

Convenient categories of asynchronous processes and simulations

Duško Pavlović

COGS, University of Sussex, Brighton, England; e-mail: duskop@cogs.susx.ac.uk

Abstract

In [9] we have argued for a representation of processes taking into account computationally relevant morphisms. It has been shown that the category of synchronous processes, modulo strong bisimulation, with the bisimilarity preserving simulations, is isomorphic with a particular subcategory of transition systems with graph morphisms.

In the present paper, we extend this representation to asynchronous processes, modulo the weak and the branching bisimulations and congruences. They are shown to correspond to further interesting subcategories of the category of transition systems. The form of the representatives in the case of the branching bisimilarity suggests possible connections with game theory. An abstract construction of a category of processes in a general setting is presented in the appendix.

1 Introduction

As a first approximation, processes are presented as directed graphs of states and transitions: a computation is a directed path of transitions, a *run* from state to state, starting from the initial one. The transitions are labelled by the actions taken.

However, many different graphs remain indistinguishable when only such runs are observed. The computationally irrelevant properties of graphs are factored out by defining processes as classes of observationally equivalent, *bisimilar* graphs. But large classes of graphs are not very convenient to work with, and one tries to pick a canonical representative from each of them. This boils down to extracting a class of graphs that display only the computational behaviour of a process, free of geometric redundancies. In [9], we have described a category of *irredundant* transition systems, which couniversally represent synchronous processes, i.e. strong bisimilarity classes. The treatment is now extended to asynchronous processes, induced by weaker notions of bisimilarity. Such representations are necessary for logical studies of processes [10]. The new task will be reduced to the previously solved one. In each weaker bisimilarity class we shall find a strong subclass, in fact a retract, for which we already have a representative. This retraction is worked out in section 3, after the categorical framework needed for it has been introduced in section 2 (based on a universal construction from the appendix). The actual representation of asynchronous processes is in 4, while 5 outlines modifications needed for capturing processes with respect to the generated congruences.

2 From graphs to processes

We begin with the category of *reachable transition systems*, as described in [9]. The only additional feature is the distinguished label $\tau \in \Sigma$, denoting *silent* actions [7, sec. 2.3]. It has no repercussions for the category, but allows conceptual refinements, leading to richer notions of (bi)simulation and of process. While the strong bisimulations [8, 7] take all actions into account, the weak (or observational) equivalence [4, 7] discards the silent actions from the output, and only takes their power to preempt other actions into account. However, it evens out different trees of silent actions, and remains too crude for some situations. The suitable refinement is the notion of branching bisimulation [3, ch. 3], which fully respects the branching structure of a process, including its silent parts. We shall now align the three notions formally, and derive three categories of processes.

To analyse them in categories, we shall treat bisimulations as *internal full* binary relations in various categories of transition systems. An *internal* relation between transition systems P and Q consists of a binary relation on their states and a binary relation on their transitions. The former always relates the initial states. The latter only relates transitions with the same label. Finally, whenever two transitions are related, their source states must also be related, as well as their target states. An internal relation is *full* if the converse holds as well: two transitions are related if and only if they have the same label, and related sources and targets. Obviously, such a relation is completely determined by its state component. Our categorical bisimulations are thus equivalent to the original ones, which are just relations on states (given in terms of transitions).

As usually, $x \xrightarrow{a} x'$ denotes a transition, or the statement that it exists; $x_0 \xrightarrow{*} x_n$ abbreviates a silent run $x_0 \xrightarrow{\tau} x_1 \xrightarrow{\tau} x_2 \xrightarrow{\tau} \dots \xrightarrow{\tau} x_n$ of length $n \geq 0$. An empty run $x_0 \xrightarrow{*} x_0$.

Definition 2.1 *An internal relation $P \leftarrow R \rightarrow Q$ in the category of transition systems is a strong simulation if it satisfies*

$$x \xrightarrow{a} x' \wedge xRy \implies \exists y'. y \xrightarrow{a} y' \wedge x'Ry' \quad (1)$$

$$\left. \begin{array}{l} x \xrightarrow{a} x' \wedge xRy \\ \wedge y \xrightarrow{a} y' \wedge x'Ry' \end{array} \right\} \implies \left(x \xrightarrow{a} x' \right) R \left(y \xrightarrow{a} y' \right) \quad (2)$$

for all states $x, x' \in P$ and $y \in Q$. It is a weak [resp. branching] simulation if it satisfies (2) and

$$\begin{aligned} x \xrightarrow{a} x' \wedge xRy \implies & \exists uu'y'. y \xrightarrow{*} u \xrightarrow{a} u' \xrightarrow{*} y' \wedge x'Ry' \quad [\wedge xRu \wedge x'Ru'] \\ & \vee (a = \tau \wedge x'Ry) \end{aligned} \quad (3)$$

A strong (resp. weak, branching) bisimulation is a strong (weak, branching) simulation $P \leftarrow R \rightarrow Q$ such that the dual $Q \leftarrow R^o \rightarrow P$ is a strong (...) simulation too. The transition systems P and Q are strongly (...) bisimilar if there is a strong (...) bisimulation between them. The strong, weak and branching bisimilarities are respectively denoted by \sim , \approx and \cong .

In pictures, the above definitions say that each span

$$\begin{array}{c} x \\ \vdots \\ \hat{R} \\ \vdots \\ y \end{array} \xrightarrow{a} x' \quad \text{extends to} \quad \begin{array}{ccc} x & \xrightarrow{a} & x' \\ \vdots & & \vdots \\ \hat{R} & & \hat{R} \\ \vdots & & \vdots \\ y & \xrightarrow{a} & y' \end{array} \quad \text{in the strong case,} \quad (4)$$

$$\text{to} \quad \begin{array}{ccc} x & \xrightarrow{a} & x' \\ \vdots & & \vdots \\ \hat{R} & & \hat{R} \\ \vdots & & \vdots \\ y & \xrightarrow{*} u & \xrightarrow{a} u' & \xrightarrow{**} y' \end{array} \quad \text{or} \quad \begin{array}{ccc} x & \xrightarrow{a=\tau} & x' \\ \vdots & & \vdots \\ \hat{R} & & \hat{R} \\ \vdots & & \vdots \\ y & & y' \end{array} \quad \text{in the weak case,} \quad (5)$$

$$\text{and to} \quad \begin{array}{ccc} x & \xrightarrow{a} & x' \\ \vdots & & \vdots \\ \hat{R} & & \hat{R} \\ \vdots & & \vdots \\ y & \xrightarrow{*} u & \xrightarrow{a} y' \end{array} \quad \text{or} \quad \begin{array}{ccc} x & \xrightarrow{a=\tau} & x' \\ \vdots & & \vdots \\ \hat{R} & & \hat{R} \\ \vdots & & \vdots \\ y & & y' \end{array} \quad \text{in the branching case.} \quad (6)$$

Note that omitting u' in the branching case does not change anything: one can always take $u' = y'$.

These notions now induce three poset-enriched categories, with the reachable transition systems as objects and the simulations as morphisms. Restricting to the sober simulations, we get the categories C^- , C^∞ and C^\cong , to which we apply the construction from the appendix, and get the categories of processes \mathcal{P}_- , \mathcal{P}_∞ and \mathcal{P}_\cong . The components of the families \sim , \approx and \cong are in each case the largest bisimulations.

Let us summarize what the obtained categories of processes look like. For any family of arrows/equivalence relation $\psi \in \{\sim, \approx, \cong\}$, the objects of \mathcal{P}_ψ are the ψ -bisimilarity classes of reachable transition systems. Given two such classes, Π and Θ , a morphism $\Pi \leftarrow \Xi \rightarrow \Theta$ in \mathcal{P}_ψ will be a class

$$\Xi = \{P \leftarrow R \rightarrow Q \mid P \in \Pi, Q \in \Theta\},$$

of ψ -simulations, such that for any $P \leftarrow R \rightarrow Q$ and $P' \leftarrow R' \rightarrow Q'$ holds

$$x\psi x' \wedge xRy \wedge y\psi y' \implies x'R'y' \quad (7)$$

$$x\psi x' \wedge xRy \wedge x'R'y' \implies y\psi y' \quad (8)$$

for all $x \in P, y \in Q, x' \in P', y' \in Q'$. Respectively, these conditions say that Ξ is *saturated* and *sober*. The latter says that the components of Ξ jointly preserve the ψ -bisimilarity, i.e. take the computationally equivalent states to the computationally equivalent ones. The saturation condition, on the other hand, says that if x and y are related, then everything equivalent to x must be related to everything equivalent to y . The formal consequences of these conditions are explained in the appendix. See also section 2.3 of the first part.

The obvious implications (4) \Rightarrow (6) \Rightarrow (5), induce the quotient functors $\mathcal{P}_- \rightarrow \mathcal{P}_\cong \rightarrow \mathcal{P}_\infty$. Note however, that the morphisms will fit only if the implications \sim -sober \Rightarrow \cong -sober \Rightarrow \approx -sober are valid as well. This is not immediate, but it will follow from proposition 3.2.

3 Relating simulations

In [9], we have described a subcategory \mathcal{I} of *irredundant* transition systems, and shown that it is strongly equivalent with \mathcal{P}_\sim . In fact, its skeleton is even isomorphic with \mathcal{P}_\sim . By definition, an irredundant transition system must be reachable, and such that $x \sim x'$ implies $x = x'$ for any pair x, x' of states. This irredundant representation, known in many forms, is actually universal in a formal sense [9, sec. 5].

This picture of \mathcal{P}_\sim will now be used for representing \mathcal{P}_\approx and \mathcal{P}_\cong . A transition system P will be transformed into transition systems WP and BP , weakly resp. branching bisimilar to P , and such that the weak resp. the branching simulations to and from P exactly correspond to the strong ones on WP resp. BP . The idea how for such WP and BP follows from (4), (5) and (6). To reduce (5) to (4), we must add in WP a transition $x \xrightarrow{a} x'$ whenever a path $x \xrightarrow{*} v \xrightarrow{a} v' \xrightarrow{*} x'$ occurs in P ; and a transition $x \xrightarrow{\tau} x$ for every state x . The obtained transition system WP is thus the closure of P under the “composition” with τ -transitions.

The construction of BP is bound to be more complicated, since it must expand the trapezoid from (6) into *two* squares (4). The idea is that the transition $x \xrightarrow{c} x'$, $c \neq \tau$, should be expanded in two new transitions, corresponding to $y \xrightarrow{*} u$, and $u \xrightarrow{c} y'$ respectively. Moreover, τ should be “closed under composition” with itself. The construction B itself will thus be the composite of the constructions C and D , where

- C replaces each $x \xrightarrow{c \neq \tau} x'$ with $x \xrightarrow{\tau} \left(\begin{smallmatrix} x \\ c \\ x' \end{smallmatrix} \right) \xrightarrow{c} x'$, where $\left(\begin{smallmatrix} x \\ c \\ x' \end{smallmatrix} \right)$ is a new state; while
- D adds $x \xrightarrow{\tau} x'$ whenever there is a path $x \xrightarrow{*} x'$.

All the described constructions induce endofunctors on the category of reachable transition systems. W and D are moreover idempotent monads, while C extends a comonad G , which for all $x \xrightarrow{c \neq \tau} x'$ adds $\left(\begin{smallmatrix} x \\ c \\ x' \end{smallmatrix} \right) \xrightarrow{c} x'$ but not $x \xrightarrow{\tau} \left(\begin{smallmatrix} x \\ c \\ x' \end{smallmatrix} \right)$. The units $\eta : P \rightarrow WP$ and $\eta : P \rightarrow DP$ are given by the identity maps on the states, and the inclusions on the transitions. The counit $\varepsilon : GP \rightarrow P$ maps both x and $\left(\begin{smallmatrix} x \\ c \\ x' \end{smallmatrix} \right)$ to x , the transitions $\left(\begin{smallmatrix} x \\ c \\ x' \end{smallmatrix} \right) \xrightarrow{c} x'$ to $x \xrightarrow{c} x'$, and each τ -transition to itself. These data now provide a weak bisimulation $P \xleftarrow{\text{id}} P \xrightarrow{\eta} WP$ and a branching bisimulation $P \xleftarrow{\varepsilon} GP \hookrightarrow CP \xrightarrow{\eta} DCP = BP$. This is proved simply by the inspection of definitions.

Proposition 3.1 *For every reachable transition system P , the equality relation on the states yields a weak bisimulation between P and WP . A branching bisimulation between P and BP is obtained by extending the equality on the states of P with the pairs of the form $\left\langle x, \left(\begin{smallmatrix} x \\ c \\ x' \end{smallmatrix} \right) \right\rangle$.*

Since a simulation is determined by its state component, and W does not change the states, the W -image of a full relation $P \leftarrow R \rightarrow Q$ can be defined as the full relation $WP \leftarrow WR \rightarrow WQ$ induced by the state component of the original. The B -image, on the other hand, will be the full relation $BP \leftarrow BR \rightarrow BQ$ spanned by the state component of $P \leftarrow R \rightarrow Q$ extended on the new states by

$$\left(\begin{smallmatrix} x \\ c \\ x' \end{smallmatrix} \right) BR \left(\begin{smallmatrix} y \\ d \\ y' \end{smallmatrix} \right) \iff c = d \wedge xRy \wedge vRu. \quad (9)$$

In this way, we get enriched functors $W : C^\approx \rightarrow C^\sim$ and $B : C^\approx \rightarrow C^\sim$, which turn out to be full and faithful, as a consequence of the following proposition.

Proposition 3.2 *A full relation $P \leftarrow R \rightarrow Q$ on reachable transition systems is a weak simulation if and only if $WP \leftarrow WR \rightarrow WQ$ is a strong simulation. It is a branching simulation if and only if $BP \leftarrow BR \rightarrow BQ$ is a strong simulation.*

Proof. We only prove the second statement, since the weak case is straightforward.

(\Rightarrow) Assuming that $P \leftarrow R \rightarrow Q$ satisfies (6) we derive that $BP \leftarrow BR \rightarrow BQ$ satisfies (4). In BP , there are clearly three kinds of transitions to be simulated:

- (i) $x \xrightarrow{\tau} x'$, or
- (ii) $x' \xrightarrow{\tau} \begin{pmatrix} x \\ c \\ v \end{pmatrix}$, or
- (iii) $\begin{pmatrix} x \\ c \\ v \end{pmatrix} \xrightarrow{c} v$,

where x , x' and v are old states, coming from P , while $c \neq \tau$.

To discuss (i), suppose xRy . Since R is a branching simulation, there is $y \xrightarrow{*} y'$ in Q , with $x'Ry'$. Hence $y \xrightarrow{\tau} y'$ in BQ .

Towards case (ii), note that in CP , the state $\begin{pmatrix} x \\ c \\ v \end{pmatrix}$ can only be reached through x . The transition $x' \xrightarrow{\tau} \begin{pmatrix} x \\ c \\ v \end{pmatrix}$ thus comes in BP from a path $x' \xrightarrow{*} x \xrightarrow{\tau} \begin{pmatrix} x \\ c \\ v \end{pmatrix}$ in CP . On the other hand, $x \xrightarrow{\tau} \begin{pmatrix} x \\ c \\ v \end{pmatrix}$ comes from $x \xrightarrow{c} v$ in P . So there must have been $x' \xrightarrow{*} x \xrightarrow{c} v$ in P . Given $x'Ry'$, the assumption that R is a branching simulation yields $y' \xrightarrow{*} y \xrightarrow{*} u \xrightarrow{c} w$ in Q , with xRu , xRu and vRw . But the transition $u \xrightarrow{c} w$ from Q becomes $u \xrightarrow{\tau} \begin{pmatrix} u \\ c \\ w \end{pmatrix} \xrightarrow{c} w$ in CQ . The path $y' \xrightarrow{*} y \xrightarrow{*} u \xrightarrow{\tau} \begin{pmatrix} u \\ c \\ w \end{pmatrix}$ in CQ now induces $y' \xrightarrow{\tau} \begin{pmatrix} u \\ c \\ w \end{pmatrix}$ in BQ . This transition simulates $x' \xrightarrow{\tau} \begin{pmatrix} x \\ c \\ v \end{pmatrix}$, since by (9), xRu and vRw imply $\begin{pmatrix} x \\ c \\ v \end{pmatrix} BR \begin{pmatrix} u \\ c \\ w \end{pmatrix}$.

Finally, for case (iii), suppose $\begin{pmatrix} x \\ c \\ v \end{pmatrix} BR z$. By the definition of BR again, the state z must be in the form $\begin{pmatrix} u \\ c \\ w \end{pmatrix}$, for some u, w with xRu and vRw . The transition $\begin{pmatrix} x \\ c \\ v \end{pmatrix} \xrightarrow{c} v$ is thus simulated by $\begin{pmatrix} u \\ c \\ w \end{pmatrix} \xrightarrow{c} w$.

(\Leftarrow) Now assume that $BP \leftarrow BR \rightarrow BQ$ is a strong simulation, and take $x \xrightarrow{a} x'$ in P , with xRy . If $a = \tau$, (4) gives $y \xrightarrow{\tau} y'$ in BQ , with $x'Ry'$. Hence $y \xrightarrow{*} y'$ in Q . If this is an empty path, and $y = y'$, we have the triangle from (6); otherwise we have the trapezoid.

If $a = c \neq \tau$, the transition $x \xrightarrow{c} x'$ becomes $x \xrightarrow{\tau} \begin{pmatrix} x \\ c \\ v \end{pmatrix} \xrightarrow{c} x'$ in BP . Since BR is a strong simulation, this is simulated by $y \xrightarrow{\tau} z \xrightarrow{c} y'$ in BQ , with $\begin{pmatrix} x \\ c \\ v \end{pmatrix} BR z$ and $x'BRy'$. By the definition of BR , y' must be a state from Q , such that $x'Ry'$, while z must be in the form $\begin{pmatrix} u \\ c \\ w \end{pmatrix}$, for some Q -state u with xRu . As pointed out before, the transition $x' \xrightarrow{\tau} \begin{pmatrix} u \\ c \\ w \end{pmatrix}$ in BQ must have originated from $x' \xrightarrow{*} u \xrightarrow{\tau} \begin{pmatrix} u \\ c \\ w \end{pmatrix}$ in CQ . In Q , there is thus $x' \xrightarrow{*} u \xrightarrow{c} y'$, with xRu and $x'Ry'$, just as required by (6). This completes the proof. \square

The construction W induces a poset isomorphism between the weak simulations from P to Q and the strong simulations from WP to WQ , because the state component of $WP \leftarrow WR \rightarrow WQ$ is the same as the state component of $P \leftarrow R \rightarrow Q$ again. Therefore, W is a full and faithful enriched functor from the category of weak simulations to the category of strong simulations. Since two states in P are weakly bisimilar if and only if they are strongly bisimilar in WP , the functor W preserves and reflects the sobriety, and thus restricts to a full and faithful enriched functor $W : C^\approx \rightarrow C^\sim$.

A similar reasoning leads to the same conclusion for $B : C^\cong \rightarrow C^\sim$. The fullness may seem not as obvious this time, though. But note that the sources of the transitions labelled by $c \neq \tau$ in BP are always the new states $(c \downarrow_x)$; and that all of them appear as such sources. Hence, a strong simulation from BP to BQ must relate the new states among themselves. It must further satisfy (9) — and thus appear in the form BR for some $P \leftarrow R \rightarrow Q$.

4 The representation

The described functors on the categories of simulations now induce the full and faithful functors $W : \mathcal{P}_\approx \rightarrow \mathcal{P}_\sim$ and $B : \mathcal{P}_\cong \rightarrow \mathcal{P}_\sim$. Proposition 3.1 implies that they are right inverse, respectively, to the quotient functors $\mathcal{P}_\sim \rightarrow \mathcal{P}_\approx$ and $\mathcal{P}_\sim \rightarrow \mathcal{P}_\cong$. Although the construction B on transition systems was not idempotent, the endofunctor that it induces on \mathcal{P}_\cong is. \mathcal{P}_\cong and \mathcal{P}_\approx are thus retracts of \mathcal{P}_\sim ; the latter is even a reflective subcategory. They are thus also retracts of the category of reachable transition systems and sober morphisms, since \mathcal{P}_\sim can be viewed as its reflective category (cf. sec. 5 of the first part [9, sec. 5]). This is based on presenting \mathcal{P}_\sim as the skeleton of the category \mathcal{I} of irredundant transition systems [9, thm. 4.4].

Our next task is characterising the retracts of \mathcal{I} corresponding to \mathcal{P}_\approx and \mathcal{P}_\cong . They are spanned by the images of the constructions W and B on \mathcal{I} . The scheme of the representation is:

$$\begin{array}{ccc}
 P \longleftarrow R \longrightarrow Q & & P \longleftarrow R \longrightarrow Q \\
 \uparrow & & \uparrow \\
 \approx & & \cong \\
 \downarrow & & \downarrow \\
 WP \longleftarrow WR \longrightarrow WQ & & BP \longleftarrow BR \longrightarrow BQ \\
 \uparrow & & \uparrow \\
 \sim & & \sim \\
 \widetilde{WP} = \widetilde{WR} \rightarrow \widetilde{WQ} & & \widetilde{BP} = \widetilde{BR} \rightarrow \widetilde{BQ}
 \end{array} \tag{10}$$

The upper squares depict the contents of the previous section: the vertical bisimulations are those from proposition 3.1, while 3.2 relates the horizontal sober simulations. The lower squares are from [9, sec. 5]. The functor $\widetilde{(-)}$ assigns to each reachable transition system its irredundant quotient, and to each sober \sim -simulation a graph morphism. Although \mathcal{I} is not closed under W or B , direct inspection of the definitions of W and B shows that their images are closed under $\widetilde{(-)}$. The representatives \widetilde{WP} and \widetilde{BP} will thus be in the following forms.

Definition 4.1 *A transition system is τ -replete if whenever there is a path $x \xrightarrow{*} y$, there is also*

a transition $x \xrightarrow{\tau} y$ in it. It is c -replete, for $c \neq \tau$ if every path $x \xrightarrow{*} u \xrightarrow{c} v \xrightarrow{*} y$ in it can be “shortcut” by a transition $x \xrightarrow{c} y$. The replete transition systems are a -replete for all $a \in \Sigma$. The irredundant replete transition systems span the subcategory \mathcal{I}_{\approx} of \mathcal{I} .

A τ -strategy is a transition system in which the source of a visible transition is always the target of exactly one transition, necessarily silent. In other words, if $y \xrightarrow{c \neq \tau} z$, then there is $x \xrightarrow{\tau} y$, and no other transitions to y . Irredundant, τ -replete τ -strategies span the subcategory \mathcal{I}_{\cong} of \mathcal{I} .

Theorem 4.2 $\mathcal{P}_{\psi} \cong \mathcal{I}_{\psi}$, where $\psi \in \{\approx, \cong\}$.

Remarks. In a way, the repleteness embodies asynchrony: any action can be delayed by idling, and any action that can be taken after some idling can also be taken immediately. However, such “shortcuts” may erase a part of the branching structure. In a τ -strategy, this is prevented by separating visible actions from each other by silent actions. The idea behind its name is that each visible move is given a unique silent response. The game is played *towards* the initial state, which can be thought of as the final position. The side which plays a last move, wins. The silent side wins using the τ -strategies.

Regardless of the relevance of this game theoretic picture, the point of the τ -replete τ -strategies is to allow waiting to be extended at will, or reduced to one silent step — but not completely eliminated. The branching structure is protected from these deformations by keeping the visible and the silent layers of actions separated.

5 Congruences

Finally, let us turn to a conceptual shortcoming of the weak and the branching bisimilarities: they are not congruences with respect to all of the process operations. In particular, the nondeterministic sum $+$ does not preserve them [7]. Fortunately, the reason for this turns out to be localized at the initial state. Bergstra and Klop [2] have observed that *rooted* bisimulations — where the initial states are *only* related with each other — are always preserved under $+$. The following lemma is probably folklore.

Lemma 5.1 A (strong, weak, branching) simulation $P \leftarrow R \rightarrow Q$ is rooted if and only if the relations $P + M \leftarrow R + M \rightarrow Q + M$ are (...) simulations for all transition systems M .

Although simple and elegant, the root condition is not well suited for categorical treatment. However, there is a weaker condition, due to Milner [7, ch. 7, def. 2], which yields the same processes as the rooted bisimulations, and is easily captured in categories. It requires that a transition in the form $\iota \xrightarrow{\tau} x'$ is never simulated by an empty path. The simulations satisfying this requirement may not be rooted, hence they are not stable under sums; but *if* there is a bisimulation between P and Q satisfying Milner’s requirement, *then* a rooted one must exist as well. Roughly, the *only* reason why it may be impossible to simply drop a pair $\langle x, \iota \rangle$ (or $\langle \iota, y \rangle$) from a bisimulation is that the only transition simulating $\iota \xrightarrow{\tau} x$ (or $\iota \xrightarrow{\tau} y$) may be

$\iota \xrightarrow{\tau} \iota$. The following definition thus yields processes as the rooted bisimulations. Similarly, the induced process morphisms do not consist of rooted simulations, but each of them contains rooted components (e.g., the tree morphisms), and this suffices for +-stability.

Definition 5.2 A relation $P \leftarrow R \rightarrow Q$ in \mathcal{A}_Σ is a ι -weak [ι -branching] simulation if it satisfies

$$x \xrightarrow{a} x' \wedge xRy \implies \exists uu'y'. y \xrightarrow{*} u \xrightarrow{a} u' \xrightarrow{*} y' \wedge x'Ry' [\wedge xRu \wedge x'Ru'] \vee \left(\underline{x \neq \iota} \wedge a = \tau \wedge x'Ry \right) \quad (11)$$

and (2). A weak [branching] congruence is a ι -weak [ι -branching] simulation, the dual of which is ι -weak [ι -branching] simulation as well.

The ι -weak and ι -branching simulations differ from the weak and the branching ones only by the underlined part of (11); the rest is exactly like (3). The ι -simulations can thus be analyzed along the same lines as the ordinary ones — just slightly modifying the constructions W and B . Namely, everything remains the same, but *no* τ -cycles $x \xrightarrow{\tau} x$ must be added at $x = \iota$. On the representatives, this exception is expressed simply by restricting the τ -repleteness at the root.

6 Future work

The main point of our approach is to capture processes dynamically, in a category, with computation preserving morphisms. Concurrency may have sailed well without such morphisms, but the forest of its operations looks more and more like tensor calculus in the time when it was based on bright physical intuitions, but the universal property of the tensor product had not yet been understood.

1. P. Aczel, *Non-Well-Founded Sets*, Lecture Notes 14 (CSLI 1988)
2. J.A. Bergstra and J.W. Klop, Algebra of communicating processes with abstraction, *Theoret. Comput. Sci.* 37(1985) 77–121
3. R.J. van Glabbeek, *Comparative Concurrency Semantics and Refinement of Actions*, thesis (CWI, 1990)
4. M. Hennessy and R. Milner, On observing nondeterminism and concurrency, in: J. de Bakker and J. van Leeuwen, eds., *Proceedings of ICALP 80*, Lecture Notes in Computer Science 85 (Springer 1980) 299–309
5. A. Joyal, M. Nielsen and G. Winskel, Bisimulation and open maps, *Proceedings of the Eight Symposium on Logic in Computer Science* (IEEE 1993) 418–427
6. R. Milner, *A Calculus of Communicating Systems*, Lecture Notes in Computer Science 92 (Springer 1980)
7. R. Milner, *Communication and Concurrency*, Internat. Ser. in Comp. Sci. (Prentice Hall, 1989)
8. D. Park, *Concurrency and Automata on Infinite Sequences*, Lecture Notes in Computer Science 104 (Springer 1980)

9. D. Pavlovic, Convenient category of synchronous processes and simulations, *in*: D. Pitt et al. (eds.), *Category Theory in Computer Science '95*, Lecture Notes in Computer Science 953 (Springer 1995) 3–24
10. D. Pavlovic, Categorical logic of concurrency and interaction I: Synchronous processes, *in*: C. Hankin et al. (eds.), *Proceedings of the Second TFM Workshop* (World Press 1995) 105–141

Appendix A: Quotient of a poset-enriched category

Let C be a poset-enriched category and Ψ a composition closed family of arrows in it, at most one from each hom-poset. Write $P\psi Q$ for the Ψ -arrow from P to Q . It is further required that Ψ contains for each $P \in C$ an endomorphism $P\psi P$, greater than id_P ; and also an arrow $Q\psi P$ whenever there is $P\psi Q$ in it. The family Ψ can thus be construed as an equivalence relation ψ on the objects of C , *realized* by its arrows.

We want to form a quotient category $\mathcal{P} = C/\Psi$. The objects should be the ψ -equivalence classes of objects from C , the morphisms — the families of ψ -equivalent arrows between their elements. Here the equivalence boils down to the requirement that all components of such a family of arrows can be obtained by extending any of them along ψ , just like any element of an equivalence class determines all of it. This seems to be a precondition of the existence of the quotient map, a functor $C \rightarrow \mathcal{P}$. But it is not hard to satisfy.

Let Π, Θ be some ψ -equivalence classes of objects. A morphism $\Pi \rightarrow \Theta$ in \mathcal{P} will now be a class Ξ of arrows $P \rightarrow Q$ in C , one for each pair $P \in \Pi, Q \in \Theta$, such that for any $R, R' \in \Xi$ holds $\psi R \psi \subseteq R'$, or diagrammatically

$$\begin{array}{ccc}
 P & \leftarrow \psi & P' \\
 | & & | \\
 R & \subseteq & R' \\
 \downarrow & & \downarrow \\
 Q & \leftarrow \psi & Q'
 \end{array} \tag{12}$$

This is the *saturation* condition. Using the above assumptions about Ψ , we get $R' \subseteq \psi R' \psi = \psi \psi R' \psi \subseteq \psi R \psi$, so that for all components any \mathcal{P} -morphism actually holds

$$R' = \psi R \psi. \tag{13}$$

In particular, each of them is *saturated*, i.e. $R = \psi R \psi$. At any rate, each component pins down a \mathcal{P} -morphism by formula (13).

The quotient map $C \rightarrow \mathcal{P}$ will, of course, take each object P to its equivalence class Π , and each arrow $P \rightarrow Q$ to the family of arrows $P' \rightarrow Q'$, one for each $P' \in \Pi$ and $Q' \in \Theta$, which are obtained as in (13). Clearly, this family will consist of *saturated* morphisms, satisfying $\bar{R} = \psi \bar{R} \psi$. In fact, the image of each morphism R along $C \rightarrow \mathcal{P}$ is determined by the *saturation* $\bar{R} = \psi R \psi$. Clearly, saturation can be viewed as a functor $C \rightarrow \bar{C}$, where \bar{C} is

the subcategory of C , consisting of the same objects but only the saturated morphisms. \overline{C} is the *saturation* of C .

To spell out the universal properties of the quotient \mathcal{P} and the saturation C , we shall say that a functor *annihilates ψ on objects* if it takes any two ψ -related objects to the same image; that it *annihilates ψ on arrows* if each R has the same image as its saturation $\overline{R} = \psi R \psi$; and that it *annihilates ψ* if it does so both on objects and arrows.

Proposition .1 *The quotient $C \rightarrow \mathcal{P}$ is initial among those that annihilate ψ . The saturation $C \rightarrow \overline{C}$ is initial among those that annihilate ψ on arrows. The induced functor $\overline{C} \rightarrow \mathcal{P}$ is full and faithful, and surjective on objects — hence a weak equivalence.*

Appendix B: Categories of processes, abstractly

Now we want to focus on situations when C is a category of simulations, while Ψ consists of maximal bisimulations. The quotient construction should thus yield the corresponding category of processes.

In order to be able to express some special properties of simulations, one additional operation will be needed, namely the dualising

$$(P -R \succ Q) \quad \mapsto \quad (Q -R^o \succ P). \quad (14)$$

A relation can always be dualized, but the dual of a simulation need not be a simulation. The category C will thus not be closed under dualizing, but we shall assume that it is couched a “category of relations”, providing the needed operation. For instance, C can be viewed as a poset-enriched subcategory of an allegory or of a cartesian bicategory.

There are two additional requirements that need to be imposed on the arrows of C . Firstly, each of them should be *total*, or formally:

$$\text{id} \subseteq RR^o. \quad (15)$$

Furthermore, they should all preserve the bisimilarity: indeed, a sound computational morphism should preserve computational equivalence [9, sec. 2.3]. Formally, this means that all $R \in C$ should satisfy the *sobriety* condition

$$R^o \psi R \subseteq \psi. \quad (16)$$

But this is easily enforced: in specific cases, we simply restrict consideration to the *sober* simulations.

A poset-enriched category C with total and sober morphisms can be construed as an *abstract category of simulations*. Remarkably, the enrichment of such a category degenerates under saturation: \overline{C} turns out to be an *ordinary* category.

Lemma .2 *If $P \dashv R \rhd Q$ is a sober morphism, its saturation $\overline{R} = \psi R \psi$ is sober too. If, furthermore, all morphisms from P to Q are total, the saturation \overline{R} is the largest among them which is both sober and contains R .*

Proof. \overline{R} is sober because $\overline{R} \psi \overline{R} = \psi R^o \psi \psi R \psi = \psi R^o \psi R \psi \subseteq \psi \psi R \psi = \psi$. For the second statement, suppose that $\hat{R} \supseteq R$ is sober. The totality now yields $\hat{R} \subseteq RR^o \hat{R} \subseteq R\hat{R}^o \hat{R} \subseteq R\hat{R}^o \psi \hat{R} \subseteq R\psi \subseteq \psi R \psi$. \square

Corollary .3 *The saturation \overline{C} of an abstract category of simulations C is an ordinary category.*

Proof. If $R \subseteq \hat{R}$, and \hat{R} is sober, lemma yields $\hat{R} \subseteq \overline{R}$. When R is saturated, this yields $\hat{R} \subseteq R$. \square

The quotient \mathcal{P} , constructed from a category of simulations by the method from appendix 6, is an *abstract category of processes*. The universal properties described in proposition .1 remain valid. The additional properties make \mathcal{P} into an ordinary category, though, since it is weakly equivalent to \overline{C} . Another point worth emphasizing about \mathcal{P} is that the bisimilarity is preserved not just by the individual components of its morphisms, but also jointly, which renders an apparently stronger, global notion of sobriety.

Proposition .4 *If each component of a \mathcal{P} -morphism $\Pi \dashv \Xi \rhd \Theta$ is sober, then any two of them $R, R' \in \Xi$ satisfy $R^o \psi R' \subseteq \psi$.*

$$\begin{array}{ccc}
 P & \leftarrow R^o & \longrightarrow Q \\
 | & & | \\
 \psi & \subseteq & \psi \\
 \Psi & & \Psi \\
 P' & \longrightarrow R' & \longrightarrow Q'
 \end{array} \tag{17}$$

Proof. $R^o \psi R' = R^o \psi \psi R \psi = R^o \psi R \psi \subseteq \psi \psi = \psi$ \square

Condition (17) is the naturality requirement more suitable for a relational setting than the ordinary one. When the relations involved in it are total ($\text{id} \subseteq RR^o$) and single-valued ($R^o R \subseteq \text{id}$), this condition is equivalent with the usual lax naturality $\psi R' \subseteq R \psi$. In general, though, the two are incomparable, and (17) captures the ψ -preservation. Conditions (12) and (17), fulfilled in each category of processes, were explained from a computational point of view in [9, sec. 2.3]. A morphism satisfying them actually induces a partial map between the bisimilarity classes of states at its domain and codomain. Since all simulations are total by assumption (15), this map yields an honest morphism of transition systems.