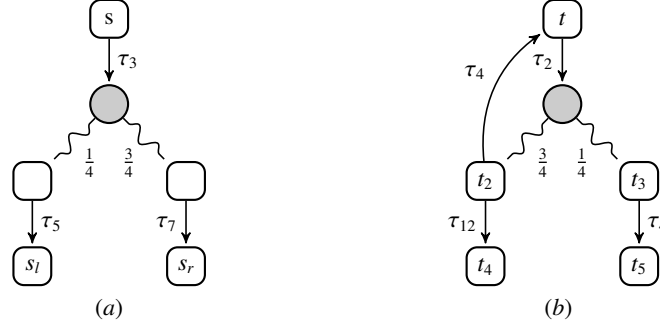


Fig. 2. Testing systems

of a fixpoint operator. For example,  $\overline{s_2} \not\sim_0 \overline{s_0}$ , where these are defined in the Introduction, because of the distinguishing formula  $\langle \text{up} \rangle_1 (\langle \text{down} \rangle_6 \text{tt} \frac{1}{4} \oplus \langle \text{down} \rangle_2 \text{tt})$ .

Our logic is expressive enough so that the whole behaviour of a state in a bounded wMDP can be captured by one formula in the logic:

**Theorem 6 (Characteristic formula).** *In a bounded wMDP, for every state  $s$  there is a characteristic formula  $\phi(s) \in \mathcal{L}_Q$  such that  $s \prec_r \Theta$  if and only if  $\langle r, \Theta \rangle \models \phi(s)$ .  $\square$*

For example, the state  $s_1$  in the Introduction has the following characteristic formula:  $\phi(s_1) = \max X. \langle \text{up} \rangle_2 (\langle \text{down} \rangle_1 X \frac{1}{4} \oplus \langle \text{down} \rangle_3 X)$ .

## 4 Benefits based testing

Standard theories of testing involve the idea of applying tests to processes and seeing if the result is a *success*. With the presence of weights in wMDPs we have a more elementary way of testing; we run them in parallel with other wMDPs and calculate the possible benefits which can be accrued. Then two wMDPs can be compared by examining the resulting sets of possible accrued benefits.

Consider the simple fully probabilistic wMDP in Figure 2(a), which results from running the test  $T = \overline{\text{up}}_1 \cdot \text{down}_4 \cdot \mathbf{0}$  in parallel with the system  $s_1$  from the Introduction. Formally this is the sub-wMDP of the wMDP  $(\overline{s_1} \mid T)$  obtained by concentrating on the internal actions  $\tau_w$ ; this is just the wMDP represented by  $(\overline{s_1} \mid T) \setminus \text{Act}$  that we denote by  $\overline{s_1} \parallel T$ . Every time the experiment runs we get the initial benefit 3; three-quarters of the time we also get the benefit 7 while a quarter of time we get 5. So the total benefit is  $3 + \frac{3}{4} \cdot 7 + \frac{1}{4} \cdot 5 = 9.5$ . In the presence of nondeterminism there will in general be a set of possible benefits, depending on the way in which the nondeterminism is resolved. Traditionally this resolution is expressed in terms of a scheduler, or adversary, which for each state decides which of its successors is chosen for execution, with the resulting set of benefits consequently depending on the choice of scheduler. Here we take a more abstract approach, following [5], and essentially allow arbitrary schedulers.

**Definition 8 (Extreme derivatives).** *For any  $\Delta$  in a wMDP we write  $\Delta \xRightarrow{\tau}_w \Phi$  if*

- $\Delta \xRightarrow{\tau}_w \Phi$ , that is  $\Phi$  is a hyper-derivative of  $\Delta$
- $\Phi$  is stable, that is  $s \xrightarrow{\tau}$  for every  $s$  in  $[\Phi]$

where  $s \xrightarrow{\tau}$  means that  $s$  cannot enable any  $\tau$ -transition. We say  $\Phi$  is an extreme derivative of  $\Delta$ , with weight  $w$ .  $\square$

Intuitively every extreme derivative  $\Delta \xrightarrow{w} \Phi$  represents a computation from the initial distribution  $\Delta$  guided by some implicit scheduler. For example, consider the hyper-derivation of an extreme derivative:

$$\begin{aligned}
\Delta &= \Delta_0^{\rightarrow} + \Delta_0^{\times} \\
\Delta_0^{\rightarrow} &\xrightarrow{w_0} \Delta_1^{\rightarrow} + \Delta_1^{\times} \\
&\vdots \\
\Delta_k^{\rightarrow} &\xrightarrow{w_k} \Delta_{k+1}^{\rightarrow} + \Delta_{k+1}^{\times} \\
&\vdots \\
\Phi &= \sum_{k=0}^{\infty} \Delta_k^{\times}
\end{aligned} \tag{3}$$

where  $w = \sum_{k \geq 0} w_k$ . Initially, since  $\Delta_0^{\times}$  is stable,  $\Delta_0^{\rightarrow}$  contains (in its support) all states which can proceed with the computation. The implicit scheduler decides for each of these states which step to take, cumulating in the first move,  $\Delta_0^{\rightarrow} \xrightarrow{w_0} \Delta_1^{\rightarrow} + \Delta_1^{\times}$ . At an arbitrary stage,  $\Delta_k^{\rightarrow}$  contains all states which can continue; the scheduler decides which step to take for each individual state and the overall result of the scheduler's decision for this stage is captured in the step  $\Delta_k^{\rightarrow} \xrightarrow{w_k} \Delta_{k+1}^{\rightarrow} + \Delta_{k+1}^{\times}$ .

**Example 5.** Referring to Figure 2(a) it is easy to see that  $\bar{s}$  has a unique (degenerate) extreme derivative,  $\bar{s}_1 \xrightarrow{9.5} (\frac{1}{4}\bar{s}_l + \frac{3}{4}\bar{s}_r)$ , intuitively representing the unique weighted computation from  $\bar{s}_1$ . However, consider the wMDP in Figure 2(b), in which there is a nondeterministic choice from state  $t_2$ ; here the extreme derivatives generated from  $\bar{t}$ , and their associated weights, will depend on the choices made during the computation by the implicit scheduler.

First suppose that the scheduler uses the static policy which maps  $t_2$  to  $\langle 12, \bar{t}_4 \rangle$ . Then it is easy to see that the generated extreme derivative, which is degenerate, is  $\bar{t} \xrightarrow{12} (\frac{3}{4}\bar{t}_4 + \frac{1}{4}\bar{t}_5)$ . However using the static policy which maps  $t_2$  to  $\langle 4, \bar{t}_1 \rangle$  we generate, using (3), a non-degenerate extreme derivative; after some calculations this can be seen to be  $\bar{t}_1 \xrightarrow{24} \bar{t}_5$ .

However there are many other possible implicit schedulers, for example at different times in the computations employing either of these static policies, or even choosing nondeterministically between them. But these are the only static policies and therefore we know from Theorem 1 that if  $\bar{t}_1 \xrightarrow{w} \Delta$  then  $w$  must take the form  $p \cdot 12 + (1-p) \cdot 24$  for some  $0 \leq p \leq 1$ . That is the set of benefits which can be generated from  $\bar{t}_1$  is  $\{24 - 12 \cdot p \mid 0 \leq p \leq 1\}$ .  $\square$

**Definition 9 (May testing).** In a wMDP, for any  $\Delta \in \mathcal{D}(S)$ , let

$$\text{Benefits}(\Delta) = \{w \in \mathbb{R}_{\geq 0} \mid \Delta \xrightarrow{w} \Phi, \text{ for some } \Phi \in \mathcal{D}_{\text{sub}}(S)\}$$

Benefit sets are compared as follows:

$$B_1 \leq_{\text{Ho}}^r B_2 \text{ if for every } r_1 \in B_1 \text{ there exists some } r_2 \in B_2 \text{ such that } r_1 \leq r + r_2$$

For any two distributions  $\Delta$ ,  $\Theta$  we write  $\Delta \sqsubseteq_{\text{may}}^r \Theta$  if for every finite (testing) process  $T$ ,  $\mathbf{Benefits}(\Delta \parallel T) \leq_{\text{Ho}}^r \mathbf{Benefits}(\Theta \parallel T)$ . We write  $\Delta \sqsubseteq_{\text{may}} \Theta$  to mean that there is some  $r \in \mathbb{R}_{\geq 0}$  such that  $\Delta \sqsubseteq_{\text{may}}^r \Theta$ .  $\square$

This interpretation of processes is optimistic;  $\Delta \sqsubseteq_{\text{may}}^r \Theta$  means that, given the investment  $r$ , every possible benefit produced by  $\Delta$  can in principle be improved upon by  $\Theta$ .

Note that in a bounded wMDP  $\mathbf{Benefits}(\Delta)$  cannot contain  $\infty$ . Moreover we can show that the parallel composition of a bounded wMDP with a finite wMDP is also bounded. This means that if we confine our attention to bounded wMDPs then benefit sets will always only contain real numbers. One way of restricting to bounded wMDPs is, by Theorem 2, to only use finitary convergent wMDPs.

Our first result shows that simulations can be used as a sound proof technique for this semantics:

**Theorem 7 (Soundness).**  $\Delta \bar{\sqsubseteq}_r \Theta$  implies  $\Delta \sqsubseteq_{\text{may}}^r \Theta$ .  $\square$

The converse is not true in general:

**Example 6.** Consider the two distributions  $\Delta = \mathbf{0} \frac{1}{2} \oplus a_1 \cdot \mathbf{0}$  and  $\Theta = \tau_2 \cdot \mathbf{0} \frac{1}{2} \oplus a_0 \cdot \mathbf{0}$ . It is easy to see that  $\Delta \bar{\mathcal{A}}_0 \Theta$  because there is no way to decompose  $\Theta$  into  $\Theta_1 \frac{1}{2} \oplus \Theta_2$  for some  $\Theta_1, \Theta_2$  such that  $a_1 \cdot \mathbf{0} \triangleleft_0 \Theta_2$ . However, one can show that  $\Delta \sqsubseteq_{\text{may}}^0 \Theta$ . This follows from the observations below:

- (i) For all weights  $w$  and tests  $T$ ,  $\mathbf{Benefits}(\tau_w \cdot \mathbf{0} \parallel T) = \{v + w \mid v \in \mathbf{Benefits}(\mathbf{0} \parallel T)\}$ .
- (ii) For all weights  $w$  and tests  $T$ ,  $\mathbf{Benefits}(a_w \cdot \mathbf{0} \parallel T) \leq_{\text{Ho}}^w \mathbf{Benefits}(a_0 \cdot \mathbf{0} \parallel T)$ .

Both assertions can be proved by structural induction on  $T$ .

Now suppose  $w \in \mathbf{Benefits}(\Delta \parallel T)$  for an arbitrary test  $T$ . There is some stable derivative  $\Gamma$  such that  $\Delta \parallel T \xrightarrow{\tau}_w \Gamma$ . It can be shown that there are some  $w_1, w_2, \Gamma_1, \Gamma_2$  with  $\mathbf{0} \parallel T \xrightarrow{\tau}_{w_1} \Gamma_1$ ,  $a_1 \cdot \mathbf{0} \parallel T \xrightarrow{\tau}_{w_2} \Gamma_2$ ,  $w = \frac{1}{2}w_1 + \frac{1}{2}w_2$ , and  $\Gamma = \frac{1}{2} \cdot \Gamma_1 + \frac{1}{2} \cdot \Gamma_2$ , where both  $\Gamma_1$  and  $\Gamma_2$  are stable. In other words, we have  $w_1 \in \mathbf{Benefits}(\mathbf{0} \parallel T)$  and  $w_2 \in \mathbf{Benefits}(a_1 \cdot \mathbf{0} \parallel T)$ . By (i) above,  $w_1 + 2 \in \mathbf{Benefits}(\tau_2 \cdot \mathbf{0} \parallel T)$ ; by (ii) above, there exists some  $w'_2 \in \mathbf{Benefits}(a_0 \cdot \mathbf{0} \parallel T)$  with  $w_2 \leq w'_2 + 1$ . Thus, we can infer that

$$\begin{aligned} w &= \frac{1}{2}w_1 + \frac{1}{2}w_2 \\ &< \frac{1}{2}(w_1 + 2) + \frac{1}{2}(w_2 - 1) \\ &\leq \frac{1}{2}(w_1 + 2) + \frac{1}{2}w'_2 \end{aligned}$$

It turns out that  $\frac{1}{2}(w_1 + 2) + \frac{1}{2}w'_2 \in \mathbf{Benefits}(\Theta \parallel T)$ . Therefore, we have

$$\mathbf{Benefits}(\Delta \parallel T) \leq_{\text{Ho}}^0 \mathbf{Benefits}(\Theta \parallel T).$$

Since this reasoning is carried out for an arbitrary test  $T$ , it follows that  $\Delta \sqsubseteq_{\text{may}}^0 \Theta$ .  $\square$

Nevertheless we do have a testing characterisation for the unannotated simulation preorder:

**Theorem 8 (Testing characterisation).** *In a bounded wMDP,  $\bar{s} \sqsubseteq_{\text{may}} \Theta$  if and only if  $\bar{s} \sqsubseteq_{\text{sim}} \Theta$ .*

*Proof (Outline).* One direction follows from Theorem 7. For the converse we carry out the proof in two steps: we first prove that  $\bar{s} \sqsubseteq_{\text{may}}^r \Theta$  implies the existence of some compensation  $r' \geq r$  with  $\mathcal{L}_Q(0, \bar{s}) \subseteq \mathcal{L}_Q(r', \Theta)$ , then appeal to Theorem 5. In the first step we proceed by constructing, for each formula  $\phi$ , a *characteristic test*  $T(\phi)$ , such that if a process satisfies  $\phi$  then it passes test  $T(\phi)$  with some threshold benefit.  $\square$

An alternative approach to testing would be to use one special action  $\omega$  in a test to report success and when applying such a test to a system to report the weighted average of the weight of each path leading to an occurrence of the success action; this we refer to as *expected benefits testing*. Here we will not give the formal definition of how these *expected benefits* are calculated, which is provided in [4], but simply give an informal argument to show that our simulation preorder is not sound with respect to it.

**Example 7 (Simulation is unsound for expected benefits testing).** Consider the following processes:

$$\begin{aligned} P &= \tau_2.(\mathbf{0}_{\frac{1}{4}} \oplus a_0. \mathbf{0}) \\ Q &= \tau_1.(\tau_2.(\mathbf{0}_{\frac{1}{2}} \oplus a_0. \mathbf{0})_{\frac{1}{2}} \oplus a_0. \mathbf{0}) \end{aligned}$$

It is easy to see that  $P \prec_0 Q$  as the transition  $P \xrightarrow{\tau} \mathbf{0}_{\frac{1}{4}} \oplus a_0. \mathbf{0}$  can be simulated by the hyper-derivative  $Q \xrightarrow{\tau} \mathbf{0}_{\frac{1}{4}} \oplus a_0. \mathbf{0}$ . Now let  $T$  be the test  $\bar{a}_0.\omega$ . Both  $P \parallel T$  and  $Q \parallel T$  give rise to fully probabilistic wMDPs. The unique expected benefit resulted from  $P \parallel T$  is  $\frac{1}{4} \cdot 0 + \frac{3}{4} \cdot 2$ , i.e.  $\frac{3}{2}$ . On the other hand, the unique expected benefit obtained from  $Q \parallel T$  is  $\frac{1}{2} \cdot 1 + \frac{1}{2}(\frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 3)$ , i.e.  $\frac{5}{4}$ . As  $\{\frac{3}{2}\} \not\subseteq_{\text{Ho}}^0 \{\frac{5}{4}\}$ , we have that  $P$  is not related to  $Q$  under expected benefits may testing; thus  $\prec$  is unsound for expected benefits testing. Note that if we consider total benefits, then  $\mathbf{Benefits}(P \parallel T) = \{2\} = \mathbf{Benefits}(Q \parallel T)$ .  $\square$

## 5 Concluding remarks

We have proposed the model of weighted Markov decision processes for compositional reasoning about the behaviour of systems with uncertainty. Amortised weighted simulation is coinductively defined to be a behavioural preorder for comparing different wMDPs. It is shown to be a precongruence relation with respect to all structural operators for constructing wMDPs from components. For finitary convergent wMDPs, we have also given logical and testing characterisations of the simulation preorder: it can be completely determined by a qualitative probabilistic logic and for each system we can find a characteristic formula to capture its behaviour; the simulation preorder also coincides with a notion of may testing preorder.

The dual of may testing is must testing. It would be interesting to investigate the must preorder given by our testing approach. We leave it as future work to provide a coinductive formulation of the preorder and study its logical characterisations.

There is a very limited literature on compositional theories of Markov decision processes particularly in the presence of weights. There is however an extensive literature on probabilistic variations of bisimulation equivalence for Markov chains; see Chapter 10 of [1] for an elementary introduction and [10] for a survey. Bisimulation equivalence has also been defined in [8] for *Interactive Markov Chains (IMCs)*, and it is shown to be compositional, in the sense of our Theorem 4: it is preserved by the operators of a process calculus interpreted as IMCs. Bisimulation and testing equivalence for Markovian process algebras are also investigated in [9, 2], but the analysis was mainly restricted to models free of nondeterminism. Recently a combination of probabilistic automata and IMCs has been studied in [7], where a notation of weak bisimulation is proposed. Since time rates are treated essentially as action names, some intuitively equivalent processes are differentiated by the weak bisimulation.

There is also an extensive literature on weighted automata [6], and probabilistic variations have also been studied [3]. However there the focus is on traditional language theoretic issues, rather than our primary concern, compositionality.

## References

1. C. Baier and J.-P. Katoen. *Principles of Model Checking*. The MIT Press, 2008.
2. M. Bernardo and R. Cleaveland. A theory of testing for Markovian processes. In Proc. *CONCUR'00*, LNCS 1877, pp. 305-319. Springer, 2000.
3. K. Chatterjee, L. Doyen, and T.A. Henzinger. Probabilistic Weighted Automata. In Proc. *CONCUR'09*, LNCS 5710, pp. 244–258. Springer, 2009.
4. Y. Deng and M. Hennessy. Compositional reasoning for Markov decision processes (full version), 2010. Available at <http://basics.sjtu.edu.cn/~yuxin/temp/mdp.pdf>.
5. Y. Deng, R. van Glabbeek, M. Hennessy, and C. Morgan. Testing finitary probabilistic processes. In Proc. *CONCUR'09*, LNCS 5710, pp. 274–288. Springer, 2009.
6. M. Droste, W. Kuich, and H. Vogler (Eds.) *Handbook of Weighted Automata*. Springer, 2009.
7. E. Eisentraut, H. Hermanns, and L. Zhang. On probabilistic automata in continuous time. In Proc. *LICS'10*, pp. 342-351. IEEE Computer Society, 2010.
8. H. Hermanns. *Interactive Markov Chains: The Quest for Quantified Quality*, LNCS 2428. Springer, 2002.
9. J. Hillston. *A Compositional Approach to Performance Modelling*. Cambridge University Press, 1996.
10. B. Jonsson, K.G. Larsen, and Y. Wang. Probabilistic Extensions of Process Algebras. In *Handbook of Process Algebra*, pp. 685–710. Elsevier, 2001.
11. A. Kiehn and S. Arun-Kumar. Amortised bisimulations. In Proc. *FORTE'05*, LNCS 3731, pp. 320–334. Springer, 2005.
12. M. Puterman. *Markov Decision Processes*. Wiley, 1994.
13. J. Rutten, M. Kwiatkowska, G. Norman, and D. Parker. *Mathematical Techniques for Analyzing Concurrent and Probabilistic Systems*. American Mathematical Society, 2004.
14. R. Segala. Modeling and verification of randomized distributed real-time systems. Technical Report MIT/LCS/TR-676, PhD thesis, MIT, Dept. of EECS, 1995.