Cellular Monads from Positive GSOS Specifications

Tom Hirschowitz*

Univ. Grenoble Alpes, Univ. Savoie Mont Blanc, CNRS, LAMA
73000 Chambéry, France

`tom.hirschowitz@univ-smb.fr`

We give a leisurely introduction to our abstract framework for operational semantics based on cellular monads on transition categories. Furthermore, we relate it for the first time to an existing format, by showing that all Positive GSOS specifications generate cellular monads whose free algebras are all compositional. As a consequence, we recover the known result that bisimilarity is a congruence in the generated labelled transition system.

1 Introduction

1.1 Motivation

In the vast majority of foundational research on programming languages, although ideas are thought of as widely applicable, they are presented on *one*, simple example. Typically, there is a tension between simplicity of exposition, leading to the minimal language making the idea relevant, and significance, leading to the most expressive one. Strikingly, the scope of the idea is often mostly clear to the experts, but no attempt is made at stating it precisely. The reason for this is that the mathematical concepts needed for even only making such statements are lacking. Indeed, one needs to be able to say something like: "for all programming languages of such shape, the following holds". But there simply is no widely accepted mathematical notion of programming language.

Such a general notion should account for both

(i) the interaction between syntax and dynamics, as involved in, e.g., structural operational semantics [21], or in the statement of results like type soundness, congruence of program equivalence, or compiler correctness, and

(ii) denotational semantics, in the sense of including not only operational, syntactic models but also others, typically ones in which program equivalence is coarser.

Typically, standard formats [19] elude denotational semantics, and are exclusively syntactic. To our knowledge, the only such proposals meeting all these criteria are *functorial operational semantics*, a.k.a. *bialgebraic semantics* [25], and a few variants [4, 24]. This approach has been deeply developed, and shown to extend smoothly to various settings, e.g., non-deterministic and probabilistic languages. However, two important extensions have proved more difficult.

- The treatment of languages with variable binding is significantly more technical than the basic setting [9, 8, 24].

- More importantly, the bialgebraic study of higher-order languages like the λ-calculus or the higher-order π-calculus is only in its infancy [20].

This leaves some room for exploring potential alternatives.

---

*Thanks to Jorge Pérez and Jurriaan Rot for the invitation, and to the referees for helpful comments.

1.2   Context

In recent work [12], a new approach to abstract operational semantics was proposed, and its expressive power was demonstrated by proving for the first time an abstract soundness result for *bisimulation up to context* in the presence of variable binding. Bisimulation up to context is an efficient technique [23, Chapter 6] for proving program equivalences, which had previously been proved correct in the bialgebraic setting [2], but only without binding.

Its novelty mainly resides in the following two technical features.

Transition categories   First, while standard operational semantics is based on *labelled transition systems*, this is both generalised and abstracted over in the framework.

Generalisation   Indeed, in the examples, instead of standard labelled transition systems, we use a slight generalisation similar in spirit to [6], essentially from relations to graphs, i.e., possibly with several transitions between two states. This simple, harmless generalisation brings in a lot of useful structure, typically that of a *topos* [17], which is unavailable at this level in bialgebraic operational semantics.

Abstraction   In full generality, the framework takes as a parameter a *transition category*, a typical example of which is given by such generalised transition systems. For any object of a given transition category, bisimulation may be defined by lifting, following an idea from [14].

Combinatorial category theory   A second technical innovation is the use of advanced combinatorial category theory. To start with, *familial monads* [3], or rather their recent *cellular* variant [10], provide a notion of evaluation context for both programs and transitions, at the abstract level. Standard reasoning by induction on context thus becomes simple algebraic calculation. A second, crucial notion is cofibrantly generated factorisation systems, a notion from homotopy theory [13, 22] which, together with cellularity, allows for a conceptually simple, yet relevant characterisation of well-behaved transition contexts.

Each instance of the framework is then constructed as follows.

Type of transition system   The first step is the choice of a type of transition system, which may involve different kinds of states (e.g., initial or final ones), the set of labels to be put on transitions, etc. Technically, this amounts to fixing a transition category $\mathscr{C}$. This also fixes the relevant notion of bisimulation, hence bisimilarity.

Transition rules   The second step consists in defining the dynamics of the considered language, which is usually specified through a set of inference rules. This comes in as a monad T on $\mathscr{C}$, whose algebras are essentially the transition systems satisfying the given inference rules. The standard, syntactic transition system is typically the free algebra $T(0)$. This fixes the relevant notion of context closure. In this setting, congruence of bisimilarity $\sim_X$ on a T-algebra X is the fact that $T(\sim_X) \to X^2$ factors through $\sim_X \to X^2$ (see (1) on page 7).

One of the main results [12, Corollary 4.30] is that if the considered algebra is *compositional*, in the sense that its structure map $T(X) \to X$ is a functional bisimulation, and if the monad T satisfies an additional condition, then bisimilarity is indeed a congruence. The latter condition is called $\mathbf{T}_s$-*familiality* in [12], but we will here call it *cellularity*, because it is a specialisation of cellularity in the sense of [10] to familial functors. As mentioned above, a second main result [12, Corollary 5.15] is that under a different condition called $\mathbf{T}_s^\vee$-*familiality*, bisimulation up to context is sound.

1.3 Contribution

One of the main issues with cellular monads T on transition categories $\mathscr{C}$ is the lack of an efficient generation mechanism, i.e., a mathematical construction that produces pairs $(\mathscr{C}, T)$ from more basic data. In this paper, we initiate the search for such generating constructions by showing that an existing simple format, *Positive GSOS* [1], always produces cellular monads whose free algebras are compositional. As a consequence, we recover (Theorem 10.4) the known result that bisimilarity is a congruence in all free algebras.

As this is an invited contribution, we briefly introduce the approach at an expository, rather concrete level. In particular, the only considered transition category is the one of generalised labelled transition systems in the sense alluded to above. Finally, our proofs are meant to be instructive rather than fully detailed.

1.4 Plan

In §2, we explain our generalisation of labelled transition systems, and bisimulation by lifting. In §3, we recall Positive GSOS specifications Σ and show how they generate monads T_Σ. In §4, we argue that algebras for the obtained monad T_Σ are a good notion of model for the considered Positive GSOS specification. We then state congruence of bisimilarity in categorical terms, and quickly reduce it to two key properties: (i) compositionality of the considered algebra and (ii) preservation of functional bisimulations by T_Σ.

We deal with (i) in §5, where we show that when T_Σ is obtained from a Positive GSOS specification, all free algebras are compositional. In §6, we then attack (ii), by further reducing it to familiality and cellularity. The remaining sections develop these ideas.

In §7, we define familiality for functors (as opposed to monads), and show that T_Σ is a familial functor. In §8, we establish some factorisation properties of familial functors which were announced and used in §4 to reduce congruence of bisimilarity to compositionality and preservation of bisimulation. We then introduce cellularity in §9, and show that T_Σ is indeed cellular. Finally, we wrap up in §10 by defining familiality for monads (which is slightly more demanding than for mere functors), and showing that T_Σ does form a familial monad. This fills a hole left open in §5, thus allowing us to state the main theorem.

Finally, we conclude and give some perspective in §11.

1.5 Prerequisites

We assume familiarity with basic category theory [16, 15], including categories, functors, natural transformations, monads and their algebras, and the Yoneda lemma.

2 Labelled transition systems as presheaves

2.1 Generalised transition systems

A standard SOS specification is given by a signature, plus a family of transition rules over a fixed set $\mathbb{A}$ of labels. The set $\mathbb{A}$ fixes the relevant kind of transition system, and we interpret this by constructing a corresponding category of (generalised) transition systems. Given any set $\mathbb{A}$, let $\Gamma_\mathbb{A}$ denote the graph with

- vertex set $\mathbb{A} + 1$, i.e., vertices are elements of $\mathbb{A}$, denoted by $[a]$ for $a \in \mathbb{A}$, plus a special vertex $\star$,

- two edges $s^a, t^a : \star \to [a]$, for all $a \in \mathbb{A}$.

Pictorially, $\Gamma_{\mathbb{A}}$ looks like this:

$$\ldots \qquad \begin{array}{c} [a] \\ {}_{s^a}\Big(\ \ \Big){}_{t^a} \\ \star. \end{array} \qquad \ldots \qquad\qquad (a \in \mathbb{A})$$

There are no composable edges in Γ_A, so, adding formal identity arrows, it readily forms a category, which we also denote by Γ_A.

Definition 2.1. The *category of transition systems induced by* $\mathbb{A}$ is $\widehat{\Gamma_{\mathbb{A}}}$, the category of presheaves over $\Gamma_{\mathbb{A}}$.

To see what presheaves over $\Gamma_{\mathbb{A}}$ have to do with transition systems, let us observe that a presheaf $X \in \widehat{\Gamma_{\mathbb{A}}}$ consists of a set $X(\star)$ of *states*, together with, for each $a \in \mathbb{A}$, a set of *transitions* $e \in X[a]$ with source and target maps $X(s^a), X(t^a) : X[a] \to X(\star)$. Our notion is thus only slightly more general than standard labelled transition systems over $\mathbb{A}$, in that it allows several transitions with the same label between two given states.

Remark 2.2. The category $\widehat{\Gamma_{\mathbb{A}}}$ may be viewed as a category of labelled graphs. Indeed, letting $\Omega_{\mathbb{A}}$ denote the one-vertex graph with $\mathbb{A}$ loops on it, we have by well-known abstract nonsense an equivalence $\mathbf{Gph}/\Omega_{\mathbb{A}} \simeq \widehat{\Gamma_{\mathbb{A}}}$ of categories. (This is due to the fact that $\Gamma_{\mathbb{A}}$ is isomorphic to the *category of elements* of $\Omega_{\mathbb{A}}$, see Definition 7.1 below.)

Notation 2.3. For any $X \in \widehat{\Gamma_{\mathbb{A}}}$, we denote the action of morphisms in $\Gamma_{\mathbb{A}}$ with a dot. E.g., if $e \in X[a]$, then $e \cdot s^a \in X(\star)$ is its source. We also sometimes abbreviate s^a and t^a to just s and t.

Example 2.4. For languages like CCS [18], we let $\mathbb{A} = \mathcal{N} + \mathcal{N} + 1$ denote the disjoint union of a fixed set $\mathcal{N}$ of *channel names* with itself and the singleton 1. Elements of the first term are denoted by $\overline{a}$, for $a \in \mathcal{N}$, and are used for output transitions, while elements of the second term, simply denoted by a, are used for input transitions. Finally, the unique element of the third term is denoted by τ and used for silent transitions. E.g., the labelled transition system

$$x \xleftarrow{\ \overline{a}\ } y \overset{b}{\underset{b}{\rightrightarrows}} z \,\overset{a}{\underset{}{\circlearrowleft}}$$

is modelled by the presheaf X with

$$X(\star) = \{x, y, z\} \qquad\qquad \begin{aligned} X(\overline{a}) &= \{e\} \\ X(b) &= \{f, f'\} \\ X(a) &= \{g\} \end{aligned} \qquad\qquad \begin{aligned} x &= e \cdot t \\ y &= e \cdot s = f \cdot s = f' \cdot s \\ z &= f \cdot t = f' \cdot t = g \cdot s = g \cdot t. \end{aligned}$$
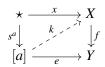
2.2 Bisimulation

Returning to generalised transition systems, we may define bisimulation categorically in the following way. Morphisms $f : X \to Y$, i.e., natural transformations, are the analogue in our setting of standard functional simulations. Indeed, given any transition $e : x \xrightarrow{a} x'$ in X, then $f(x)$ sure has an a transition to some state related to x': this is simply $f(e)$! The next step is to define an analogue of functional bisimulation. For this, let us observe that the base category $\Gamma_{\mathbb{A}}$ embeds into the presheaf category $\widehat{\Gamma_{\mathbb{A}}}$ – this is just the Yoneda embedding $\mathbf{y} : \Gamma_{\mathbb{A}} \to \widehat{\Gamma_{\mathbb{A}}}$, directly specialised to our setting for readability:

- the state object $\star$ embeds as the one-vertex graph $\mathbf{y}_\star$ with no transition;

- any transition object $[a]$ embeds as the graph $\mathbf{y}_{[a]}$ with one a-transition between two distinct vertices;

- the morphisms $s^a, t^a : \star \to [a]$ embed as the morphisms $\mathbf{y}_\star \to \mathbf{y}_{[a]}$ picking up the source and target, respectively, of the given transition.

Notation 2.5. We often omit $\mathbf{y}$, treating it as an implicit coercion.

Definition 2.6. Let $f : X \to Y$ be a *functional bisimulation* whenever all commuting squares as the solid part below, admit a (potentially non-unique) *lifting k* as shown, i.e., a morphism making both triangles commute.

$$
\begin{array}{ccc}
\star & \xrightarrow{\ x\ } & X \\
{\scriptstyle s^a}\big\downarrow & \overset{k}{\nearrow} & \big\downarrow{\scriptstyle f} \\
[a] & \xrightarrow[\ e\ ]{} & Y
\end{array}
$$

Let us explain why this matches the standard definition. In any such square, x is essentially the same as just a state in X, while e is just an a-transition in Y. Furthermore, the composite $\star \xrightarrow{s^a} [a] \xrightarrow{e} Y$ picks the source of e, so commutation of the square says that the source $e \cdot s^a$ of e is in fact $f(x)$ (a.k.a. $f \circ x$). So we are in the situation described by the solid part below.

$$
\begin{array}{ccc}
x & \xmapsto{\ f\ } & f(x) \\
{\scriptstyle k}\big\downarrow & & \big\downarrow{\scriptstyle e} \\
x' & \dashrightarrow[\ f\ ] & y'
\end{array}
$$

Finding a lifting k then amounts to finding an antecedent to e whose source is x, as desired.

We finally recover the analogue of standard bisimulation relations.

Definition 2.7. A *bisimulation relation* on X is a subobject $R \hookrightarrow X^2$ (= isomorphism class of monomorphisms into X^2) whose projections $R \to X$ are both functional bisimulations.

In this case, the above diagram specialises to

$$
\begin{array}{ccc}
(x_1, x_2) & \xmapsto{\ \pi_i\ } & x_i \\
{\scriptstyle (e_1, e_2)}\big\downarrow & & \big\downarrow{\scriptstyle e_i} \\
(x'_1, x'_2) & \dashrightarrow[\ \pi_i\ ] & x'_i,
\end{array}
$$

where $(x_1, x_2), (x'_1, x'_2) \in R(\star)$, and $(e_1, e_2) \in R[a]$.

Now, $\widehat{\Gamma_\mathbb{A}}$, as a presheaf category, is very well-behaved, namely it is a Grothendieck topos [17]. In particular, subobjects of X^2 form a (small) complete lattice, in which the union of a family $R_i \hookrightarrow X^2$ is computed by first taking the copairing $\sum_i R_i \to X^2$, which is generally not monic, and then taking its image. Furthermore, bisimulation relations are closed under unions and so admit a maximum element, *bisimilarity* [12, Proposition 3.14].

The presheaf category $\widehat{\Gamma_\mathbb{A}}$ is thus only a slight generalisation of standard labelled transition systems over $\mathbb{A}$, in which we have an analogue of bisimulation, conveniently defined by lifting, and bisimilarity. Let us now consider the case where states are terms in a certain language, and transitions are defined inductively by a set of transition rules, i.e., operational semantics.

3 Positive GSOS specifications as monads

Let us briefly recall the Positive GSOS format. We fix a set $\mathbb{A}$ of labels, and start from a *signature* $\Sigma_0 = (O_0, E_0)$ on **Set**, i.e., a set O_0 equipped with a map $E_0 : O_0 \to \mathbb{N}$.

Definition 3.1. A *Positive GSOS rule* over Σ_0 consists of

- an operation $f \in O_0$, say of arity $n = E_0(f)$,
- a label $a \in \mathbb{A}$,
- n natural numbers $m_1, \ldots, m_n$,
- for all $i \in n$, m_i labels $a_{i,1}, \ldots, a_{i,m_i}$, and
- a term t with $n + \sum_{i=1}^{n} m_i$ free variables.

In more standard form, such a rule is just

$$\frac{\ldots \qquad x_i \xrightarrow{a_{i,j}} y_{i,j} \qquad \ldots \qquad (i \in n, j \in m_i)}{f(x_1, \ldots, x_n) \xrightarrow{a} t}$$

where the x_i's and $y_{i,j}$'s are all distinct and denote the potential free variables of t.

Definition 3.2. A *Positive GSOS specification* is a signature Σ_0, together with a set Σ_1 of Positive GSOS rules.

Let us now describe how any Positive GSOS specification Σ induces a monad T_Σ on $\widehat{\Gamma_\mathbb{A}}$, starting with the action of T_Σ on objects. Given any $X \in \widehat{\Gamma_\mathbb{A}}$, the set $T_\Sigma(X)(\star)$ of states consists of all Σ_0-terms with variables in $X(\star)$, as defined by the grammar

$$M, N ::= (\!|u|\!) \mid f(M_1, \ldots, M_n),$$

where u ranges over $X(\star)$. Similarly, each $T_\Sigma(X)[a]$ consists of all transition proofs following the rules in Σ_1, with axioms in all $X[a']$'s. Formally, such proofs are constructed inductively from the following rules,

$$\frac{}{(\!|e|\!) :_X (\!|e \cdot s|\!) \xrightarrow{a} (\!|e \cdot t|\!)} (e \in X[a]) \qquad \frac{\ldots \qquad R_{i,j} :_X M_i \xrightarrow{a_{i,j}} M_{i,j} \qquad \ldots \qquad (i \in n, j \in m_i)}{\rho(R_{i,j})_{i \in n, j \in m_i} :_X f(M_1, \ldots, M_n) \xrightarrow{a} t[(x_i \mapsto M_i, (y_{i,j} \mapsto M_{i,j})_{j \in m_i})_{i \in n}]}$$

where in the second rule $f \in O_0$, $E_0(f) = n$, $\rho = (f, a, (m_i, (a_{i,j})_{j \in m_i})_{i \in n}, t) \in \Sigma_1$. When $m_i = 0$, we want to keep track of M_i in the transition proof, so by convention the family $(R_{i,j})_{j \in m_i}$ denotes just M_i. In the sequel we simply call *transitions* such transition proofs.

Example 3.3. Let us consider the following simple CCS transition of depth > 1, in any $T_{CCS}(X)[\tau]$.

$$\frac{\dfrac{}{(\!|e_1|\!) :_X (\!|x_1|\!) \xrightarrow{\bar{a}} (\!|y_1|\!)}}{lpar((\!|e_1|\!), (\!|x_2|\!)) :_X (\!|x_1|\!)|(\!|x_2|\!) \xrightarrow{\bar{a}} (\!|y_1|\!)|(\!|x_2|\!)} \qquad \dfrac{}{(\!|e_2|\!) :_X (\!|x_3|\!) \xrightarrow{a} (\!|y_2|\!)}}{sync(lpar((\!|e_1|\!), (\!|x_2|\!)), (\!|e_2|\!)) :_X ((\!|x_1|\!)|(\!|x_2|\!))|(\!|x_3|\!) \xrightarrow{\tau} ((\!|y_1|\!)|(\!|x_2|\!))|(\!|y_2|\!)} ,$$

where *lpar* and *sync* denote the left parallel and synchronisation rules, $(e_1 : x_1 \xrightarrow{\bar{a}} y_1) \in X[\bar{a}]$, $x_2 \in X(\star)$, and $(e_2 : x_3 \xrightarrow{a} y_2) \in X[a]$.

The source and target of a transition $R :_X M \xrightarrow{a} N$ are M and N, respectively, which ends the definition of T_Σ on objects. On morphisms $f : X \to Y$, $T_\Sigma(f)$ merely amounts to renaming variables $(\!|x|\!)$ and $(\!|e|\!)$ to $(\!|f(x)|\!)$ and $(\!|f(e)|\!)$, respectively. It thus remains to show that T_Σ has monad structure. The unit $\eta_X : X \to T_\Sigma(X)$ is obviously given by $(\!|-|\!)$, while multiplication $\mu_X : T_\Sigma(T_\Sigma(X)) \to T_\Sigma(X)$ is given inductively by removing the outer layer of $(\!|-|\!)$'s

$$
\begin{array}{rrcl}
\text{on states} & \mu_X(\!|M|\!) & = & M \\
& \mu_X(f(M_1, \ldots, M_n)) & = & f(\mu_X(M_1), \ldots, \mu_X(M_n)) \\
\text{and on transitions} & \mu_X(\!|R|\!) & = & R \\
& \mu_X(\rho(R_{i,j})_{i \in n, i \in m_i}) & = & \rho(\mu_X(R_{i,j}))_{i \in n, i \in m_i}.
\end{array}
$$

Lemma 3.4. *The natural transformations η and μ equip T_Σ with monad structure.*

Proof. A straightforward induction. □

4 Models as algebras and congruence of bisimilarity

Algebras for T_Σ readily give the right notion of model for the transition rules:

Definition 4.1. An *algebra* for a monad T, or a T-*algebra*, consists of an object X, equipped with a morphism $\alpha : T(X) \to X$ such that the following diagrams commute.

$$
\begin{array}{ccc}
T(T(X)) & \xrightarrow{T(\alpha)} & T(X) \\
{\scriptstyle \mu_X}\downarrow & & \downarrow{\scriptstyle \alpha} \\
T(X) & \xrightarrow{\alpha} & X
\end{array}
\qquad\qquad
\begin{array}{ccc}
X & \xrightarrow{\eta_X} & T(X) \\
 & {\scriptstyle}\searrow{\scriptstyle \alpha} \\
 & X
\end{array}
$$

Thus, intuitively, a T_Σ-algebra is a transition system which is stable under the given operations and transition rules.

We now would like to show that, under suitable hypotheses, bisimilarity for any given T_Σ-algebra $\alpha : T_\Sigma(X) \to X$ is a congruence. We may state this categorically by saying that the canonical morphism $T_\Sigma(\sim_X) \to X^2$ factors through $m : (\sim_X) \hookrightarrow X^2$, as in

$$
\begin{array}{ccc}
T_\Sigma(\sim_X) & \dashrightarrow & \sim_X \\
{\scriptstyle T_\Sigma(m_X)}\downarrow & & \downarrow{\scriptstyle m} \\
T_\Sigma(X^2) & \xrightarrow[\langle T_\Sigma(\pi_1), T_\Sigma(\pi_2)\rangle]{} (T_\Sigma(X))^2 \xrightarrow[\alpha^2]{} & X^2.
\end{array}
\tag{1}
$$

Indeed, an element of $T_\Sigma(\sim_X)$ is a term M whose free variables are pairs of bisimilar elements of X, which we write as $M((x_1, y_1), \ldots, (x_n, y_n))$, with $x_i \sim_X y_i$ for all $i \in n$. The morphism $\langle T_\Sigma(\pi_1), T_\Sigma(\pi_2)\rangle$ maps this to the pair

$$
(M(x_1, \ldots, x_n), M(y_1, \ldots, y_n)),
$$

which α^2 then evaluates componentwise. The given factorisation thus boils down to

$$
\alpha(M(x_1, \ldots, x_n)) \sim_X \alpha(M(y_1, \ldots, y_n))
$$

for all M and $x_1 \sim_X y_1, \ldots, x_n \sim_X y_n$, i.e., bisimilarity is a congruence.

In order to prove such a property, it is sufficient to prove that T_Σ preserves *all* bisimulation relations, in the sense that if $m : R \hookrightarrow X^2$ is a bisimulation relation, then so is

$$
T_\Sigma(R) \xrightarrow{T_\Sigma(m)} T_\Sigma(X^2) \xrightarrow{\langle T_\Sigma(\pi_1), T_\Sigma(\pi_2)\rangle} (T_\Sigma(X))^2 \xrightarrow{\alpha^2} X^2
$$

(in the slightly generalised sense that its image is). Equivalently, an easy diagram chasing shows that it all boils down to

$$
T_\Sigma(R) \xrightarrow{T_\Sigma(m)} T_\Sigma(X^2) \xrightarrow{T_\Sigma(\pi_i)} T_\Sigma(X) \xrightarrow{\alpha} X
$$

being a functional bisimulation for $i \in \{1, 2\}$.

Finally, $\pi_i \circ m$ is a functional bisimulation by definition, and functional bisimulations are stable under composition, so it is sufficient to prove that

(i) the considered algebra is *compositional*, in the sense that its structure map $\alpha : T_\Sigma(X) \to X$ is a functional bisimulation, and

(ii) T_Σ preserves all functional bisimulations.

Compositionality essentially means that transitions of any $\alpha(M(x_1, \ldots, x_n))$ are all obtained by assembling transitions of the x_i's. This is not always the case, even for free algebras:

Example 4.2. Consider a specification Σ consisting of the unique rule

$$\frac{x \xrightarrow{a} y}{f(g(x)) \xrightarrow{a} f(g(y))} \ ,$$

say ρ, where f and g are two unary operations. Then the free algebra $\mu_1 : T_\Sigma(T_\Sigma(1)) \to T_\Sigma(1)$ is not compositional. Indeed, 1 contains a unique vertex, say $\star$, and a transition $b : \star \xrightarrow{b} \star$ for all labels b. Thus, $T_\Sigma(1)$ contains a transition $\rho(\!(a)\!) : f(g(\!(\star)\!)) \xrightarrow{a} f(g(\!(\star)\!))$. But the term $f(g(\!(\star)\!))$ is the image under μ_1 of $f(\!(g(\!(\star)\!))\!)$, which has no transition.

Summing up, we have proved:

Lemma 4.3. *If T_Σ preserves functional bisimulations, then bisimilarity in any compositional T_Σ-algebra is a congruence.*

5   Compositionality

Let us first consider compositionality. For a general algebra, we cannot do more than taking compositionality as a hypothesis. However, we can say something when the considered algebra is free:

Lemma 5.1. *The multiplication $\mu_X : T_\Sigma(T_\Sigma(X)) \to T_\Sigma(X)$ is a functional bisimulation.*

Proof. We will see below (Lemma 10.3) that all naturality squares of μ are pullbacks. In particular, we have a pullback

$$\begin{array}{ccc}
T_\Sigma(T_\Sigma(X)) & \xrightarrow{T_\Sigma(T_\Sigma(!))} & T_\Sigma(T_\Sigma(1)) \\
\mu_X \downarrow & & \downarrow \mu_1 \\
T_\Sigma(X) & \xrightarrow[T_\Sigma(!)]{} & T_\Sigma(1).
\end{array}$$

But functional bisimulations are easily seen to be stable under pullback, so it is enough to show that μ_1 is a functional bisimulation. We thus consider any term $\mathbf{M}$ whose free variables are in $T_\Sigma(1)(\star)$, i.e., are themselves terms over a single free variable, say $\star$, together with a transition $R : \mu_1(\mathbf{M}) \xrightarrow{a} N$. And we need to show that there exists a transition $\mathbf{R} : \mathbf{M} \xrightarrow{a} \mathbf{N}$ whose free variables and axioms are in $T_\Sigma(1)$, such that $\mu_{[a]}(\mathbf{R}) = R$. We proceed by induction on $\mathbf{M}$:
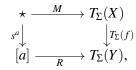
- If $\mathbf{M} = (\!(M)\!)$, then taking $\mathbf{R} = (\!(R)\!)$ does the job.

- Otherwise, $\mathbf{M} = f(\mathbf{M}_1, \ldots, \mathbf{M}_n)$, so $M = \mu_1(\mathbf{M}) = f(M_1, \ldots, M_n)$, with $M_i = \mu_1(\mathbf{M}_i)$ for all $i \in n$. But then, R must have the form $\rho(R_{i,j})_{i \in n, j \in m_i}$, for a certain rule $\rho = (f, a, (m_i, (a_{i,j})_{j \in m_j})_{i \in n}, t)$ of Σ. By induction hypothesis, we find for all $i \in n$ and $j \in m_i$ a transition

$$\mathbf{R}_{i,j} : \mathbf{M}_i \xrightarrow{a_{i,j}} \mathbf{N}_{i,j},$$

  such that $\mu_1(\mathbf{R}_{i,j}) = R_{i,j}$. Thus, $\mathbf{R} = \rho(\mathbf{R}_{i,j})_{i \in n, j \in m_i}$ does have $\mathbf{M}$ as its source, and furthermore satisfies $\mu_1(\mathbf{R}) = R$, as desired. $\square$

6   Preserving bisimulations through familiality and cellularity

Let us now consider *(ii)*, i.e., the fact that T_Σ preserves functional bisimulations. So we need to find a lifting to any commuting square of the form

$$
\begin{array}{ccc}
\star & \xrightarrow{\ M\ } & T_\Sigma(X) \\
{\scriptstyle s^a}\downarrow & & \downarrow{\scriptstyle T_\Sigma(f)} \\
[a] & \xrightarrow{\ R\ } & T_\Sigma(Y),
\end{array}
$$

for any functional bisimulation f.

We will proceed in two steps: we will require T_Σ to be first familial, and then cellular. Familiality will allow us to factor the given square as the solid part below left, while cellularity will ensure existence of a lifting k as on the right.

$$
(2)
$$

The composite $T_\Sigma(k) \circ R'$ will thus give the desired lifting for the original square.

At this stage, both steps may seem mysterious to the reader. In fact, as we will see, factorisation as above left follows directly from the fact that T_Σ may be expressed as a sum of representable functors. Let us first explain intuitively why this latter fact holds. We will then prove it more rigorously in §7, to eventually return to factorisation in §8.

To start with, let us observe that the set $T_\Sigma(1)(\star)$ consists of terms over a single free variable, say $\star$. For any such term M, we may count the number of occurrences of $\star$, say n_M. Thus, any term in any $T_\Sigma(X)(\star)$ is entirely determined by an $M \in T_\Sigma(1)(\star)$, together with a map $n_M \to X(\star)$ assigning an element of $X(\star)$ to each occurrence of $\star$ in M. But maps $n_M \to X(\star)$ in **Set** are in 1-1 correspondence with maps $n_M \cdot \mathbf{y}_\star \to X$ in $\widehat{\Gamma_\mathbb{A}}$, where $n_M \cdot \mathbf{y}_\star$ denotes the n_M-fold coproduct $\mathbf{y}_\star + \cdots + \mathbf{y}_\star$ of $\mathbf{y}_\star$ with itself. In other words, letting $E^\star(M) = n_M \cdot \mathbf{y}_\star$, we have

$$
T_\Sigma(X)(\star) \cong \sum_{M \in T_\Sigma(1)(\star)} \widehat{\Gamma_\mathbb{A}}(E^\star(M), X). \tag{3}
$$

Clearly, for any $f : X \to Y$, the action of $T_\Sigma(f)$ at $\star$ is given by postcomposing with f, i.e., we have

$$
\begin{array}{ccc}
T_\Sigma(X)(\star) & \xrightarrow{\ \simeq\ } & \sum_{M \in T_\Sigma(1)(\star)} \widehat{\Gamma_\mathbb{A}}(E^\star(M), X) \\
{\scriptstyle T_\Sigma(f)_\star}\downarrow & & \downarrow{\scriptstyle \sum_{M \in T_\Sigma(1)(\star)} \widehat{\Gamma_\mathbb{A}}(E^\star(M), f)} \\
T_\Sigma(Y)(\star) & \xrightarrow{\ \simeq\ } & \sum_{M \in T_\Sigma(1)(\star)} \widehat{\Gamma_\mathbb{A}}(E^\star(M), Y).
\end{array}
$$

The family (3) of isomorphisms is thus natural in X. We will see shortly that this extends to objects other than $\star$. Indeed, any transition in $T_\Sigma(X)[a]$ may be decomposed into a transition R in $T_\Sigma(1)[a]$, together with a morphism $E^a(R) \to X$, where $E^a(R)$ is obtained from occurrences of term and transition variables in R.

We have seen that our isomorphisms are natural in X, so it seems natural to try to express some naturality constraint in the second argument of T_Σ. But this requires making the right-hand side of (3)

functorial in this variable in the first place! In fact, for any transition $R : M \xrightarrow{a} N$, we will construct morphisms

$$E^\star(M) \xrightarrow{E(s^a \restriction R)} E^a(R) \xleftarrow{E(t^a \restriction R)} E^\star(N)$$

(see Notation 7.2 below). Thus, e.g., precomposing by the left-hand map yields the desired functorial action

$$\sum_{R \in T_\Sigma(1)[a]} \widehat{\Gamma_{\mathbb{A}}}(E^a(R), X) \to \sum_{M \in T_\Sigma(1)(\star)} \widehat{\Gamma_{\mathbb{A}}}(E^\star(M), X),$$

of $s^a : \star \to [a]$, sending any $\varphi : E^a(R) \to X$ to the composite

$$E^\star(M) \xrightarrow{E(s^a \restriction R)} E^a(R) \xrightarrow{\varphi} X. \tag{4}$$

7 Familiality for functors

Let us now state more rigorously the definition of familiality and the fact that T_Σ is familial. In the next section, we will explain how this entails the desired factorisation (2).

Definition 7.1. The *category of elements* $el(X)$ of any presheaf $X \in \widehat{\mathscr{C}}$ on any category $\mathscr{C}$ has

- as objects all pairs (c, x) with $c \in \mathbf{ob}(\mathscr{C})$ and $x \in X(c)$,

- and as morphisms $(c, x) \to (c', x')$ all morphisms $f : c \to c'$ such that $x' \cdot f = x$.

Notation 7.2. The morphism f, viewed as a morphism $(c, x) \to (c', x')$, is entirely determined by f and x'. We denote it by $f \restriction x'$.

Definition 7.3. An endofunctor $F : \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ on a presheaf category is *familial* iff there is a functor $E : el(F(1)) \to \widehat{\mathbb{C}}$ such that

$$F(X)(c) \cong \sum_{o \in F(1)(c)} \widehat{\mathbb{C}}(E(c, o), X), \tag{5}$$

naturally in $X \in \widehat{\mathbb{C}}$ and $c \in \mathbb{C}$.

And indeed, we have:

Lemma 7.4. *The endofunctor T_Σ is familial.*

Proof. We need to do two things: (1) extend the isomorphisms (3) to objects of the form $[a]$, and (2) define the morphisms $E(s^a \restriction R)$ and $E(t^a \restriction R)$ rendering our isomorphisms natural also in the second argument of T_Σ. In fact, we will do almost everything simultaneously by induction: we define $E^a(R)$ and $E(s^a \restriction R) : E^\star(R \cdot s^a) \to E^a(R)$ by induction on R. By convention, as we did for the unique element $\star \in 1(\star)$, we denote the unique element of $1[a]$ by a itself.

- If $R = (a)$, then its source is $M = (\star)$ and we put $E^a(R) = \mathbf{y}_{[a]}$ and $E(s^a \restriction R) = s^a : \mathbf{y}_\star \to \mathbf{y}_{[a]}$.

- If $R = \rho(R_{i,j})_{i \in n, j \in m_i} :_1 M \xrightarrow{a} N$, then for all i and $j \in m_i$, by induction hypothesis, we get morphisms

$$E(s^{a_{i,j}} \restriction R_{i,j}) : E^\star(M_i) \to E^{a_{i,j}}(R_{i,j}),$$

where $R_{i,j} :_1 M_i \xrightarrow{a_{i,j}} N_{i,j}$ for all $i \in n$ and $j \in m_i$. Let us temporarily fix any $i \in n$. For all $j, j' \in m_i$, we have $R_{i,j} \cdot s^a = R_{i,j'} \cdot s^a = M_i$, so we take the wide pushout $E_i = \bigoplus_{E^\star(M_i)} E^{a_{i,j}}(R_{i,j})$, i.e., the colimit of the following diagram.

$$
\begin{array}{ccc}
& E^\star(M_i) & \\
{\scriptstyle E(s^{a_{i,j}} \restriction R_{i,j})} \nearrow & & \searrow {\scriptstyle E(s^{a_{i,j'}} \restriction R_{i,j'})} \\
\dots E^{a_{i,j}}(R_{i,j}) & \dots & E^{a_{i,j'}}(R_{i,j'}) \dots
\end{array}
\tag{6}
$$

If $m_i = 0$, this reduces to just $E^\star(M_i)$, which is exactly what we want. Finally, we let $E^a(R)$ be the coproduct $\sum_i E_i$ of all the E_i's, and observe that $E^\star(f(M_1, \ldots, M_n)) = \sum_i E^\star(M_i)$ by definition, so that we may define $E(s^a \restriction R)$ to be the coproduct $\sum_i S_i$ of all canonical injections $S_i : E^\star(M_i) \to E_i$.

This ends the inductive definition of $E^a(R)$ and $E(s^a \restriction R)$. We now need to construct the morphisms $E(t^a \restriction R)$. We again proceed inductively. When $R = (\!|a|\!)$, the desired morphism is clearly t^a itself. When $R = \rho(R_{i,j})_{i \in n, j \in m_i}$, the target is $N = t[x_i \mapsto M_i, (y_{i,j} \mapsto N_{i,j})_{j \in m_i}]$. Now, by construction, occurrences $occ_\star(N)$ of the unique variable $\star$ in N are in 1-1 correspondence with

$$V_N = \sum_i \left((occ_\star(M_i))^{occ_{x_i}(t)} + \sum_{j \in m_i} (occ_\star(N_{i,j}))^{occ_{y_{i,j}}(t)} \right),$$

and our map $E(t^a \restriction R)$ should reflect the intended correspondences. Since $E^\star(N) = V_N \cdot \mathbf{y}_\star$, $E(t^a \restriction R)$ is entirely determined by choosing a map $E_u : \mathbf{y}_\star \to E^a(R)$ for all $u \in V_N$:

- If u denotes an occurrence of $\star$ in M_i, for some occurrence of x_i in t, we let E_u denote the composite

$$\mathbf{y}_\star \to E^\star(M_i) \to E^\star(R),$$

  where the latter map denotes injection into the colimit of (6).

- If u denotes an occurrence of $\star$ in $N_{i,j}$, for some occurrence of $y_{i,j}$ in t, we let E_u denote the composite

$$\mathbf{y}_\star \to E^\star(N_{i,j}) \xrightarrow{E(t^{a_{i,j}} \restriction R_{i,j})} E^{a_{i,j}}(R_{i,j}) \to E^\star(R),$$

  where the latter map again denotes injection into the colimit of (6). $\qquad\square$

Rather than a full formal proof, let us illustrate that our construction satisfies the isomorphisms (3) on a few examples.

Example 7.5. In the case of the transition of Example 3.3, familiality means that this transition is determined by picking the following transition in $T_{CCS}(1)[\tau]$,

$$\cfrac{\cfrac{(\!|\bar{a}|\!) :_1 (\!|\star|\!) \xrightarrow{\bar{a}} (\!|\star|\!)}{lpar((\!|\bar{a}|\!), (\!|\star|\!)) :_1 (\!|\star|\!) | (\!|\star|\!) \xrightarrow{\bar{a}} (\!|\star|\!) | (\!|\star|\!)} \qquad (\!|a|\!) :_1 (\!|\star|\!) \xrightarrow{a} (\!|\star|\!)}{sync(lpar((\!|\bar{a}|\!), (\!|\star|\!)), (\!|a|\!)) :_1 ((\!|\star|\!) | (\!|\star|\!)) | (\!|\star|\!) \xrightarrow{\tau} ((\!|\star|\!) | (\!|\star|\!)) | (\!|\star|\!)}$$

together with a morphism $E^\tau(sync(lpar((\!|\bar{a}|\!), (\!|\star|\!)), (\!|a|\!))) \to X$. Let us start with $E^\star(lpar((\!|\bar{a}|\!), (\!|\star|\!)))$: it is given by the colimit of

$$\begin{array}{ccc} \mathbf{y}_\star & & \mathbf{y}_\star \\ {\scriptstyle s^{\bar{a}}}\downarrow & & \\ \mathbf{y}_{[\bar{a}]} & & \end{array}$$

(one $\mathbf{y}_\star$ for each argument x_i of $lpar$, and for each x_i one $\mathbf{y}_{[a_{i,j}]}$ for each premise $x_i \xrightarrow{a_{i,j}} y_{i,j}$). Equivalently, this is just the coproduct $\mathbf{y}_{[\bar{a}]} + \mathbf{y}_\star$, and $E(s^{\bar{a}} \restriction lpar((\!|\bar{a}|\!), (\!|\star|\!)))$ and $E(t^{\bar{a}} \restriction lpar((\!|\bar{a}|\!), (\!|\star|\!)))$ are given by

$$\mathbf{y}_\star + \mathbf{y}_\star \xrightarrow{s^{\bar{a}} + \mathbf{y}_\star} \mathbf{y}_{[\bar{a}]} + \mathbf{y}_\star \xleftarrow{t^{\bar{a}} + \mathbf{y}_\star} \mathbf{y}_\star + \mathbf{y}_\star.$$

It is then clear that $E(s^\tau \restriction sync(lpar((\![\overline{a}]\!), (\![\star]\!)), (\![a]\!)))$ is given by

$$\mathbf{y}_\star + \mathbf{y}_\star + \mathbf{y}_\star \xrightarrow{s^{\overline{a}}+\mathbf{y}_\star+s^a} \mathbf{y}_{[\overline{a}]} + \mathbf{y}_\star + \mathbf{y}_{[a]}.$$

Now how about $E(t^\tau \restriction sync(lpar((\![\overline{a}]\!), (\![\star]\!)), (\![a]\!)))$? As the term t occurring in the rule is here linear in the $y_{i,j}$'s, an easy computation leads to

$$\mathbf{y}_{[\overline{a}]} + \mathbf{y}_\star + \mathbf{y}_{[a]} \xleftarrow{t^{\overline{a}}+\mathbf{y}_\star+t^a} \mathbf{y}_\star + \mathbf{y}_\star + \mathbf{y}_\star.$$

On this example, the isomorphism (5) thus boils down to the transition $sync(lpar((\![e_1]\!), (\![x_2]\!)), (\![e_2]\!))$ above being entirely determined by picking $R = sync(lpar((\![\overline{a}]\!), (\![\star]\!)), (\![a]\!)) \in T_{CCS}(1)[\tau]$, and giving a morphism

$$\varphi : E^\tau(sync(lpar((\![\overline{a}]\!), (\![\star]\!)), (\![a]\!))) = \mathbf{y}_{[\overline{a}]} + \mathbf{y}_\star + \mathbf{y}_{[a]} \longrightarrow X,$$

which holds by universal property of coproduct and the Yoneda lemma. Naturality of (5) in c says that the source of (R, φ) is given up to this correspondence by $((\![\star]\!)|(\![\star]\!))|(\![\star]\!)$ and the composite

$$E^\star(((\![\star]\!)|(\![\star]\!))|(\![\star]\!)) = \mathbf{y}_\star + \mathbf{y}_\star + \mathbf{y}_\star \xrightarrow{s^{\overline{a}}+\mathbf{y}_\star+s^a} \mathbf{y}_{[\overline{a}]} + \mathbf{y}_\star + \mathbf{y}_{[a]} \longrightarrow X,$$

and likewise for the target.

Example 7.6. Let us now illustrate the treatment of branching, in the sense of a rule having several premises involving the same x_i. An example from CCS is the 'replicated synchronisation' rule

$$\frac{x_1 \xrightarrow{\overline{a}} y_{1,1} \qquad x_1 \xrightarrow{a} y_{1,2}}{!x_1 \xrightarrow{\tau} !x_1|(y_{1,1}|y_{1,2})} \ ,$$

say *rsync*. First, $E^\tau(rsync((\![[\overline{a}]]\!), (\![[a]]\!)))$ is simply the pushout

$$
\begin{array}{ccc}
\mathbf{y}_\star & \xrightarrow{\ s^a\ } & \mathbf{y}_{[a]} \\
{\scriptstyle s^{\overline{a}}}\downarrow & & \downarrow \\
\mathbf{y}_{[\overline{a}]} & \longrightarrow & E^\tau(rsync((\![[\overline{a}]]\!), (\![[a]]\!))),
\end{array}
$$

which rightly models the fact that a transition $rsync((\![e_1]\!), (\![e_2]\!)) \in T_{CCS}(X)[\tau]$ is entirely determined by picking $rsync((\![[\overline{a}]]\!), (\![[a]]\!)) \in T_{CCS}(1)[\tau]$, together with elements e_1 and e_2 of $X[\overline{a}]$ and $X[a]$ with a common source. The morphism $E(s^\tau \restriction rsync((\![[\overline{a}]]\!), (\![[a]]\!)))$ is then straightforwardly given by the diagonal. The target morphism $E(t^\tau \restriction rsync((\![[\overline{a}]]\!), (\![[a]]\!)))$ is a bit more complex to compute. Indeed, the target $t = !x_1|(y_{1,1}|y_{1,2})$ has three free variables. The first, x_1, should yield a morphism $\mathbf{y}_\star \to E^\tau(rsync((\![[\overline{a}]]\!), (\![[a]]\!)))$ that is determined by the source morphism $E^\star(M_1) \to E^{\overline{a}}(R_{1,1})$. Here, we get

$$\mathbf{y}_\star \xrightarrow{s^{\overline{a}}} E^{\overline{a}}((\![[\overline{a}]]\!)) = \mathbf{y}_{[\overline{a}]} \to E^\tau(rsync((\![[\overline{a}]]\!), (\![[a]]\!))).$$

On the other hand, $y_{1,1}$ and $y_{1,2}$ should be determined by the target morphisms $E^\star(M_{1,1}) \to E^{\overline{a}}(R_{1,1})$ and $E^\star(M_{1,2}) \to E^a(R_{1,2})$, in our case

$$\mathbf{y}_\star \xrightarrow{t^{\overline{a}}} E^{\overline{a}}((\![[\overline{a}]]\!)) = \mathbf{y}_{[\overline{a}]} \to E^\tau(rsync((\![[\overline{a}]]\!), (\![[a]]\!))) \quad \text{and} \quad \mathbf{y}_\star \xrightarrow{t^a} E^a((\![a]\!)) = \mathbf{y}_a \to E^\tau(rsync((\![[\overline{a}]]\!), (\![[a]]\!))).$$

8   Familiality and factorisation

Let us now return to our proof sketch (2), and explain the properties of familiality that allow us to factor the original square as indicated. The crucial observation is that elements of the form

$$(R, id_{E^c(R)}) \in \sum_{R \in T_\Sigma(1)(c)} \widehat{\Gamma_\mathbb{A}}(E^c(R), E^c(R)) \cong T_\Sigma(E^c(R))(c)$$

have the special property that any other element of the form $(R, \varphi) \in T_\Sigma(X)(c)$ may be obtained uniquely as the image of (R, id) by the action of

$$T_\Sigma(\varphi)_c : T_\Sigma(E^c(R))(c) \to T_\Sigma(X)(c).$$

Having the same first component R is equivalent to having the same image in $T_\Sigma(1)(c)$. So by Yoneda, having two elements of $T_\Sigma(X)(c)$ and $T_\Sigma(Y)(c)$ with common first component is the same as having a commuting square of the form below left.

$$
\begin{array}{ccc}
\mathbf{y}_c & \longrightarrow & T_\Sigma(X) \\
\downarrow & & \downarrow \\
T_\Sigma(Y) & \longrightarrow & T_\Sigma(1)
\end{array}
\qquad\qquad
\begin{array}{ccc}
\mathbf{y}_c & \overset{p}{\longrightarrow} & T_\Sigma(X) \\
\xi \downarrow & \overset{T_\Sigma(k)}{\nearrow} & \downarrow T_\Sigma(!) \\
T_\Sigma(E^c(R)) & \underset{T_\Sigma(!)}{\longrightarrow} & T_\Sigma(1)
\end{array}
$$

The special property of (R, id) is thus equivalently that any commuting square as above right (solid part) admits a unique (dashed) lifting k as shown, making the non-trivial triangle commute. In fact, this holds more generally by replacing $\mathbf{y}_c$ and 1 by arbitrary objects:

Definition 8.1. Given a functor $F : \mathscr{C} \to \mathscr{D}$, a morphism $\xi : D \to F(C)$ is F-*generic*, or *generic* for short, when any commuting square as the solid part of

$$
\begin{array}{ccc}
D & \overset{\chi}{\longrightarrow} & F(B) \\
\xi \downarrow & \overset{F(l)}{\nearrow} & \downarrow F(h) \\
F(C) & \underset{F(k)}{\longrightarrow} & F(A)
\end{array}
$$

admits a unique *strong* lifting l as shown, in the sense that $F(l) \circ \xi = \chi$ and $h \circ l = k$.

Lemma 8.2 ([26, Remark 2.12]). *A functor* $F : \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ *is familial iff any morphism* $Y \to F(X)$ *factors as*

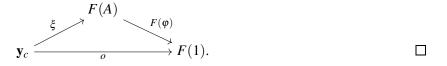$$Y \overset{\xi}{\to} F(A) \overset{F(\varphi)}{\longrightarrow} F(X),$$

where ξ *is F-generic. This is called a* generic-free *factorisation.*

Proof sketch. ($\Rightarrow$) Passing from $\mathbf{y}_c$ to any Y goes by observing that generic morphisms are stable under colimits in the comma category $\widehat{\mathbb{C}} \downarrow F$, remembering that any presheaf Z is a colimit of the composite

$$el(Z) \overset{\mathbf{p}_Z}{\to} \mathbb{C} \overset{\mathbf{y}}{\to} \widehat{\mathbb{C}},$$

where $\mathbf{p}_Z$ denotes the obvious projection functor.
($\Leftarrow$) Conversely, $E(c, o)$ is given by A, for any choice of generic-free factorisation

$$
\begin{array}{ccc}
 & F(A) & \\
\overset{\xi}{\nearrow} & & \overset{F(\varphi)}{\searrow} \\
\mathbf{y}_c & \underset{o}{\longrightarrow} & F(1).
\end{array}
$$
$\qquad\square$

Lemma 8.2 thus accounts for the factorisation of the original square as on the left in (2): M and R respectively factor as

$$\star \xrightarrow{M'} T_\Sigma(A) \xrightarrow{T_\Sigma(\varphi)} T_\Sigma(X) \qquad \text{and} \qquad \star \xrightarrow{R'} T_\Sigma(B) \xrightarrow{T_\Sigma(\psi)} T_\Sigma(Y),$$

with M' and R' generic. But genericness of M' yields the strong lifting γ in

$$
\begin{array}{ccc}
\star & \xrightarrow{\ s^a\ } [a] \xrightarrow{\ R'\ } & T_\Sigma(B) \\
{\scriptstyle M'}\downarrow \quad {\scriptstyle T_\Sigma(\gamma)} \quad \raise2pt\hbox{$\scriptstyle{-}$} & & \downarrow{\scriptstyle T_\Sigma(\psi)} \\
T_\Sigma(A) \xrightarrow{\ T_\Sigma(\varphi)\ } & T_\Sigma(X) \xrightarrow{\ T_\Sigma(f)\ } & T_\Sigma(Y).
\end{array}
$$

9 Cellularity

We have now factored the original square as promised, but for the moment we have no guarantee that the 'inner' square

$$
\begin{array}{ccc}
A & \xrightarrow{\ \varphi\ } & X \\
{\scriptstyle \gamma}\downarrow & & \downarrow{\scriptstyle f} \\
B & \xrightarrow{\ \psi\ } & Y
\end{array}
\tag{7}
$$

will admit a lifting. The point of cellularity is precisely this. For once, let us start from the abstract viewpoint and explain how directly relevant it is in this case.

The starting point is the observation that our definition of bisimulation by lifting is based on a Galois connection. Indeed, for any class $\mathscr{L}$ of morphisms, let $\mathscr{L}^{\boxtimes}$ denote the class of maps $f : X \to Y$ such that for any $l : A \to B$ in $\mathscr{L}$, any commuting square as below left admits a (not necessarily unique) lifting.

$$
\begin{array}{ccc}
A & \xrightarrow{\ u\ } & X \\
{\scriptstyle \mathscr{L}\ni l}\downarrow & & \downarrow{\scriptstyle f\in\mathscr{L}^{\boxtimes}} \\
B & \xrightarrow{\ v\ } & Y
\end{array}
\qquad\qquad
\begin{array}{ccc}
X & \xrightarrow{\ u\ } & A \\
{\scriptstyle {}^{\boxtimes}\mathscr{R}\ni f}\downarrow & & \downarrow{\scriptstyle r\in\mathscr{R}} \\
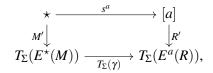Y & \xrightarrow{\ v\ } & B
\end{array}
$$

Conversely, given a class $\mathscr{R}$ of morphisms, let $^{\boxtimes}\mathscr{R}$ denote the class of morphisms $f : X \to Y$ such that for any $r : A \to B$ in $\mathscr{R}$, any commuting square as above right admits a lifting. Clearly, letting $\mathscr{S}$ denote the set of all maps of the form $s^a : \star \to [a]$, $\mathscr{S}^{\boxtimes}$ catches exactly all functional bisimulations. But what is $^{\boxtimes}(\mathscr{S}^{\boxtimes})$? In other words, which maps will admit a lifting against all functional bisimulations? This is very relevant to us, because finding a lifting for our inner square (7) is obviously equivalent to showing that $\gamma \in {}^{\boxtimes}(\mathscr{S}^{\boxtimes})$! Fortunately, the theory of weak factorisation systems gives a precise characterisation [13, Corollary 2.1.15], of which we only need the following very special cases:

Lemma 9.1. *Maps in $^{\boxtimes}(\mathscr{S}^{\boxtimes})$ are closed under composition and pushout, in the sense that*

- *for any composable $f, g \in {}^{\boxtimes}(\mathscr{S}^{\boxtimes})$, $g \circ f \in {}^{\boxtimes}(\mathscr{S}^{\boxtimes})$, and*
- *for any $f : X \to Y$ in $^{\boxtimes}(\mathscr{S}^{\boxtimes})$ and $u : X \to X'$, the pushout f' of f along u, as below, is again in $^{\boxtimes}(\mathscr{S}^{\boxtimes})$.*

$$
\begin{array}{ccc}
X & \xrightarrow{\ f\in{}^{\boxtimes}(\mathscr{S}^{\boxtimes})\ } & Y \\
{\scriptstyle u}\downarrow & & \downarrow{\scriptstyle u'} \\
X' & \xrightarrow{\ f'\in{}^{\boxtimes}(\mathscr{S}^{\boxtimes})\ } & Y'
\end{array}
$$

This is useful to us because the map γ that we want to show is in $^{\boxtimes}(\mathscr{S}^{\boxtimes})$ may be obtained as a finite composite of pushouts of maps in $\mathscr{S}$, which allows us to conclude. Indeed, γ occurs in

$$\begin{array}{ccc}
\star & \xrightarrow{\;\;s^a\;\;} & [a] \\
{\scriptstyle M'}\downarrow & & \downarrow{\scriptstyle R'} \\
T_\Sigma(E^\star(M)) & \xrightarrow[\;T_\Sigma(\gamma)\;]{} & T_\Sigma(E^a(R)),
\end{array}$$

with M' and R' generic. So (M', γ) is the generic-free factorisation of $R' \circ s^a$ as in Lemma 8.2, hence, because generic-free factorisations are unique up to canonical isomorphism, we can actually compute γ. Indeed, letting $M' = (M'', id)$ and $R' = (R'', id)$, for suitable $M'' \in T_\Sigma(1)(\star)$ and $R'' \in T_\Sigma(1)[a]$, by (4) $R' \circ s^a$ is the pair $(R'' \cdot s^a, E(s^a \upharpoonright R''))$, where

$$E(s^a \upharpoonright R'') : E^\star(R'' \cdot s^a) \to E^a(R'')$$

is obtained by familiality of T_Σ. We thus get

$$(M'', \gamma) = (R'' \cdot s^a, E(s^a \upharpoonright R'')),$$

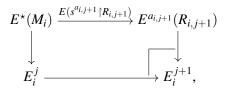hence in particular

$$\gamma = E(s^a \upharpoonright R'').$$

It is thus sufficient to show that each $E(s^a \upharpoonright R)$ is in $^\boxempty(\mathscr{S}^\boxempty)$. This goes by induction on R, following an incremental construction of $E(s^a \upharpoonright R)$. The base case is clear. When $R = \rho\,(R_{i,j})_{i \in n, j \in m_i}$, remember from the proof of Lemma 7.4 that $E^a(R)$ is the coproduct $\sum_i E_i$ for $i \in n$, each E_i being constructed as the wide pushout of all

$$E(s^{a_{i,j}} \upharpoonright R_{i,j}) : E^\star(M_i) \to E^{a_{i,j}}(R_{i,j}).$$

Coproducts may be constructed by pushout along 0, so it suffices to show that each diagonal $E^\star(M_i) \to E_i$ is in $^\boxempty(\mathscr{S}^\boxempty)$ if each $E(s^{a_{i,j}} \upharpoonright R_{i,j})$ is. This in turn boils down to incrementally constructing the diagonal $E^\star(M_i) \to E_i$ by successively pushing out each $E(s^{a_{i,j}} \upharpoonright R_{i,j})$: assuming that we have constructed the diagonal $E^\star(M_i) \to E_i^j$ up until $j < m_i$, we can incorporate $R_{i,j+1}$ by composing with the bottom morphism of

$$\begin{array}{ccc}
E^\star(M_i) & \xrightarrow{E(s^{a_{i,j+1}} \upharpoonright R_{i,j+1})} & E^{a_{i,j+1}}(R_{i,j+1}) \\
\downarrow & & \downarrow \\
E_i^j & \xrightarrow{\hspace{3cm}} & E_i^{j+1},
\end{array}$$

which is indeed in $^\boxempty(\mathscr{S}^\boxempty)$ by Lemma 9.1. Clearly, the obtained $E_i^{m_i}$ is canonically isomorphic to E_i, so we have shown:

Lemma 9.2. *The monad T_Σ is* cellular, *in the sense that in any commuting square of the form*

$$\begin{array}{ccc}
\star & \xrightarrow{\;\;s^a\;\;} & [a] \\
{\scriptstyle \xi}\downarrow & & \downarrow{\scriptstyle \chi} \\
T_\Sigma(A) & \xrightarrow[\;T_\Sigma(\gamma)\;]{} & T_\Sigma(B),
\end{array}$$

with ξ and χ generic, we have $\gamma \in {}^\boxempty(\mathscr{S}^\boxempty)$.

10   Familiality for monads

We have now almost proved:

Lemma 10.1. T_Σ *preserves functional bisimulations.*

The only remaining bit is the hole we left in the proof of Lemma 5.1, when we claimed that all naturality squares of μ were pullbacks. Let us prove this now, as part of the following upgrade of Definition 7.3 and Lemma 7.4.

Definition 10.2. A monad is *familial* when its underlying functor is, and its unit and multiplication are *cartesian* natural transformations, i.e., their naturality squares are pullbacks.

In a case like ours, where the underlying category has a terminal object, by the pullback lemma, it is sufficient to verify that squares of the following form are pullbacks.

$$
\begin{array}{ccc}
T_\Sigma^2(X) & \xrightarrow{\ T_\Sigma^2(!)\ } & T_\Sigma^2(1) \\
{\scriptstyle \mu_X}\downarrow\ \ & & \ \ \downarrow{\scriptstyle \mu_1} \\
T_\Sigma(X) & \xrightarrow[\ T_\Sigma(!)\ ]{} & T_\Sigma(1)
\end{array}
\qquad\qquad
\begin{array}{ccc}
X & \xrightarrow{\ T_\Sigma^2(!)\ } & 1 \\
{\scriptstyle \eta_X}\downarrow\ \ & & \ \ \downarrow{\scriptstyle \eta_1} \\
T_\Sigma(X) & \xrightarrow[\ T_\Sigma(!)\ ]{} & T_\Sigma(1)
\end{array}
$$

The following will conclude our proof of congruence of bisimilarity:

Lemma 10.3. T_Σ *is a familial monad.*

Proof. Pullbacks in presheaf categories being pointwise, we just need to check that a few types of squares are pullbacks in **Set**: for μ and η, and for each type of label. Let us treat the most interesting one, namely the left one below, assuming that the right one has already been covered.

$$
\begin{array}{ccc}
T^2(X)[a] & \xrightarrow{\ T^2(!)_{[a]}\ } & T^2(1)[a] \\
{\scriptstyle \mu_{X,[a]}}\downarrow\ \ & & \ \ \downarrow{\scriptstyle \mu_{1,[a]}} \\
T(X)[a] & \xrightarrow[\ T(!)_{[a]}\ ]{} & T(1)[a]
\end{array}
\qquad\qquad
\begin{array}{ccc}
T^2(X)(\star) & \xrightarrow{\ T^2(!)_\star\ } & T^2(1)(\star) \\
{\scriptstyle \mu_{X,\star}}\downarrow\ \ & & \ \ \downarrow{\scriptstyle \mu_{1,\star}} \\
T(X)(\star) & \xrightarrow[\ T(!)_\star\ ]{} & T(1)(\star)
\end{array}
$$

Let $R[!]$ denote $T(!)(R)$ and $\mathbf{R}[\![!]\!]$ denote $T^2(!)(\mathbf{R})$, for all $R \in T(X)[a]$ and $\mathbf{R} \in T^2(X)[a]$. We must show that for any $a \in \mathbb{A}$, given any $\mathbf{R} \in T^2(1)[a]$ and $R \in T(X)[a]$ such that $R[!] = \mu_1(\mathbf{R})$, there exists a unique $\mathbf{R}^0 \in T^2(X)[a]$ satisfying

$$
\mu_X(\mathbf{R}^0) = R \qquad\qquad\text{and}\qquad\qquad \mathbf{R}^0[\![!]\!] = \mathbf{R}.
$$

We proceed by induction on $\mathbf{R}$. The base case is easy. For the induction step, if $\mathbf{R} = \rho(\mathbf{R}_{i,j})_{i\in n, j\in m_i}$, then because $\mu_1(\mathbf{R}) = R[!]$, we have $R = \rho(R_{i,j})_{i\in n, j\in m_i}$ with $R_{i,j}[!] = \mu_1(\mathbf{R}_{i,j})$ for all i, j[1]. By induction hypothesis, we find a family $\mathbf{R}_{i,j}^0 \in T^2(X)[a_{i,j}]$ such that

$$
\mu_X(\mathbf{R}_{i,j}^0) = R_{i,j} \qquad\qquad\text{and}\qquad\qquad \mathbf{R}_{i,j}^0[\![!]\!] = \mathbf{R}_{i,j} \qquad\qquad\text{for all } i, j.
$$

Letting now $\mathbf{R}^0 = \rho(\mathbf{R}_{i,j}^0)_{i\in n, j\in m_i}$, we get as desired

$$
\mu_X(\mathbf{R}^0) = \rho(\mu_X(\mathbf{R}_{i,j}^0))_{i,j} = \rho(R_{i,j})_{i,j} = R \text{ and } \mathbf{R}^0[\![!]\!] = \rho(\mathbf{R}_{i,j}^0[\![!]\!])_{i,j} = \rho(\mathbf{R}_{i,j})_{i,j} = \mathbf{R}. \qquad\qquad \square
$$

This ends the proof of:

Theorem 10.4. *For all $X \in \widehat{\Gamma_\mathbb{A}}$ and Positive GSOS specifications Σ, bisimilarity in the free algebra $T_\Sigma(X)$ is a congruence.*

---

[1] For all i such that $m_i = 0$, we in fact deal with some term M_i, using the corresponding square. Let us ignore this detail for readability.

11   Conclusion and perspectives

In this paper, we have introduced the familial approach to programming language theory [12] at the rather concrete level of generalised labelled transition systems ($\widehat{\Gamma_{\mathbb{A}}}$). Notably, we have recalled the notions of cellular monad and compositional algebra, and recalled that bisimilarity is always a congruence in a compositional algebra for a cellular monad (Lemma 4.3).

We have also shown that all monads T_Σ generated from a Positive GSOS specification Σ are cellular (Lemma 9.2) and that free T_Σ-algebras are always compositional (Lemma 5.1). Putting all three results together, we readily recover (Theorem 10.4) the known result that bisimilarity is a congruence for all free T_Σ-algebras. In particular, this is the case for the standard, syntactic transition system, which is the initial algebra $T_\Sigma(0)$.

This result constitutes a first generic tool for constructing instances of the framework of [12]. However, its scope is rather limited, and we plan to refine the construction to cover other formats like *tyft/tyxt* [11]. A striking and promising observation here is that the well-foundedness condition demanded of a tyft/tyxt specification for bisimilarity to be a congruence is clearly covered by our approach based on weak factorisation systems (see §9). Cellularity thus provides a semantic criterion for well-foundedness, whose precise relationship with the original, syntactic one seems worth investigating.

Beyond the task of showing by hand that existing formats yield cellular monads whose free algebras are compositional, we also plan to investigate a more categorical understanding of the generating process. The main motivation here is to design a construction that would cover variable binding. The theory developed by Fiore and his colleagues [7, 5] seems like a good starting point.

References

[1] B. Bloom, S. Istrail & A. Meyer (1995): *Bisimulation can't be traced*. Journal of the ACM 42, pp. 232–268, doi:10.1145/200836.200876.

[2] Filippo Bonchi, Daniela Petrisan, Damien Pous & Jurriaan Rot (2014): *Coinduction up-to in a fibrational setting*. In: *Proc. 29th Symposium on Logic in Computer Science*, ACM, pp. 20:1–20:9, doi:10.1145/2603088.2603149.

[3] Aurelio Carboni & Peter Johnstone (1995): *Connected Limits, Familial Representability and Artin Glueing*. Mathematical Structures in Computer Science 5(4), pp. 441–459, doi:10.1017/S0960129500001183.

[4] Andrea Corradini, Reiko Heckel & Ugo Montanari (2002): *Compositional SOS and beyond: a coalgebraic view of open systems*. Theoretical Computer Science 280(1-2), pp. 163–192, doi:10.1016/S0304-3975(01)00025-1.

[5] Marcelo Fiore & Chung-Kil Hur (2009): *On the construction of free algebras for equational systems*. Theoretical Computer Science 410, pp. 1704–1729, doi:10.1016/j.tcs.2008.12.052.

[6] Marcelo P. Fiore (2000): *Fibred Models of Processes: Discrete, Continuous, and Hybrid Systems*. In: *IFIP TCS*, *LNCS* 1872, Springer, pp. 457–473, doi:10.1007/3-540-44929-9_32.

[7] Marcelo P. Fiore (2008): *Second-Order and Dependently-Sorted Abstract Syntax*. In: *LICS*, IEEE, pp. 57–68, doi:10.1109/LICS.2008.38.

[8] Marcelo P. Fiore & Sam Staton (2006): *A Congruence Rule Format for Name-Passing Process Calculi from Mathematical Structural Operational Semantics*. In: *Proc. 21st Symposium on Logic in Computer Science*, IEEE, pp. 49–58, doi:10.1109/LICS.2006.7.

[9] Marcelo P. Fiore & Daniele Turi (2001): *Semantics of Name and Value Passing*. In: *Proc. 16th Symposium on Logic in Computer Science*, IEEE, pp. 93–104, doi:10.1109/LICS.2001.932486.

[10] Richard H. G. Garner & Tom Hirschowitz (2018): *Shapely monads and analytic functors*. Journal of Logic and Computation 28(1), pp. 33–83, doi:10.1093/logcom/exx029.

[11] Jan Friso Groote & Frits Vaandrager (1992): *Structured Operational Semantics and Bisimulation as a Congruence*. Information and Computation 100, pp. 202–260, doi:10.1016/0890-5401(92)90013-6.

[12] Tom Hirschowitz (2019): *Familial monads and structural operational semantics*. PACMPL 3(POPL), pp. 21:1–21:28, doi:10.1145/3290334.

[13] Mark Hovey (1999): *Model Categories. Mathematical Surveys and Monographs, Volume 63, AMS (1999)* 63, American Mathematical Society, doi:10.1090/surv/063.

[14] André Joyal, Mogens Nielsen & Glynn Winskel (1993): *Bisimulation and open maps*. In: *Proc. 8th Symposium on Logic in Computer Science*, IEEE, pp. 418–427, doi:10.1109/LICS.1993.287566.

[15] Tom Leinster (2014): *Basic Category Theory. Cambridge Studies in Advanced Mathematics* 143, Cambridge University Press, doi:10.1017/CBO9781107360068.

[16] Saunders Mac Lane (1998): *Categories for the Working Mathematician*, 2nd edition. *Graduate Texts in Mathematics* 5, Springer, doi:10.1007/978-1-4757-4721-8.

[17] Saunders Mac Lane & Ieke Moerdijk (1992): *Sheaves in Geometry and Logic: A First Introduction to Topos Theory*. Universitext, Springer, doi:10.1007/978-1-4612-0927-0.

[18] Robin Milner (1980): *A Calculus of Communicating Systems. LNCS* 92, Springer, doi:10.1007/3-540-10235-3.

[19] MohammadReza Mousavi, Michel A. Reniers & Jan Friso Groote (2007): *SOS Formats and Meta-Theory: 20 Years After*. Theoretical Computer Science 373(3), pp. 238–272, doi:10.1016/j.tcs.2006.12.019.

[20] Marco Peressotti (2017): *Coalgebraic Semantics of Self-Referential Behaviours*. Ph.D. thesis, University of Udine, doi:10.13140/rg.2.2.26899.07203.

[21] Gordon D. Plotkin (1981): *A Structural Approach to Operational Semantics*. DAIMI Report FN-19, Computer Science Department, Aarhus University.

[22] Emily Riehl (2014): *Categorical Homotopy Theory. New Mathematical Monographs* 24, Cambridge University Press, doi:10.1017/CBO9781107261457.

[23] Davide Sangiorgi & Jan Rutten, editors (2011): *Advanced Topics in Bisimulation and Coinduction. Cambridge Tracts in Theoretical Computer Science* 52, Cambridge University Press, doi:10.1017/CBO9780511792588.

[24] Sam Staton (2008): *General Structural Operational Semantics through Categorical Logic*. In: *Proc. 23rd Symposium on Logic in Computer Science*, pp. 166–177, doi:10.1109/LICS.2008.43.

[25] Daniele Turi & Gordon D. Plotkin (1997): *Towards a Mathematical Operational Semantics*. In: *Proc. 12th Symposium on Logic in Computer Science*, pp. 280–291, doi:10.1109/LICS.1997.614955.

[26] Mark Weber (2007): *Familial 2-functors and parametric right adjoints*. Theory and Applications of Categories 18(22), pp. 665–732.