

A coalgebraic semantics for causality in Petri nets

Roberto Bruni^a, Ugo Montanari^a, Matteo Sammartino^{b,*}

^a*University of Pisa, Computer Science Department, Largo Bruno Pontecorvo 3, 56127 Pisa, Italy*

^b*Radboud University, Institute for Computing and Information Sciences, Faculty of Science, Heyendaalseweg 135, 6525AJ Nijmegen, The Netherlands*

Abstract

In this paper we revisit some pioneering efforts to equip Petri nets with compact operational models for expressing causality. The models we propose have a bisimilarity relation and a minimal representative for each equivalence class, and they can be fully explained as coalgebras on a presheaf category on an index category of partial orders. First, we provide a set-theoretic model in the form of a *causal case graph*, that is a labelled transition system where states and transitions represent markings and firings of the net, respectively, and are equipped with causal information. Most importantly, each state has a poset representing causal dependencies among past events. Our first result shows the correspondence with behavior structure semantics as proposed by Trakhtenbrot and Rabinovich. Causal case graphs may be infinitely-branching and have infinitely many states, but we show how they can be refined to get an equivalent finitely-branching model. In it, states only keep the most recent causes for each token, are up to isomorphism, and are equipped with a symmetry, i.e., a group of poset isomorphisms. Symmetries are essential for the existence of a minimal, often finite-state, model. This first part requires no knowledge of category theory. The next step is constructing a coalgebraic model. We exploit the fact that events can be represented as names, and event generation as name generation. Thus we can apply the Fiore-Turi framework, where the semantics of nominal calculi are modelled as coalgebras over presheaves. We model causal relations as a suitable category of posets with action labels, and generation of new events with causal dependencies as an endofunctor on this category. Presheaves indexed by labelled posets represent the functorial association between states and their causal information. Then we define a well-behaved category of coalgebras. Our coalgebraic model is still infinite-state, but we exploit the equivalence between coalgebras over a class of presheaves and *History Dependent automata* to derive a compact representation, which is equivalent to our set-theoretical compact model. Remarkably, state reduction is automatically performed along the equivalence.

Keywords: Petri nets, causal case graphs, causal semantics, behavior structures, presheaves, coalgebras, HD-automata.

[☆]Research supported by the EU Integrated Project 257414 ASCENS and by the Italian MIUR Project CINA (PRIN 2010).

*Corresponding author. Tel.: (+31) 0243652642

Email addresses: bruni@di.unipi.it (Roberto Bruni), ugo@di.unipi.it (Ugo Montanari), m.sammartino@cs.ru.nl (Matteo Sammartino)

URL: <http://www.di.unipi.it/~bruni> (Roberto Bruni), <http://www.di.unipi.it/~ugo> (Ugo Montanari), <http://www.cs.ru.nl/M.Sammartino/> (Matteo Sammartino)

1. Introduction

Petri Nets are a well-known graphical and formal notation for representing concurrent computations. An interesting aspect of Petri Nets is that they allow for the representation of causal dependencies among actions. This kind of information can be useful for debugging distributed systems or for tracing expected or unwanted causal dependencies, and it is usually not provided by interleaving models.

In order to carry out verification on Petri nets, it is convenient to have an *operational* model, that is a model representing single steps of computation and their observable actions. In Petri nets, steps are typically firings and actions are action labels of transitions. One important class of operational models for Petri Nets are *behavior structures* [27]. They are automata where each state is equipped with a partial order over events: events represent different occurrences of actions and the poset describes causal dependencies among such occurrences. Behavior structures come with a notion of behavioral equivalence, which later has been called *history preserving bisimilarity* [14].

Other causal models, such as *event structures* [20], do not come with a built-in operational notion of bisimilarity. Such a notion is essential to compute minimal models, where all states with the same behavior are identified. Open maps [16] can be used to derive *hereditary history preserving bisimulations* (HHPBs), but the existence of minimal representatives is not guaranteed by that theory. Indeed, the general agreement is that HHPB is more suited to capture concurrency, whereas the non-hereditary version deals better with causality. The latter equivalence is coarser, but still causality is informative enough to characterize key security properties, such as non-interference [4]. Moreover, the non-hereditary equivalence has better decidability properties than the hereditary one [14].

The main issue with causal operational models is that they often have infinitely many states, so model checking is unfeasible. This is indeed the case of behavior structures, where posets of states are enlarged at each transition, because a new event for the corresponding action is generated. Even if we minimize w.r.t. bisimilarity, there is no way of throwing away “useless” events or decreasing the size of posets.

In this paper we present an approach to obtain compact, and in many cases finite, operational models for causality in Petri nets. They will be presented in two “flavors”: a set-theoretic and a categorical one, based on coalgebras [22, 1]. In addition to the theoretical and practical interest of reconducting our problem to unifying and well studied models such as coalgebras, we emphasize that our coalgebraic model is simpler than the set theoretical one. In fact, even if deriving a naive set-theoretic model from a Petri net is not difficult, the technical development required to obtain a compact model is quite involved and requires some ingenuity. Instead, in a categorical setting, this machinery will become remarkably simpler and natural. Actually, in a precise sense, the construction of the compact model will be automatic, thus providing a mathematical justification of the otherwise ad hoc set-theoretic constructions.

1.1. Set-theoretic models

After some preliminaries on Petri nets and the presentation of a running example in section 2, in section 3 we model the behavior of a labelled Petri net as a *causal case graph* (CG). Recall that a case graph is a labelled transition graph where states are markings and transitions are steps, representing many firings happening simultaneously. In causal case graphs, transitions are single firings, and causal data are used to encode information about concurrency. More precisely (see Definition 3.3, where CGs are called “concrete” as opposed to “abstract” CGs, introduced later):

- states are of the form $O \triangleright c$, where: O is a poset describing causal dependencies among a finite collection of events; c is a marking where each token is decorated with its *causes*, i.e. the set of events that led to its creation (included in O);
- the transition relation is written $\xrightarrow{K \vdash e_a}$, where: K is the set of most recent causes of tokens that enabled the firing; e is a *fresh* event, different from all those occurring in the source state; and a is the action label of the fired transition.

We define a notion of bisimilarity for CGs where causal information plays a key role: only states with the same causal dependencies among past events, namely the same poset, are compared. This fact is crucial for the equivalence with history preserving bisimilarity described in section 4.

Another important aspect is that transitions draw fresh events from an infinite set of event names. For each firing, we have *infinitely many* transitions in the CG, one for each possible fresh event. In this way we implement *event generation* in the same way name generation is represented, e.g., in nominal calculi. This fact will be crucial for our categorical models.

We, then, derive three consecutive refinements of the CG, described in Table 1, each improving the CG on one aspect:

CG_{AC} (Definition 3.8): the transition relation becomes *finitely branching*, because we don't distinguish between posets with the same structure. In fact, it is enough to generate one canonical event, instead of all possible ones, for each firing. Consequently, states contain canonical representatives of events and only the action label of the new event is recorded in the transition.

CG_{IC} (Definition 3.17): removing all but immediate causes, and identifying isomorphic states, may significantly reduce the state space, and even make it finite.

CG_{ICS} (Definition 3.27): we equip each state with a set of isomorphisms acting as the identity on the state. These isomorphisms must form a *symmetry*, i.e., a group of automorphisms, on the state's poset. Transitions are reduced accordingly: we select one representative for each collection of "symmetric" transitions. Two transitions are symmetric whenever they can be obtained from each other via isomorphisms belonging to the symmetries of source and target states. Symmetries allow for the computation of minimal models, because CGs that are not isomorphic, but bisimilar under a given isomorphism, have a unique minimal realization, where that isomorphism becomes part of the symmetry of a state.

These steps do not change the overall semantics (Theorems 3.12 and 3.22).

Finally, in Theorem 4.6 we establish a connection between CGs and behavior structures.

1.2. Categorical models

In the second part of the paper (Sections 5-7) we assume the reader has some familiarity with category theory. Some preliminaries about presheaves and coalgebras are recalled in section 5.

Coalgebras are convenient models of dynamic systems. Their theory is rich and well-developed, and many kinds of systems have been characterized in this setting. Coalgebras are also of practical interest: minimization procedures such as *partition refinement* [17] can be defined in coalgebraic terms (see, e.g., [2]). This further motivates the coalgebraic framework: algorithms implemented at this level of abstraction can be instantiated to many classes of systems.

Our coalgebraic causal model of Petri nets, presented in section 6, is based on the fact that we represent events as names and event generation as name generation, in the style of nominal calculi.

States	Transition relation
Causal case graph (CG)	
$O \triangleright c$	$\xrightarrow{K \vdash e_a}$
<ul style="list-style-type: none"> • O is a finite poset describing causal dependencies among events • c is a marking including causes for each token 	<ul style="list-style-type: none"> • K is the set of most recent causes of tokens consumed by the transition • e is a fresh event • a is the fired transition's action label
Abstract CG (CG_{AC})	
$O \triangleright c$	$\xRightarrow{K \vdash a}$
<ul style="list-style-type: none"> • O is a <i>canonical representative</i> of isomorphic posets • c contains canonical events 	<ul style="list-style-type: none"> • K as in CG • a is the action label for the canonical fresh event
Immediate causes CG (CG_{IC})	
$O \blacktriangleright c$	$\xrightarrow[h]{K \vdash a}$
<ul style="list-style-type: none"> • O and c contain only the most recent causes w.r.t. each token (<i>immediate causes</i>) • each state is a canonical representative of isomorphic states 	<ul style="list-style-type: none"> • K and a as in CG_{AC} • h is a map telling how events in the target state correspond to those of the source state
Immediate causes CG with symmetries (CG_{ICS})	
$O \blacktriangleright_{\Phi} c$	$\xRightarrow[h]{K \vdash a}$
<ul style="list-style-type: none"> • O and c as in CG_{IC} • Φ is a symmetry on O 	<ul style="list-style-type: none"> • K, a and h as in CG_{IC} • transitions are canonical representatives of “symmetric” ones

Table 1: Set-theoretic models.

This allows us to construct a coalgebra where states are equipped with nominal structures, namely causal relations between events, and event generation is explicit, along the lines of [13]. The key idea is to define coalgebras over *presheaves*, that are functors from a certain *index category* \mathbf{C} to \mathbf{Set} , the category of sets and functions. Presheaves formalize the association between a collection of names, seen as an objects of \mathbf{C} , and a set of processes within \mathbf{Set} , indexed by names of the collection. Fresh name generation can be formalized as an endofunctor on \mathbf{C} , that is lifted to presheaves and used in the definition of coalgebras.

We take as index category for presheaves a suitable category of *labelled posets* up to isomorphism, representing causal relations between events decorated with actions. This category provides us with the needed structure to model operations over causal relations. In fact, we use colimits to implement a *well-behaved functorial* model of event generation, which augments a given poset with fresh events and relations to their causes. Our definition ensures that its lifting to presheaves, when used to define coalgebras, yields a category of coalgebras with a final object and a final semantics in agreement with coalgebraic bisimilarity. This is essential for a correct notion of minimal model. Then, we define a presheaf of causal markings, yielding, for each poset, the set of causal markings whose causes are “compatible” with that poset. We construct a *causal coalgebra* by translating the abstract CG. The important result is that coalgebraic and ordinary bisimilarity are equivalent (Theorem 6.16).

The infinite state issue still exists in the causal coalgebra, because the poset of a causal marking keeps growing along transitions. However, if the presheaf of states is “well-behaved”, according to [10], it is always possible to recover the *support* of a causal marking, that is the minimal poset including all and only events that appear in the marking. This is the key condition for the equivalence between presheaf-based coalgebras and *History Dependent (HD) automata* [21].

HD-automata are coalgebras with states in *named-sets* [11], that are sets whose elements are equipped with symmetry groups over finite collections of names. They have two main features:

- a single state can represent the whole orbit of its symmetry, namely all the states reachable via poset isomorphisms;
- the names of each state are *local*, related to those of other states via suitable mappings.

Both features are important for applying finite state methods, such as minimization and model-checking, to nominal calculi. In particular, the latter point captures *deallocation*: maps between states can discard unused names and “compact” remaining ones, much like *garbage collectors* do for memory locations. A minimization procedure for HD-automata for the (finite-control) π -calculus has been shown and implemented in [12].

Interestingly, we are able to define the presheaf of causal markings in a way that computing the support corresponds to discarding all but the immediate causes. Therefore, in section 7 we show that the aforementioned equivalence amounts to deriving the immediate causes CG. Actually, it also equips states with symmetries, achieving the last refinement step. We emphasize that such equivalence is completely standard in the theory of nominal calculi. In our case, it is extended to labelled posets and allows the automatic derivation of an HD-automaton over a named set of minimal causal markings.

2. Basic definitions and running example

Given a set of labels L , we call *L-labeled poset* (or just labeled poset, when L is clear from the context) on a set S a triple $O = (X_O, \leq_O, l_O)$, where $X_O \subseteq S$, \leq_O is a reflexive, transitive

and antisymmetric relation on X_O and $l_O: X_O \rightarrow L$ is a labeling function. A morphism of labeled posets $O \rightarrow O'$ is a function $\sigma: X_O \rightarrow X_{O'}$ that preserves order and labeling, namely $x \preceq_O y$ implies $\sigma(x) \preceq_{O'} \sigma(y)$ and $l_O = l_{O'} \circ \sigma$. We say that σ *reflects order* whenever $\sigma(x) \preceq_{O'} \sigma(y)$ implies $x \preceq_O y$; σ is an *order-embedding* whenever it both preserves and reflects order. Notice that isomorphisms reflect order, because their inverses preserve order, and it can be easily checked that order-embeddings are always injective. To simplify notation, we sometimes regard O as a poset on $S \times L$, we write $|O|$ for the underlying set of pairs and $x_l \in X_O \times L$ for the pair $(x, l) \in |O|$. A set $K \subseteq |O|$ is *down-closed* w.r.t. O whenever $y \in K$ and $x \preceq_O y$ implies $x \in K$. We say that a poset O is a *prefix* of O' if O is a subposet of O' and $|O|$ is down-closed w.r.t. O' .

In this paper we consider the following kind of Petri nets, which we call just *nets*.


Definition 2.1 (Net). A net is a tuple (S, T, F, l) where:

- S is a set of *places* and T is a set of *transitions*, with $S \cap T = \emptyset$;
- $F \subseteq (S \times T) \cup (T \times S)$ is the *flow relation*;
- $l: T \rightarrow Act$ is a *labeling function*, where Act is a fixed set of action labels.

If $x \in S \cup T$ then $\bullet x = \{y \mid (y, x) \in F\}$ and $x^\bullet = \{y \mid (x, y) \in F\}$ are called the *pre-set* and *post-set* of x , respectively; for all $t \in T$, we assume $\bullet t, t^\bullet \neq \emptyset$. A *marking* m is a multiset over S . A transition $t \in T$ is *enabled* at marking m if $s \in m$, for all $s \in \bullet t$, in which case it can *fire*, written $m[t]m'$, i.e., a new marking $m' = (m \setminus \bullet t) \cup t^\bullet$ is produced. We say that a net is *marked* whenever it has an *initial marking* m_0 . We denote by $[m_0]$ the set of markings reachable from m_0 by a (finite) sequence of firings.

We require that elements of initial markings have multiplicity one. This implies that m_0 is actually a set, in agreement with the fact that pre-sets and post-sets in nets are sets, meaning that they can only consume one token at a time from a given place. In typical P/T nets transitions may consume many tokens from the same place, but this difference is inessential for the development of our theory.

Running example. As a running example, we will use the marked net defined as follows: $S = \{s_1, s_2\}$, $T = \{t_1, t_2, t_3\}$, F includes $(s_i, t_i), (s_i, t_3)$ (for $i = 1, 2$) and symmetric pairs, and $l(t_1) = l(t_2) = a$, $l(t_3) = b$. The initial marking is $m_0 = \{s_1, s_2\}$. This net is depicted below: circles denote places, squares denote transitions, edges describe the flow relation, and filled circles indicate the position of tokens in m_0 . Notice that $[m_0] = \{m_0\}$.



3. Causal semantics for Petri nets

In this section we introduce our *causal labelled semantics* for nets. It will be in the form of a *causal case graph* (CG in short), that is a labelled transition graph whose states are markings with causal information and transitions represent firings. We start from a naive CG, derived from a given net in the simplest way, and then we give three subsequent refinements that will lead to a compact

and, in some cases, finite-state CG. Throughout this section we fix a net $N = (S, T, F, l)$ and we assume that an infinite set \mathcal{E} of event names (or just events) is available.

The key idea is to equip markings with information about the occurrences of actions that led to the creation of each token. An occurrence of a transition labeled by $a \in Act$ is represented as an Act -labeled event e_a . Formally, a *causal marking* c is a set of the form

$$\{K_1 \vdash s_1, \dots, K_n \vdash s_n\}$$

where $K_i \subseteq \mathcal{P}_f(\mathcal{E} \times Act)$ is the set of *causes* of $s_i \in S$, for $i = 1, \dots, n$. More specifically, if $e_a \in K_i$ then the sequence of firings that generated the token includes a transition with action label a . We write $\mathcal{K}(c)$ for $K_1 \cup \dots \cup K_n$ and $|c|$ for the underlying marking $\{s_1, \dots, s_n\}$ of c . Given a marking m and $K \subseteq \mathcal{P}_f(\mathcal{E} \times Act)$, $K \vdash m$ is the causal marking obtained by assigning causes K to each $s \in m$.

Transitions of our CGs will generate new events and their causal dependencies. In order to keep track of these data, we equip causal markings with Act -labeled posets, describing the causal relations between events which are occurrences of past actions.

Definition 3.1 (P-marking). A P-marking is a pair $O \triangleright c$, where c is a causal marking and O is a finite Act -labeled poset on \mathcal{E} such that: if $K \vdash s \in c$ then K is down-closed w.r.t. O .

Down-closure requires each set of causes to contain the whole “history” of its events, as described by O . Nevertheless, O may contain events that are unrelated to or caused by those of $\mathcal{K}(c)$, but that are not among them.

Posets will have different purposes in the different classes of CGs we are going to introduce: they will be used to record either all the events happened so far or the “most recent” ones. The shape of P-markings will not change, but there will be additional requirements on their components.

We introduce a useful operation on P-markings. Their posets can be enlarged by adding events from which existing events causally depend on, but a closure operator must be applied, in order to retain down-closure of sets of causes.

Definition 3.2 (Closure operator). Given $K \subseteq |O|$ and O' such that O is a subposet of O' , the *closure* of K w.r.t. O' is given by

$$K \downarrow_{O'} = \bigcup_{x \in K} \{y \in |O'| \mid y \preceq_{O'} x\}$$

Its extension to causal markings is $(K \vdash s) \downarrow_{O'} = K \downarrow_{O'} \vdash s$ and acts element-wise on sets.

Given a P-marking $O \triangleright c$ and $O' \supseteq O$, it can be easily verified that $O' \triangleright c \downarrow_{O'}$ is a proper P-marking.

3.1. Concrete CG

The first step is deriving a CG from the net. Its states are P-markings $O \triangleright c$ such that O contains the whole history of past events and transition labels are of the form $K \vdash e_a$, meaning that an a -labelled transition t is fired: e_a is an event fresh w.r.t. all the previous ones (i.e., those in O) and K is the set of most recent causes associated to tokens that enabled t . We call this CG *concrete* because posets with the same structure but different event names are distinguished.

Definition 3.3 (Concrete CG). The *concrete CG* ($\text{CG}_{\mathcal{C}}$) is the smallest CG generated by the following rule

$$\frac{t \in T \quad |c| = \bullet t \quad a = l(t) \quad e \in \mathcal{E} \setminus X_O \quad K = \max_O \mathcal{K}(c)}{O \triangleright c \cup c' \xrightarrow{K \vdash e_a} \delta(O, K, e_a) \triangleright (\mathcal{K}(c) \cup \{e_a\} \vdash t \bullet) \cup c'}$$

where $\max_O K$, for $K \subseteq |O|$, is the set of maximal elements in K according to O , and $\delta(O, K, x) = (O \cup (K \times \{x\}))^*$.

Given a P-marking, the rule above checks whether it includes a causal marking c such that its underlying marking is the pre-set of a transition t ($|c| = \bullet t$). If this is the case, t is turned into a CG transition whose label $K \vdash e_a$ is formed by the maximal causes K of c w.r.t. O and by a labelled event e_a , where e does not occur in the source poset ($e \notin \mathcal{E} \setminus X_O$). The target state is obtained by replacing c with the tokens produced by the firing, each equipped with the whole set of causes of c plus the new event e_a . Since e_a is causally dependent on the causes of c , the poset in the target state is updated with new pairs representing such dependencies by taking $\delta(O, K, e_a)$.

Note that event generation is similar to name generation in nominal calculi.¹ For instance, in a π -calculus extrusion transition $(y)\bar{x}y.p \xrightarrow{\bar{x}(z)} p[z/y]$ we observe a free name x and a fresh name z , which then becomes free in the continuation. Analogously, in a transition $O \triangleright c \xrightarrow{K \vdash e_a} \delta(O, K, e_a) \triangleright c'$ the elements of K are “free” events, in the sense that they occur in c , and e is a fresh one, which is then added to the continuation. As in the π -calculus, event generation causes $\text{CG}_{\mathcal{C}}$ to have infinitely-many states and to be infinitely-branching, because there are infinitely-many transitions and continuations from any state, differing only for the identity of the fresh event.

Remark 3.4. Even if initial markings are sets, firings may eventually produce a proper multiset, for instance when a transition puts a token in a place s that is already marked. Instead, our causal markings are sets: they can never contain two occurrences of $K \vdash s$, for any K . In fact, suppose the first of the described firings becomes a CG transition that goes to a P-marking including $K \vdash s$. Then, since the second transition fires later, it will generate an event $e_a \notin K$ and a target P-marking that includes both $K \vdash s$ and a new $K' \vdash s$ such that $e_a \in K'$, so $K \neq K'$.

Example 3.5. Figure 1 depicts some transitions of the $\text{CG}_{\mathcal{C}}$ for the running example. It shows only the reachable part from $\emptyset \triangleright \emptyset \vdash m_0$, up to a certain depth. Each state has three kinds of outgoing transitions, corresponding to the three net transitions. The figure only shows one transition for each kind, but there are actually infinitely many ones, one for each fresh event.

We now introduce bisimulations for $\text{CG}_{\mathcal{C}}$.

Definition 3.6 (Concrete causal bisimulation). A *concrete causal bisimulation* (\mathcal{C} -bisimulation in short) is a family of relations $\{R_O\}$ on P-markings, indexed by *Act*-labeled posets, such that:

- whenever $(O_1 \triangleright c_1, O_2 \triangleright c_2) \in R_O$ then $O_1 = O_2 = O$;
- whenever $(O \triangleright c_1, O \triangleright c_2) \in R_O$ and $O \triangleright c_1 \xrightarrow{K \vdash e_a} O' \triangleright c'_1$ then $O \triangleright c_2 \xrightarrow{K \vdash e_a} O' \triangleright c'_2$ and $(O' \triangleright c'_1, O' \triangleright c'_2) \in R_{O'}$ (and viceversa).

The concrete causal bisimilarity is the greatest such family and is denoted by $\sim_{\mathcal{C}}$.

¹The relationship between π -calculus and causality has been investigated in [6].

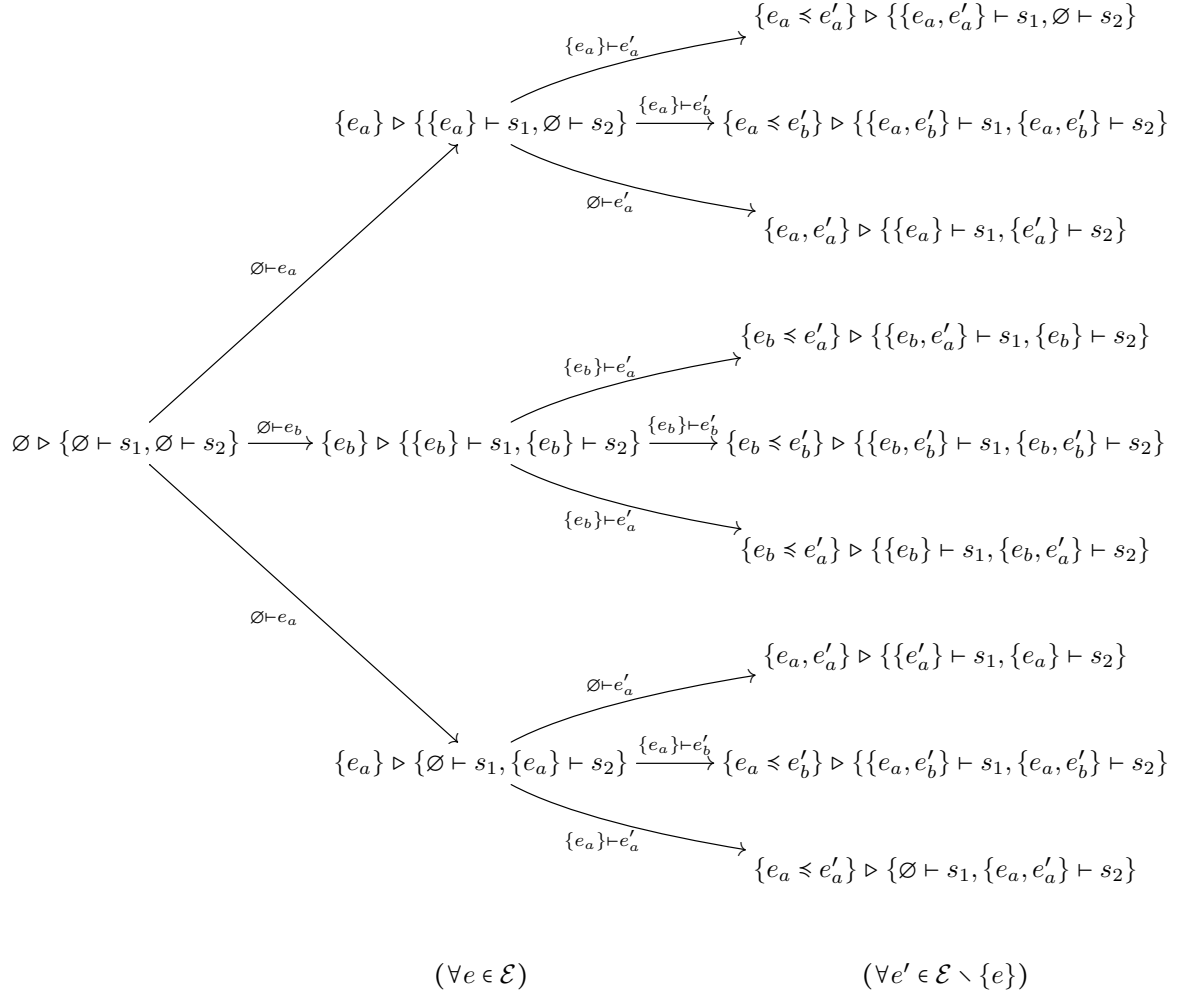


Figure 1: CG_c for the running example.

3.2. Abstract CG

We now introduce an *abstract CG*, where we only take posets up to isomorphism. We write $[O]_{\cong}$ for the isomorphism representative of O , and we call it *abstract poset*. We call *abstract* a P-marking of the form $[O]_{\cong} \triangleright c$.

Given an abstract poset O , $K \subseteq |O|$ and $a \in Act$, we assume the following operations:

- $\delta(O, K, a)$, generating $[\delta(O, K, e_a)]_{\cong}$, for any e_a ; the actual identity of e_a is not relevant, because of the quotient up to isomorphism;
- $new(O, K, a)$, giving the unique new event in $\delta(O, K, a)$;
- the morphism $old(O, K, a)$, embedding O into $\delta(O, K, a)$;

These operations can be used to define the *extension* of $\sigma: O \rightarrow O'$ (with O, O' abstract posets) to a morphism $\sigma_{K,a}^+: \delta(O, K, a) \rightarrow \delta(O', \sigma(K), a)$ given by

$$\sigma_{K,a}^+(x) = \begin{cases} new(O', \sigma(K), a) & x = new(O, K, a) \\ old(O', \sigma(K), a)(\sigma(y)) & x = old(O, K, a)(y) \end{cases}$$

The intuition is that $\sigma_{K,a}^+$ does not mix up old and new events: it acts “as” σ (modulo suitable embeddings) on events that were already in O , and maps the new event in $\delta(O, K, a)$ to the new one in $\delta(O', \sigma(K), a)$. To ease notation, we will just write σ^+ when K and a are clear from the context.

Example 3.7. Suppose $O_1 = \{x_a, x'_b\}$ and $O_2 = \{y_a, y'_b, y''_c\}$ are discrete abstract posets, and let $\sigma: O_1 \rightarrow O_2$ map x_a to y_a and x'_b to y'_b . Let \hat{x}_z (resp. \hat{y}_z) be the image of x_z via $old(O, \{x_a, x'_b\}, d)$ (resp. via $old(O', \{y_a, y'_b\}, d)$), for $z \in \{a, b\}$. Then we have

$$\delta(O_1, \{x_a, x'_b\}, d) = \begin{array}{ccc} & new(O_1, \{x_a, x'_b\}, d) & \\ \nearrow & & \nwarrow \\ \hat{x}_a & & \hat{x}_b \end{array} \quad \delta(O_2, \{y_a, y'_b\}, d) = \begin{array}{ccc} & new(O_2, \{y_a, y'_b\}, d) & \\ \nearrow & \uparrow & \\ \hat{y}_a & \hat{y}'_b & \hat{y}''_c \end{array}$$

where arrows represent ordered pairs (reflexive pairs are omitted). Then $\sigma^+: \delta(O_1, \{x_a, x'_b\}, d) \rightarrow \delta(O_2, \{y_a, y'_b\}, d)$ maps \hat{x}_a to \hat{y}_a , \hat{x}'_b to \hat{y}'_b and $new(O, \{x_a, x'_b\}, d)$ to $new(O_2, \{y_a, y'_b\}, d)$.

We now introduce the *abstract CG*. Its states are abstract P-markings and its labels have the form $K \vdash a$. Labels have the same meaning as in CG_C , but here there is no need to observe the generated event: it will always be $new(O, K, a)$, if O is the source P-marking’s poset.

In order to translate concrete P-markings, and their transitions, to their abstract counterparts in CG_{AC} , we fix an *abstraction isomorphism* $\alpha_O: O \rightarrow [O]_{\cong}$, for each poset O , giving a canonical representative of each event in O . In the following we write $\|x\|_O$ for the “abstract version” of x , namely $x\alpha_O$. We also introduce an operation $\|c\|_{O, K, e_a}$. It will be applied to causal markings c appearing in continuations of transitions of CG_C , namely those P-markings of the form $\delta(O, K, e_a) \triangleright c$. Intuitively, given a transition in CG_C , the operation $\|-\|_{O, K, e_a}$ applies the abstraction isomorphism of the source P-marking to its continuation, so that events of source and continuation are consistent with each other and the fresh event generated by the transition always becomes the canonical new one. Formally, $\|c\|_{O, K, e_a}$ is defined as follows: events in O are mapped via α_O and then embedded into $[\delta(O, K, e_a)]_{\cong}$ via $old([O]_{\cong}, \|K\|_O, a)$ (notice that $[\delta(O, K, e_a)]_{\cong} = \delta([O]_{\cong}, \|K\|_O, a)$, because they are isomorphic); and e_a is embedded into $[\delta(O, K, e_a)]_{\cong}$ as $new([O]_{\cong}, \|K\|_O, a)$.

Definition 3.8 (abstract CG). The *abstract CG* (CG_{AC}) is the smallest CG generated by the following rule

$$\frac{O \triangleright c \xrightarrow{K \vdash e_a} \delta(O, K, e_a) \triangleright c'}{[O]_{\cong} \triangleright \|c\|_O \xrightarrow{\|K\|_{O \vdash a}} \delta([O]_{\cong}, \|K\|_O, a) \triangleright \|c'\|_{O, K, e_a}}$$

The most important fact to notice is that CG_{AC} is finitely branching. In fact, even if there are infinitely-many concrete P-markings that generate the transitions of an abstract P-marking $O \triangleright c$, they are all isomorphic. To see this, take any two P-markings $O_1 \triangleright c_1$ and $O_2 \triangleright c_2$ such that $\|c_1\|_{O_1} = \|c_2\|_{O_2} = c$. Then we have $c = c_1 \alpha_{O_1}^{-1} = c_2 \alpha_{O_2}^{-1}$, so $c_2 = c_1 \sigma$, where σ is the isomorphism $\alpha_{O_2}^{-1} \circ \alpha_{O_1}$. The following lemma states the correspondence between transitions of such P-markings.

Lemma 3.9. *Let $\sigma: O_1 \rightarrow O_2$ be an isomorphism. Then $O_1 \triangleright c_1 \xrightarrow{K \vdash e_a} \delta(O_1, K, e_a) \triangleright c'_1$ if and only if $O_2 \triangleright c_1 \sigma \xrightarrow{\sigma(K) \vdash e'_a} \delta(O_2, \sigma(K), e'_a) \triangleright c'_1 \sigma[e'_a/e_a]$, for any $e' \notin X_{O_2}$.*

If we take any two transitions of $O_1 \triangleright c_1$ and $O_2 \triangleright c_2$ that correspond by this lemma, and we apply the rule in Definition 3.8 to them, it can be easily verified that we get the same transition, no matter the choice of e_a and e'_a . Therefore, all the infinitely-many P-markings whose abstract version is $O \triangleright c$ generate precisely the same transitions of $O \triangleright c$, and transitions that differ for the choice of the fresh event are all identified. This means that CG_{AC} is finitely-branching.

There is again a similarity with the π -calculus. A well-known technique to make the π -calculus LTS finitely-branching is to only take α -equivalence representatives. For instance, if $(y)\bar{x}y.p$ is such a representative, then the transition $(y)\bar{x}y.p \xrightarrow{\bar{x}(y)} p$ is enough to represent all the analogous transitions from α -equivalent processes. We can also omit y from the label, because its identity uniquely depends on the free names of $(y)\bar{x}y.p$. This is similar to the presentation of the π -calculus using abstraction and concretion operators [23, 4.3.1]. Here a transition from $(y)\bar{x}y.p$ is labelled by \bar{x} and goes to the concretion $\langle \nu y \rangle p$, where y is bound. Incidentally, this presentation naturally arises from the coalgebraic semantics of the π -calculus [13], and its implementation in logical frameworks.

Example 3.10. The CG_{AC} for the running example can be represented again by Figure 1. If we assume that depicted posets are abstract (i.e., translation maps from concrete to abstract posets are identities) then, in order to get a CG_{AC} , we just have to remove the universal quantification over events, and also remove the generated event from the label. The result is a finitely-branching CG, where each state has only one transition for each net transition. The state-space is still infinite, because posets keep growing along transitions.

Definition 3.11 (Abstract causal bisimilarity). An *abstract causal bisimulation* (**AC**-bisimulation in short) is a family of relations $\{R_O\}$, indexed by abstract posets, such that:

- whenever $(O_1 \triangleright c_1, O_2 \triangleright c_2) \in R_O$ then $O_1 = O_2 = O$;
- whenever $(O \triangleright c_1, O \triangleright c_2) \in R_O$ and $O \triangleright c_1 \xrightarrow{K \vdash a} O' \triangleright c'_1$ then $O \triangleright c_2 \xrightarrow{K \vdash a} O' \triangleright c'_2$ and $(O' \triangleright c'_1, O' \triangleright c'_2) \in R_{O'}$ (and viceversa).

The greatest such relation is denoted by \sim_{AC} .

We have the following correspondence between \sim_{C} and \sim_{AC} .




Figure 2: Example net.

Theorem 3.12. *Let $O \triangleright c_1$ and $O \triangleright c_2$ be (concrete) P-markings. Then $O \triangleright c_1 \sim_c O \triangleright c_2$ if and only if $[O]_{\cong} \triangleright \parallel c_1 \parallel_O \sim_{AC} [O]_{\cong} \triangleright \parallel c_2 \parallel_O$.*

We list some closure properties, which will be important in the following.

Proposition 3.13. *Transitions of CG_{AC} are preserved and reflected by order-embeddings $\sigma: O \rightarrow O'$, that is:*

- (i) *If $O \triangleright c \xrightarrow{K \vdash a} \delta(O, K, a) \triangleright c'$ then $O' \triangleright (c\sigma) \downarrow_{O'} \xrightarrow{\sigma(K) \vdash a} \delta(O', \sigma(K), a) \triangleright (c'\sigma^+) \downarrow_{\delta(O', \sigma(K), a)}$ (preservation);*
- (ii) *If $O' \triangleright (c\sigma) \downarrow_{O'} \xrightarrow{K' \vdash a} \delta(O', K', a) \triangleright c'$ then there are K and c'' such that $\sigma(K) = K'$, $(c''\sigma^+) \downarrow_{\delta(O', K', a)} = c'$ and $O \triangleright c \xrightarrow{K \vdash a} \delta(O, K, a) \triangleright c''$ (reflection).*

The definition of preservation and reflection are quite involved, due to the presence of event generation and the need of applying the closure operator to compute proper continuations. We will see that the categorical counterparts of these properties will be remarkably simpler.

Example 3.14. We motivate the requirement of order-reflection by showing that transitions of CG_{AC} are not reflected by functions without such property.

Consider the marked net of Figure 2. We can derive its CG_{AC} as shown for the running example. In it, from the initial P-marking $\emptyset \triangleright \{\emptyset \vdash s_1, \emptyset \vdash s_2\}$ we can reach the transition

$$\{e_a, e'_b\} \triangleright \{\{e_a\} \vdash r_1, \{e'_b\} \vdash r_2\} \xrightarrow{\{e_a, e'_b\} \vdash c} \{e_a \leq e''_c, e'_b \leq e''_c\} \triangleright \{\{e_a, e'_b, e''_c\} \vdash s_3\}$$

Consider the function $\sigma: \{e_a, e'_b\} \rightarrow \{e_a \leq e'_b\}$, mapping events to themselves. Clearly σ does not reflect posets. If we apply σ and then $\downarrow_{\{e_a \leq e'_b\}}$ to the source P-marking we get

$$\{e_a \leq e'_b\} \triangleright \{\{e_a\} \vdash r_1, \{e_a, e'_b\} \vdash r_2\}$$

but its c transition is

$$\{e_a \leq e'_b\} \triangleright \{\{e_a\} \vdash r_1, \{e_a, e'_b\} \vdash r_2\} \xrightarrow{\{e'_b\} \vdash c} \{e_a \leq e''_c \leq e'_b\} \triangleright \{\{e_a, e'_b, e''_c\} \vdash s_3\}$$

because only e'_b is maximal. However, this transition cannot be obtained from the one of $\{e_a, e'_b\} \triangleright \{\{e_a\} \vdash r_1, \{e'_b\} \vdash r_2\}$ via an application of σ .

The following theorem is a consequence of Proposition 3.13.

Theorem 3.15. *\sim_{AC} is closed under order-embeddings. Explicitly: for all order-embeddings $\sigma: O \rightarrow O'$, we have $O \triangleright c \sim_{AC} O \triangleright c'$ if and only if $O' \triangleright (c\sigma) \downarrow_{O'} \sim_{AC} O' \triangleright (c'\sigma) \downarrow_{O'}$.*

3.3. Immediate causes CG

We now introduce a further refinement of CG_{AC} , called *immediate causes CG* (CG_{IC}): we keep only *immediate causes*, i.e., causes that are maximal w.r.t. at least one of the tokens, and we identify isomorphic states. Immediate causes of a causal marking w.r.t. a poset O are given by

$$ic_O(K \vdash s) = \max_O(K) \quad ic_O(c_1 \cup c_2) = ic_O(c_1) \cup ic_O(c_2)$$

We define isomorphism of P-markings as follows: $O \triangleright c \cong O' \triangleright c'$ if and only if there is an isomorphism $\sigma: O \rightarrow O'$ such that $c\sigma = c'$. We denote by $[O \triangleright c]_{\cong}$ a chosen representative for the isomorphism class of $O \triangleright c$.

Definition 3.16 (Minimal P-marking). A *minimal* P-marking $O \blacktriangleright c$ is an abstract P-marking such that:

- $|O| = \mathcal{K}(c)$;
- for each $K \vdash s \in c$, $K \subseteq ic_O(c)$;
- it is a canonical isomorphism representative, i.e., $O \blacktriangleright c = [O \triangleright c]_{\cong}$.

Consider an abstract P-marking $O \triangleright c$. In order to compute the corresponding minimal P-marking $[O \triangleright c]_{\cong}$, we first take immediate causes for each token. Then, since the resulting P-marking may not be abstract, we take its canonical isomorphism representative. Formally, let O_{I} be O restricted to $ic_O(c)$, then

$$[[O \triangleright c]] = [O_{\text{I}} \triangleright \text{norm}_{O_{\text{I}}}(c)]_{\cong}$$

where $\text{norm}_O(K \vdash s) = K \cap |O_{\text{I}}| \vdash s$ and has an element-wise action on sets. We denote by $\langle O \triangleright c \rangle$ the map $[O_{\text{I}}]_{\cong} \rightarrow O$ obtained by composing a chosen isomorphism $[O_{\text{I}}]_{\cong} \rightarrow O_{\text{I}}$ and the embedding $O_{\text{I}} \hookrightarrow O$.

Definition 3.17 (Immediate causes CG). The *immediate causes CG* (CG_{IC}) is the smallest CG generated by the following rule

$$\frac{O \triangleright c \xrightarrow{K \vdash a} O' \triangleright c'}{O \blacktriangleright c \xrightarrow[\langle O' \triangleright c' \rangle]{K \vdash a} [[O' \triangleright c']]}_{\text{IC}}$$

This rule relies on the fact that minimal P-markings are also ordinary ones, so it takes the transition in CG_{AC} from a minimal P-marking, replaces the continuation $O' \triangleright c'$ with its minimal version $[[O' \triangleright c']]$ and, in order to keep track of the original identity of events, equips the transition with a *history map* $\langle O' \triangleright c' \rangle$, mapping canonical events to the original ones. In particular, the one with image $\text{new}(O', K, a)$ is the fresh event generated by the original transition.

The CG_{IC} has a finite state-space in many cases. We give a sufficient condition on the net from which the CG_{IC} is generated.

Proposition 3.18. *Given a net N with initial marking m_0 , if $[m_0]$ is finite then the corresponding CG_{IC} , reachable from $\emptyset \blacktriangleright \emptyset \vdash m_0$, has a finite state-space.*

Example 3.19. In order to derive a CG_{IC} for the running example, we take the P-markings of Figure 1 and we compute their minimal versions. For instance, we have

$$\begin{array}{c}
\{e_b \leq e'_b\} \triangleright \{\{e_b, e'_b\} \vdash s_1, \{e_b, e'_b\} \vdash s_2\} \\
\downarrow \text{immediate causes} \\
\{e'_b\} \triangleright \{\{e'_b\} \vdash s_1, \{e'_b\} \vdash s_2\} \\
\downarrow \text{canonical representative} \\
\{e_b\} \blacktriangleright \{\{e_b\} \vdash s_1, \{e_b\} \vdash s_2\}
\end{array}$$

because we assumed that $\{e_b\}$ is an abstract poset. Notice that the resulting P-marking is already in Figure 1. This is a crucial fact: minimization identifies many states and in some cases it even produces a finite state-space, as stated in Proposition 3.18. This is indeed the case for the running example.

Figure 3 shows the part of the running example's CG_{IC} that is reachable from $\{e_b\} \blacktriangleright \{\{e_b\} \vdash s_1, \{e_b\} \vdash s_2\}$. Most history maps are irrelevant, so they are omitted. Notice that in the CG_{AC} , from this P-marking, there are infinitely many transitions with action b . These all become a single loop over the same P-marking in the CG_{IC} ; the associated history map h_1 tells that e_b , after the transition, represents the most recent event, and that the previous event is discarded. Analogously for the two loops over $\{e_a, e'_a\} \blacktriangleright \{\{e_a\} \vdash s_1, \{e'_a\} \vdash s_2\}$. The interesting fact to notice is that our definition of h_2 and h_3 is not the only possible one. For instance, we could exchange the images of e_a and e'_a in the definition of h_2 . This is due to the fact that $\{e_a, e'_a\}$ has an automorphism that swaps e_a and e'_a .

Remark 3.20. The generation of the CG_{IC} from a net has been performed in two steps for the sake of clarity, but we can easily imagine an algorithm that performs it in a single step and incrementally. Given any P-marking, this is turned into a minimal one by taking immediate causes and then its canonical representative. Then outgoing transitions are computed from this P-marking, and the algorithm is applied to their continuations. Notice that minimizing a P-marking may yield a previously computed one: in this case the algorithm is not reapplied on that P-marking.

The notion of bisimilarity for CG_{IC} is more involved: while, given two P-markings, we may find a common poset for them (if any), which enables them to be compared w.r.t. \sim_{AC} , this is not always possible for posets of minimal P-markings. In other words, events in ordinary P-markings have a *global* identity, while those in minimal P-markings have a *local* identity. Therefore, we need to introduce an explicit correspondence between them. This correspondence can be a partial function, because some events may not be observable.

Definition 3.21 (Immediate causes bisimilarity). An *immediate causes bisimulation* R (IC-bisimulation in short) is a ternary relation such that, whenever $(O_1 \blacktriangleright c_1, \sigma, O_2 \blacktriangleright c_2) \in R$:

- σ is a partial isomorphism (i.e., an isomorphism between subposets) from O_1 to O_2 ;
- if $O_1 \blacktriangleright c_1 \xrightarrow[h_1]{K \vdash a} O'_1 \blacktriangleright c'_1$ then σ is defined on K , and there are $O_2 \blacktriangleright c_2 \xrightarrow[h_2]{\sigma(K) \vdash a} O'_2 \blacktriangleright c'_2$ and σ'

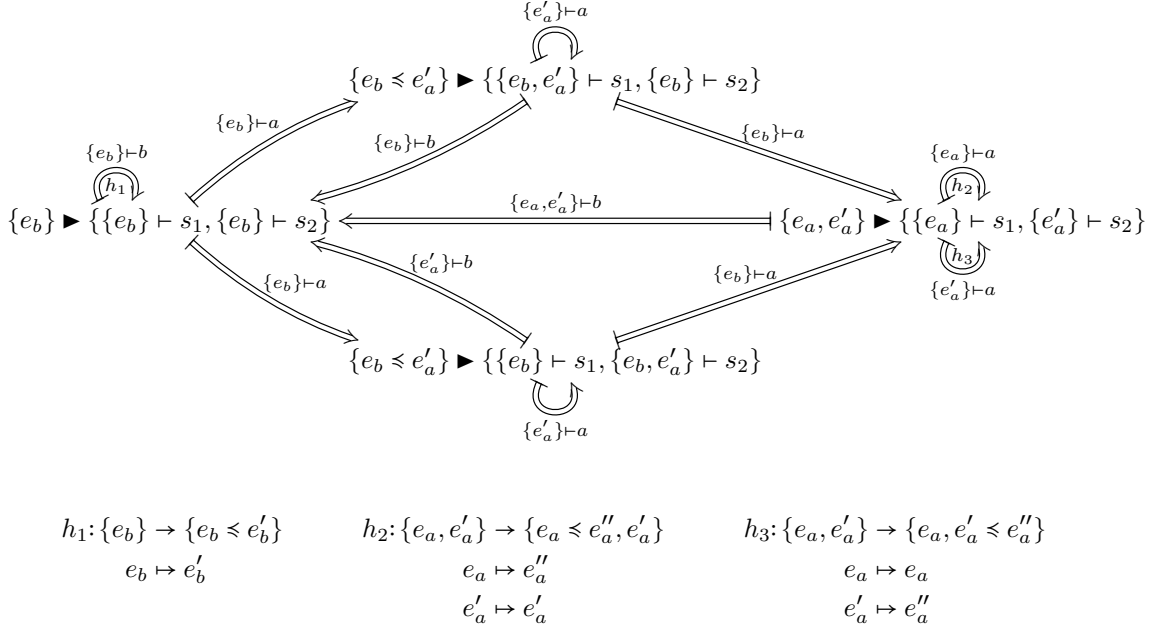


Figure 3: CGIC for the running example.

such that $(O'_1 \blacktriangleright c'_1, \sigma', O'_2 \blacktriangleright c'_2) \in R$ and the following diagram commutes

$$\begin{array}{ccc}
 O'_1 & \xrightarrow{h_1} & \delta(O_1, K, a) \\
 \sigma' \downarrow & & \downarrow \sigma^+ \\
 O'_2 & \xrightarrow{h_2} & \delta(O_2, \sigma(K), a)
 \end{array}$$

- if $O_2 \blacktriangleright c_2 \xrightarrow[h_2]{K \vdash a} O'_2 \blacktriangleright c'_2$ then σ is defined on K , and there are $O_1 \blacktriangleright c_1 \xrightarrow[h_1]{\sigma(K)^{-1} \vdash a} O'_1 \blacktriangleright c'_1$ and σ' as in the previous item.

The greatest such bisimulation is denoted \sim_{IC} . We write $O_1 \blacktriangleright c_1 \sim_{\text{IC}}^{\sigma} O_2 \blacktriangleright c_2$ to mean

$$(O_1 \blacktriangleright c_1, \sigma, O_2 \blacktriangleright c_2) \in \sim_{\text{IC}} .$$

The commuting diagram essentially says that σ' should never map old events to new ones (or viceversa). More precisely, given $x \in |O'_1|$, we have two cases:

- $h_1(x) = \text{new}(O_1, K, a)$, then, by definition, $h_1(x)$ is mapped by σ^+ to $\text{new}(O_2, \sigma(K), a)$, so $\sigma'(x) = y$ such that $h_2(y) = \text{new}(O_2, \sigma(K), a)$;
- $h_1(x) = \text{old}(O_1, K, a)(x')$, for some x' , then $h_1(x)$ is mapped by σ^+ to $\text{old}(O_2, \sigma(K), a)(\sigma(x'))$, so $\sigma'(x) = y$ such that $h_2(y) = \text{old}(O_2, \sigma(K), a)(\sigma(x'))$.

We have the following correspondence between \sim_{IC} and \sim_{AC} .

Theorem 3.22. \sim_{IC} is fully abstract w.r.t. \sim_{AC} in the following sense:

- (i) If $O \triangleright c_1 \sim_{\text{AC}} O \triangleright c_2$ then $\llbracket O \triangleright c_1 \rrbracket \sim_{\text{IC}} \llbracket O \triangleright c_2 \rrbracket$;
- (ii) If $O_1 \blacktriangleright c_1 \sim_{\text{IC}}^{\sigma} O_2 \blacktriangleright c_2$ then for all $O \triangleright \hat{c}_1$ and $O \triangleright \hat{c}_2$ such that:
 - (a) $\llbracket O \triangleright \hat{c}_1 \rrbracket = O_1 \blacktriangleright c_1$ and $\llbracket O \triangleright \hat{c}_2 \rrbracket = O_2 \blacktriangleright c_2$;
 - (b) $\langle O \triangleright \hat{c}_1 \rangle|_{\text{dom}(\sigma)} = \langle O \triangleright \hat{c}_2 \rangle \circ \sigma$;
we have $O \triangleright \hat{c}_1 \sim_{\text{AC}} O \triangleright \hat{c}_2$.

Statement (i) is self-explanatory. Statement (ii) says that if we have two equivalent minimal P-markings $O_1 \blacktriangleright c_1 \sim_{\text{IC}}^{\sigma} O_2 \blacktriangleright c_2$ and we take any two P-markings $O \triangleright \hat{c}_1$ and $O \triangleright \hat{c}_2$ whose minimal versions are $O_1 \blacktriangleright c_1$ and $O_2 \blacktriangleright c_2$ respectively ((ii)(a)), these are equivalent provided that local events matched by σ have the same global interpretation as events of O ((ii)(b)).

3.4. Immediate causes CG with symmetries

The final step is to introduce *symmetries* over states of CG. Given an abstract poset O , a *symmetry over O* is a set Φ of automorphisms $O \rightarrow O$ (called just *permutations* hereafter) such that $\text{id} \in \Phi$ and it is closed under composition. This section is an adaptation of the work in [21, 19] on the set-theoretic version of HD-automata for the π -calculus.

We now motivate the introduction of symmetries. We say that two CG_{ICs} are *isomorphic* when there is a bijective correspondence ω between their P-markings and, for each P-marking $O \blacktriangleright c$ of the former such that $\omega(O \blacktriangleright c) = O' \blacktriangleright c'$, transitions from $O' \blacktriangleright c'$ can be obtained from those of $O \blacktriangleright c$ via an isomorphism. In the case of ordinary labelled transition systems (LTSs), one can compute minimal versions w.r.t. bisimilarity, where all bisimilar states have been identified. Bisimilar LTSs have isomorphic minimal versions, so we may use any of them as canonical representative of the class of bisimilar LTSs. This cannot be done for CG_{ICs} , because of the following fact.

Proposition 3.23. *There are minimal CG_{ICs} that are \sim_{IC} -bisimilar but not isomorphic.*

Example 3.24. Consider the P-marking $\{e_a, e'_a\} \blacktriangleright \{\{e_a\} \vdash s_1, \{e'_a\} \vdash s_2\}$ of Example 3.19 and its looping transitions. Take another P-marking $\{e_a, e'_a\} \blacktriangleright \{\{e_a\} \vdash s'_1, \{e'_a\} \vdash s'_2\}$ with the following transitions

$$\begin{array}{ccc}
\begin{array}{c} \{e_a\} \vdash a \\ \text{\scriptsize } \curvearrowright \text{\scriptsize } h_4 \\ \{e_a, e'_a\} \blacktriangleright \{\{e_a\} \vdash s'_1, \{e'_a\} \vdash s'_2\} \\ \text{\scriptsize } \curvearrowleft \text{\scriptsize } h_5 \\ \{e'_a\} \vdash a \end{array} & \begin{array}{c} h_4: \{e_a, e'_a\} \mapsto \{e_a \leq e''_a, e'_a\} \\ e_a \mapsto e''_a \\ e'_a \mapsto e'_a \end{array} & \begin{array}{c} h_5: \{e_a, e'_a\} \mapsto \{e_a, e'_a \leq e''_a\} \\ e_a \mapsto e''_a \\ e'_a \mapsto e_a \end{array}
\end{array}$$

Notice that we have $h_4 = h_2$ and $h_5 = h_3 \circ \phi$, where ϕ switches e_a and e'_a .

Suppose we want to find a minimal realization of these CGs. They are not isomorphic, in the sense that there is no permutation on $\{e_a, e'_a\}$ that, applied to labels and composed with history maps, turns transitions of the former CG into those of the latter. However, we have

$$\{e_a, e'_a\} \blacktriangleright \{\{e_a\} \vdash s_1, \{e'_a\} \vdash s_2\} \sim_{\text{IC}}^{\phi} \{e_a, e'_a\} \blacktriangleright \{\{e_a\} \vdash s'_1, \{e'_a\} \vdash s'_2\} ,$$

so these states should be identified in some way. This way is provided by symmetries: minimal behavior, according to \sim_{ICS} , is invariant under ϕ , so we can identify those P-markings, provided that the resulting state is annotated with ϕ and possibly other permutations that fix the state.

The same argument applies when considering versions of the same CG_{IC} that only differ for the choice of history maps: if $s'_1 = s_1$ and $s'_2 = s_2$ in the P-marking $\{e_a, e'_a\} \blacktriangleright \{\{e_a\} \vdash s'_1, \{e'_a\} \vdash s'_2\}$ above, then the P-marking $\{e_a, e'_a\} \blacktriangleright \{\{e_a\} \vdash s'_1, \{e'_a\} \vdash s'_2\}$ is bisimilar to itself under the permutation ϕ . This has a practical consequence: when constructing the CG_{IC} for a given net, one should not spend computational effort in computing the “right” history maps, because the choice of history maps does not affect bisimilarity and thus minimal models.

Definition 3.25 (Minimal P-marking with symmetry). A *minimal P-marking with symmetry* is a triple $O \blacktriangleright_{\Phi} c$, where $O \blacktriangleright c$ is a minimal P-marking and Φ is a symmetry over O such that $c\phi = c$, for all $\phi \in \Phi$.

Symmetries allow us to remove some transitions from CG_{IC} : we can only take one representative transition among all the *symmetric* ones, i.e., those whose observable causes and history maps only differ for some permutations in the symmetries of source and target states.

Definition 3.26 (Symmetric transitions). Given $O \blacktriangleright_{\Phi} c$, $O' \blacktriangleright_{\Phi'} c'$ and two transitions

$$O \blacktriangleright c \xrightarrow[h_1]{K_1 \vdash a} O' \blacktriangleright c' \quad O \blacktriangleright c \xrightarrow[h_2]{K_2 \vdash a} O' \blacktriangleright c'$$

they are *symmetric* if and only if there are $\phi \in \Phi$ and $\phi' \in \Phi'$ such that $K_2 = \phi(K_1)$ and the following diagram commutes

$$\begin{array}{ccc} O' & \xrightarrow{h_1} & \delta(O, K_1, a) \\ \phi' \downarrow & & \downarrow \phi' \\ O' & \xrightarrow[h_2]{} & \delta(O, K_2, a) \end{array}$$

We write $\{K\}$ and $\{h\}$ for a canonical choice of K and h among those of all the symmetric transitions. Actually $\{h\}$ depends on the considered symmetries Φ and Φ' , but they are omitted to simplify notation: they will always be clear from the context.

Definition 3.27 (CG_{ICS}). The *CG_{IC} with symmetries* (CG_{ICS}) is the smallest CG generated by the following rule

$$\frac{O \blacktriangleright c \xrightarrow[h]{K \vdash a} O' \blacktriangleright c'}{O \blacktriangleright_{\Phi} c \xrightarrow[\{h\}]{\{K\} \vdash a} O' \blacktriangleright_{\Phi'} c'}$$

The notion of bisimulation is analogous to IC-bisimulation. However, P-markings are required to simulate each other only up to symmetries. More specifically, when comparing $O_1 \blacktriangleright_{\Phi_1} c_1$ and $O_2 \blacktriangleright_{\Phi_2} c_2$ under a mediating map σ , for each permutation in Φ_1 and each transition of the first P-marking, we have to find a permutation in Φ_2 and a transition of the second P-marking. The correspondence between observable causes and between history maps must be as in IC-bisimulations, but the action of mediating maps is changed according to the considered permutations.

Definition 3.28 (Immediate causes bisimulation with symmetries). An *immediate causes bisimulation with symmetries* R (ICS-bisimulation in short) is a ternary relation such that, whenever $(O_1 \blacktriangleright_{\Phi_1} c_1, \sigma, O_2 \blacktriangleright_{\Phi_2} c_2) \in R$:

- σ is a partial isomorphism from O_1 to O_2 ;
- for each $\phi_1 \in \Phi_1$ and $O_1 \blacktriangleright_{\Phi_1} c_1 \xrightarrow[h_1]{K_1 \vdash a} O'_1 \blacktriangleright_{\Phi'_1} c'_1$, σ is defined on $\phi_1(K)$ and there are $\phi_2 \in \Phi_2$ and $O_2 \blacktriangleright_{\Phi_2} c_2 \xrightarrow[h_2]{K_2 \vdash a} O'_2 \blacktriangleright_{\Phi'_2} c'_2$ such that:
 - $K_2 = \gamma(K_1)$, for $\gamma = \phi_2^{-1} \circ \sigma \circ \phi_1$;
 - there is σ' such that $(O'_1 \blacktriangleright_{\Phi'_1} c'_1, \sigma', O'_2 \blacktriangleright_{\Phi'_2} c'_2) \in R$ and the following diagram commutes

$$\begin{array}{ccc} O'_1 & \xrightarrow{h_1} & \delta(O_1, K_1, a) \\ \sigma' \downarrow & & \downarrow \gamma^* \\ O'_2 & \xrightarrow[h_2]{} & \delta(O_2, K_2, a) \end{array}$$

(and viceversa)

The greatest such relation is denoted \sim_{ICS} and we write $O_1 \blacktriangleright_{\Phi_1} c_1 \sim_{\text{ICS}}^\sigma O_2 \blacktriangleright_{\Phi_2} c_2$ whenever $(O_1 \blacktriangleright_{\Phi_1} c_1, \sigma, O_2 \blacktriangleright_{\Phi_2} c_2) \in \sim_{\text{ICS}}$.

As mentioned, symmetries allow computing minimal realizations, where all bisimilar P-markings are identified. More precisely, we can identify \sim_{ICS} -equivalent P-markings, namely $O_1 \blacktriangleright_{\Phi_1} c_1$ and $O_2 \blacktriangleright_{\Phi_2} c_2$ that are related by \sim_{ICS}^σ , for some σ . Then σ becomes part of the state symmetry. Actually, σ is a permutation between subsets of O_1 and O_2 , but it can be shown that all \sim_{ICS} -equivalent P-markings have the same poset of observable events on which σ is defined. This means that σ is indeed a permutation on that poset.

Definition 3.29 (Minimal CG_{ICS}). The minimal CG_{ICS} is defined as follows:

- states are canonical representatives of \sim_{ICS} -equivalence, namely $O \blacktriangleright_{\Phi} c$ such that $\Phi = \{\sigma \mid \exists \Phi' : O \blacktriangleright_{\Phi'} c \sim_{\text{ICS}}^\sigma O \blacktriangleright_{\Phi'} c\}$;
- transitions are derived according to Definition 3.27.

In order to compute the symmetry Φ of a canonical representative $O \blacktriangleright_{\Phi} c$, we take P-markings of the form $O \blacktriangleright_{\Phi'} c$ and we consider triples where $O \blacktriangleright_{\Phi'} c$ is bisimilar to itself. Notice that Φ may be different than Φ' : some $\phi \in \Phi$, in fact, may not act as the identity of c ; with a little abuse of notation, $O \blacktriangleright_{\Phi} c$ stands for a P-marking where every $\phi \in \Phi$ has identical action on c up to bisimilarity. It can be proved that we do not need to consider non-canonical P-markings for the computation of Φ (see, e.g., [19, 5.2]).

Example 3.30. Consider the CG_{IC} of Example 3.19. It can be regarded as a CG_{ICS} where all states have the singleton symmetry $\{id\}$. Its minimal version is depicted in Figure 4. Notice that the P-marking $\{e_a, e'_a\} \blacktriangleright_{\Phi_3} \{\{e_a\} \vdash s_1, \{e'_a\} \vdash s_2\}$ has a non trivial symmetry, because we have $\{e_a, e'_a\} \blacktriangleright_{\{id\}} \{\{e_a\} \vdash s_1, \{e'_a\} \vdash s_2\} \sim_{\text{ICS}}^{(e_a \ e'_a)} \{e_a, e'_a\} \blacktriangleright_{\{id\}} \{\{e_a\} \vdash s_1, \{e'_a\} \vdash s_2\}$.




Figure 4: Minimal CG_{ICS} for the running example.

4. Causal case graphs and behavior structures

In the pioneering work [27] of Trakhtenbrot and Rabinovich, *behavior structures* have been introduced as causal models for Petri nets. In this section we compare them with our causal models. We recall a slightly simplified definition.

Definition 4.1 (Behavior structure). Let Act be a set of action labels. A behavior structure (BS in short) is a triple $B = (M, P, \phi)$, where:

- M is an automaton such that:
 - transitions have the form $n \xrightarrow{a}_B m$, with $a \in Act$;
 - all states are reachable from the initial one r ;
 - there are no oriented cycles, i.e., sequences of transitions where the first and last state coincide;
 - there are no parallel edges, i.e., $n \xrightarrow{a}_B m$ and $n \xrightarrow{b}_B m$ implies $a = b$.
- P is a family P_n of Act -labeled posets of events, one for each state n of M (for the root state r we must have $P_r = \emptyset$);
- ϕ is a family of labeled posets morphisms: for each pair of states n and m such that $n \xrightarrow{a}_B m$
 - $\phi_{n,m}$ is an isomorphic embedding of P_n as a prefix of P_m ;
 - $|P_m| \setminus |\phi_{n,m}(P_n)| = \{e_a\}$, for some event e ;

In a BS, each state n has a poset P_n over labelled events, describing causal dependencies among occurrences of actions that led to n . For each transition $n \xrightarrow{a}_B m$ we have a map $\phi_{n,m}$ telling the correspondence between P_n and P_m : P_n is required to be isomorphic to a prefix of P_m because it should specify causal dependencies for all the previous actions. The only additional event in P_m represents an occurrence of the most recent action a .

The associated notion of behavioral equivalence is called *BS-bisimilarity*. In [27], this equivalence compares two different behavior structures. Here states belong to the same behavior structure.

Definition 4.2 (BS-bisimulation). Given a behavior structure B , a *BS-bisimulation* on B is a relation R on triples such that, whenever $(n_1, \sigma, n_2) \in R$:

- σ is an isomorphism between P_{n_1} and P_{n_2} ;
- if $n_1 \xrightarrow{a}_B m_1$ then there exist m_2, σ' such that $n_2 \xrightarrow{a}_B m_2$ with $(m_1, \sigma', m_2) \in R$ and the following diagram commutes

$$\begin{array}{ccc} P_{n_1} & \xrightarrow{\sigma} & P_{n_2} \\ \phi_{n_1, m_1} \downarrow & & \downarrow \phi_{n_2, m_2} \\ P_{m_1} & \xrightarrow{\sigma'} & P_{m_2} \end{array}$$

(and viceversa)

The greatest such relation, denoted \sim_{bs} , is called *BS-bisimilarity*.

Notice that states are related by BS-bisimulations up to an isomorphism of their posets. This is because the actual identity of events should not matter when comparing states. Only the causal dependencies between occurrences of actions are relevant. BS-bisimilarity has been called *history preserving bisimilarity* [14] in later work.

4.1. Relationship with causal case graphs

When used to represent the behavior of Petri nets, states of behavior structures are states of deterministic, non-sequential processes equipped with information about the past history of events. They can equivalently be seen as tokens equipped with causal information (see, e.g., [18]). Therefore, we will consider behavior structures over causal markings. This will enable a more direct comparison with our causal case graphs.

We characterize a sub-LTS of $CG_{\mathbb{C}}$ that is equivalent to a BS.

Definition 4.3 (Reachable $CG_{\mathbb{C}}$). The *reachable* $CG_{\mathbb{C}}$ ($CG_{\mathbb{C}}^r$) is defined as follows:

- it has an *initial* P-marking $\emptyset \triangleright \emptyset \vdash m_0$, where m_0 is an initial marking for N ;
- transitions are only those reachable from $\emptyset \triangleright \emptyset \vdash m_0$.

$CG_{\mathbb{C}}^r$ enjoys some properties that allow us to define a BS on top of it.

Lemma 4.4.

- Each state $O_c \triangleright c$ of $CG_{\mathbb{C}}^r$ has a unique possible poset, i.e., for any other state $O \triangleright c$ we have $O = O_c$; moreover, we have $|O_c| = \mathcal{K}(c)$.
- $CG_{\mathbb{C}}^r$ does not have parallel transitions and directed cycles.

Proposition 4.5. The triple $B_{\mathbb{C}} = (M^{\mathbb{C}}, \phi^{\mathbb{C}}, P^{\mathbb{C}})$ is a behavior structure, where

- $M^{\mathbf{C}}$ is the smallest automaton generated from $CG_{\mathbf{C}}^r$ via the following rule

$$\frac{O_c \triangleright c \xrightarrow{K \vdash e_a} O_{c'} \triangleright c'}{c \xrightarrow{a}_{B_{\mathbf{C}}} c'}$$

- $P^{\mathbf{C}} = \{O_c \mid O_c \triangleright c \text{ is a state of } CG_{\mathbf{C}}^r\};$
- $\phi^{\mathbf{C}} = \{\phi_{c,c'}^{\mathbf{C}} : O_c \hookrightarrow O_{c'} \mid O_c \triangleright c \xrightarrow{K \vdash e_a} O_{c'} \triangleright c'\}.$

We have the following relation between $\sim_{\mathbf{C}}$ and BS bisimilarity.

Theorem 4.6. *Let c_1, c_2 be states of $B_{\mathbf{C}}$. Then*

- (i) *If $O \triangleright c_1 \sim_{\mathbf{C}} O \triangleright c_2$ and there is an isomorphism $\sigma: O_{c_1} \rightarrow O_{c_2}$ then $c_1 \sim_{bs}^{\sigma} c_2$;*
- (ii) *$c_1 \sim_{bs}^{\sigma} c_2$ implies $O_{c_2} \triangleright c_1 \sigma \sim_{\mathbf{C}} O_{c_2} \triangleright c_2$.*

Statement (i) says that two states c_1 and c_2 in $B_{\mathbf{C}}$ with isomorphic posets are \sim_{bs} -bisimilar whenever any two P-markings over c_1 and c_2 are $\sim_{\mathbf{C}}$ -bisimilar. Statement (ii) is somewhat dual: if c_1 and c_2 are \sim_{bs} -bisimilar under an isomorphism σ , then we can use σ to turn them into $\sim_{\mathbf{C}}$ -bisimilar P-markings.

Remark 4.7. The behavior structure we have introduced has some common aspects with $CG_{\mathbf{IC}}$: for both, posets in states have local meanings; in fact, bisimilarities require explicit mappings between posets of simulating states. However, $CG_{\mathbf{IC}}$ can discard event names along transitions and go back to an already visited state, whereas this is explicitly forbidden for BSs.

5. Background on category theory

We assume that the reader is familiar with elementary category theory. In this section we recall some notions that will be needed in the following.

5.1. Functor categories

Definition 5.1 (Functor category). Let \mathbf{C} and \mathbf{D} be two categories. The *functor category* $\mathbf{D}^{\mathbf{C}}$ has functors $\mathbf{C} \rightarrow \mathbf{D}$ as objects and natural transformations between them as morphisms.

Functors from any category \mathbf{C} to \mathbf{Set} are called (covariant) *presheaves*. Hereafter we assume that the domain category \mathbf{C} for presheaves is *small*, i.e., its collection of objects is actually a set. A presheaf P can be intuitively seen as a family of sets indexed over the objects of \mathbf{C} plus, for each $\sigma: c \rightarrow c'$ in \mathbf{C} , an action of σ on Pc , which we write

$$p[\sigma]_P = P\sigma(p) \quad (p \in Pc) ,$$

omitting the subscript P in $[\sigma]_P$ when clear from the context. This notation intentionally resembles the application of a renaming σ to a process p , namely $p\sigma$: it will, in fact, have this meaning in later sections. The set $\int P$ of *elements* of a presheaf P is

$$\int P := \sum_{c \in |\mathbf{C}|} Pc$$

where the sum symbol denotes the coproduct in \mathbf{Set} , and we denote by $c \triangleright p$ a pair belonging to $\int P$. Presheaf categories have the following nice property.

Property 5.2. *For any \mathbf{C} , $\mathbf{Set}^{\mathbf{C}}$ has all limits and colimits, both computed pointwise.*

5.2. Coalgebras

The *behavior* of systems can be modeled in a categorical setting through *coalgebras* [22, 1]. Given a *behavioral endofunctor* $B: \mathbf{C} \rightarrow \mathbf{C}$, describing the “shape” of a class of systems, we have a corresponding category of coalgebras.

Definition 5.3 (*B-Coalg*). The category *B-Coalg* is defined as follows: objects are *B-coalgebras*, i.e., pairs (X, h) of an object $X \in |\mathbf{C}|$, called *carrier*, and a morphism $h: X \rightarrow BX$, called *structure map*; *B-coalgebra homomorphisms* $f: (X, h) \rightarrow (Y, g)$ are morphisms $f: X \rightarrow Y$ in \mathbf{C} making the following diagram commute

$$\begin{array}{ccc} X & \xrightarrow{h} & BX \\ f \downarrow & & \downarrow Bf \\ Y & \xrightarrow{g} & BY \end{array}$$

For instance, given a set of labels L , consider the functor

$$B_{flts} := \mathcal{P}_f(L \times -)$$

where $\mathcal{P}_f: \mathbf{Set} \rightarrow \mathbf{Set}$ is the *finite powerset functor*, defined on a set A and on a function $h: A \rightarrow A'$ as follows

$$\mathcal{P}_f A := \{B \subseteq A \mid B \text{ finite}\} \quad \mathcal{P}_f h(B) := \{h(b) \mid b \in B\}$$

B_{flts}-coalgebras (X, h) are *finitely-branching labelled transition systems*, with labels L and states X . The function $h(x)$ returns the set of labelled transitions $x \xrightarrow{a} y$ such that $(a, y) \in h(x)$. Homomorphisms of *B_{flts}*-coalgebras are functions between states that preserve and reflect transitions.

Many notions of behavioral equivalence can be defined for coalgebras (see [25]). We adopt the one by Hermida and Jacobs and we simply call it *B-bisimulation*. To introduce it, we need some preliminary notions. A (binary) *relation on* $X \in |\mathbf{C}|$ is a jointly-monic span $X \leftarrow R \rightarrow X$ in \mathbf{C} . An *image* of a morphism $f: A \rightarrow C$ is a monomorphism $m: B \rightarrow C$ through which f factors, such that if f factors through any other mono $B' \rightarrow C$, then B is a subobject of B' . The factoring morphism $A \rightarrow B$ is called *cover*. In \mathbf{Set} all these notions become the usual ones: a relation R is a binary relation on X and the span is made of left/right projections; the image of f is $f(A) \hookrightarrow C$, and its cover is f with restricted codomain $f(A)$. Given a relation R on X , the *relation lifting* \overline{BR} is the image of the morphism $BR \rightarrow B(X \times X) \rightarrow BX \times BX$, taking R to a relation on BX .

Definition 5.4 (*B-bisimulation*). Given a *B-coalgebra* (X, h) , a *B-bisimulation* on it is a relation R on X such that there is r making the following diagram commute

$$\begin{array}{ccccc} X & \longleftarrow & R & \longrightarrow & X \\ \downarrow h & & \downarrow r & & \downarrow h \\ BX & \longleftarrow & \overline{BR} & \longrightarrow & BX \end{array}$$


The greatest such relation is called *B-bisimilarity*.

A *B_{flts}*-bisimulation R on a *B_{flts}*-coalgebra is an ordinary bisimulation on the corresponding transition system. In fact, \overline{BR} is the set of pairs $(X_1, X_2) \in BX \times BX$ such that $(l, x') \in X_1$ only if there is some $(l, (x', y')) \in BR$, but then we also have $(l, y') \in X_2$ and $(x', y') \in R$ (the symmetric

statement holds if $(l, x') \in X_2$). Clearly r exists if and only if R is a bisimulation, and is given by $(x, y) \in R \mapsto (h(x), h(y))$.

An important property of categories of coalgebras is the existence of the terminal object; the unique morphism from each coalgebra to it assigns to each state its abstract semantics. The ideal situation is when the induced equivalence, relating all the states with the same abstract semantics, agrees with B -bisimilarity. A sufficient condition for this property is when B covers pullbacks.

Property 5.5 (B covers pullbacks). *Consider a cospan $X_1 \rightarrow X_3 \leftarrow X_2$, and the morphism m from the image of the pullback (the left square below) to the pullback of the image*



Then B covers pullbacks if m is always a cover.

For the best-known Aczel-Mendler bisimulations, defined as spans of coalgebras, the condition on B that guarantees the agreement of behavioral equivalences is more demanding: B should preserve weak pullbacks. The finite powerset functor on **Set** preserves weak pullbacks, but other finite powerset functors do not, for instance the one on presheaves that we will use, which instead covers pullbacks. This motivates our preference of Hermida-Jacobs bisimulations over Aczel-Mendler ones (another important reason for this will be explained in section 6).

A sufficient condition for the existence of the final coalgebra is that B is an *accessible* functor on a *locally finitely presentable* category (see [3, 29, 1] for details). A category \mathbf{C} is *filtered* if each finite diagram is the base of a cocone in \mathbf{C} ; filtered categories generalize the notion of directed preorders, that are sets such that every finite subset has an upper bound. For any category \mathbf{D} , a *filtered colimit* in \mathbf{D} is the colimit of a diagram of shape \mathbf{C} , i.e., a functor $\mathbf{C} \rightarrow \mathbf{D}$, such that \mathbf{C} is a filtered category.

Definition 5.6 (Locally finitely presentable category). An object c of a category \mathbf{C} is *finitely presentable* if the functor $\text{Hom}_{\mathbf{C}}(c, -): \mathbf{C} \rightarrow \mathbf{Set}$ preserves filtered colimits. A category \mathbf{C} is locally finitely presentable if it has all colimits and there is a set of finitely presentable objects $X \subseteq |\mathbf{C}|$ such that every object is a filtered colimit of objects from X .

For instance, locally finitely presentable objects in **Set** are precisely finite sets. **Set** is locally finitely presentable: every set is the filtered colimit, namely the union, of its finite subsets and the whole **Set** is generated by the set containing one finite set of cardinality n for all $n \in \mathbb{N}$.

For functor categories we have the following.

Proposition 5.7. *For each locally finitely presentable category \mathbf{C} and small category \mathbf{D} , the functor category $\mathbf{C}^{\mathbf{D}}$ is locally finitely presentable.*

In particular, since **Set** is locally finitely presentable, we have that the presheaf category $\mathbf{Set}^{\mathbf{D}}$ is locally finitely presentable as well.

Definition 5.8 (Accessible functor). Let \mathbf{C} and \mathbf{D} be locally finitely presentable categories. A functor $F: \mathbf{C} \rightarrow \mathbf{D}$ is *accessible* if it preserves filtered colimits.

Here are some useful properties of accessible functors: their products, coproducts and composition is accessible as well; adjoint functors between locally finitely presentable categories are accessible. Moreover, it is a well-known fact that the finite powerset functor \mathcal{P}_f introduced in section 5.2 is accessible.

5.3. Coalgebras over presheaves

Coalgebras for functors $B: \mathbf{Set}^{\mathbf{C}} \rightarrow \mathbf{Set}^{\mathbf{C}}$ are pairs (P, ρ) of a presheaf $P: \mathbf{C} \rightarrow \mathbf{Set}$ and a natural transformation $\rho: P \rightarrow BP$. The naturality of ρ imposes a constraint on behavior

$$\begin{array}{ccc}
 c & p \in Pc & \xrightarrow{\rho_c} & beh(p) \\
 f \downarrow & [f]_P \downarrow & & \downarrow [f]_{BP} \\
 c' & p[f]_P \in P(c') & \xrightarrow{\rho_{c'}} & beh(p)[\sigma]_{BP}
 \end{array}$$

Intuitively, this diagram means that, if we take a state, apply a function to it and then compute its behavior, we should get the same thing as first computing the behavior and then applying the function to it. In other words, behavior must be *preserved* and *reflected* by the index category morphisms.

B -bisimulations have a similar structure. A B -bisimulation R is a presheaf in $\mathbf{Set}^{\mathbf{C}}$ and all the legs of the bisimulation diagram in Definition 5.4 are natural transformations. In particular, the naturality of projections implies that, given $(p, q) \in Rc$ and $f: c \rightarrow c'$ in \mathbf{C} , $(p[f], q[f]) \in R(c')$, i.e., B -bisimulations are *closed under the index category morphisms*.

6. Coalgebraic semantics

In this section we construct a coalgebraic causal semantics for Petri Nets. We first show that the notions of section 3.2 have a categorical interpretation. Then we translate $\mathbf{CG}_{\mathbf{AC}}$ into a coalgebra.

We introduce two categories of *Act*-labelled posets. Recall that, given a category \mathbf{C} , a *skeletal category* is a full subcategory of \mathbf{C} such that each object is isomorphic to one of \mathbf{C} and two distinct objects cannot be isomorphic.

Definition 6.1 (Category \mathbf{O} and \mathbf{O}). Let \mathbf{O} be the skeletal category of the category of *Act*-labelled posets and their morphisms. The category \mathbf{O} is the subcategory of \mathbf{O} whose morphisms are order-embeddings.

Taking a skeletal category amounts to choosing one canonical representative of each isomorphism class of posets, i.e., using the terminology of section 3.2, the objects of \mathbf{O} and \mathbf{O} are abstract posets. The difference between \mathbf{O} and \mathbf{O} is similar to that between \mathbf{F} , the category of finite ordinals and all functions, and its subcategory \mathbf{I} , including only injective functions (indeed \mathbf{O} only includes injective morphisms). Presheaves over these categories are used in [13] to give a coalgebraic semantics for the π -calculus.

Remark 6.2. In [7] we have introduced the category \mathbf{P} of finite posets up to isomorphisms and its subcategory \mathbf{P}_m with only order-embeddings. The category \mathbf{O} can be understood as a comma category $U \downarrow \mathbf{Act}$, where $U: \mathbf{P} \rightarrow \mathbf{Set}$ takes a poset to its underlying set and \mathbf{Act} is the constant functor mapping every set to *Act*. Similarly for \mathbf{O} , whenever $U: \mathbf{P}_m \rightarrow \mathbf{Set}$.

Proposition 6.3. *The category \mathbb{O} is small and has pullbacks.*

The category \mathbb{O} lacks colimits, but the ones we are interested in can be computed in \mathbf{O} . We will be more precise when presenting such colimits.

We introduce some notation for particular objects and morphisms of \mathbf{O} . We denote by $[k]_l$ the discrete poset with k elements and labeling function l ; if $k = 1$ then we simply write $[1]_a$ to assign label a to the only event. We write $[k]_l^a$ for the poset $[k]_l$ plus a top element with label a . Two maps will be useful:

$$[k]_l \xrightarrow{b([k]_l^a)} [k]_l^a \xleftarrow{\top([k]_l^a)} [1]_a$$

the left map picks the bottom elements in $[k]_l^a$, and the right one picks the top element.

In \mathbf{O} we can use a pushout to compute $\delta(O, K, a)$, the associated maps $old(O, K, a)$ and $new(O, K, a)$, and the extension σ^+ of a morphism $\sigma: O \rightarrow O'$, all defined in section 3.2. Given $O \in |\mathbf{O}|$, let $K: [k]_l \hookrightarrow O$ be the subobject in \mathbb{O} picking K within O . Then we have

$$\begin{array}{ccc}
 [k]_l & \xrightarrow{K} & O & \xrightarrow{\sigma} & O' \\
 \downarrow b([k]_l^a) & & \downarrow old(O, K, a) & & \downarrow old(O', \sigma(K), a) \\
 [k]_l^a & \xrightarrow{K^a} & \delta(O, K, a) & \xrightarrow{\sigma^+} & \delta(O', \sigma(K), a) \\
 \downarrow id & & & & \downarrow \\
 [k]_l^a & \xrightarrow{(\sigma(K))^a} & \delta(O, \sigma(K), a) & &
 \end{array}
 \quad new(O, K, a) = K^a \circ \top([k]_l^a) \quad (1)$$

Explicitly, $\delta(O, K, a)$ is constructed as follows: the disjoint union of O and $[k]_l^a$ is made, and then the bottom elements of $[k]_l^a$ and the causes K are identified, resulting in O plus a fresh a -labelled top event for K ; the transitive closure of this relation gives $\delta(O, K, a)$. Notice that, since K reflects order, causes of the fresh event must be incomparable, i.e., they are maximal events in O . This agrees with the definition of K in Definition 3.3. The map $\sigma^+: \delta(O, K, a) \rightarrow \delta(O', \sigma(K), a)$ is induced by the universal property of pushouts: we compute $\delta(O', \sigma(K), a)$ via the pushout of

$$[k]_l \xleftarrow{b([k]_l^a)} [k]_l^a \xrightarrow{\sigma \circ K} O'$$

that is the outer pushout in (1), and then we define σ^+ as the mediating morphism between the inner and the outer pushout. It can be easily verified that σ^+ indeed acts as described in section 3.2. All these constructions has been given in \mathbf{O} but we have the following property.

Lemma 6.4. *The diagram (1) also exists in \mathbb{O} .*

Now we want to turn the computation of $\delta(O, K, a)$ into a functorial operation on \mathbb{O} . This operation can only have O as parameter. The dependency from a and K is removed by adding a new event for *each* set of independent causes and *each* action. Formally, consider all $K_1: [k_1]_{l_1} \hookrightarrow O, \dots, K_m: [k_m]_{l_m} \hookrightarrow O$. Suppose $Act = \{a_1, \dots, a_n\}$. Then we can compute $\delta(O)$ via the colimit shown in Figure 5. It is the colimit of m cospans with vertex $[k_i]_{l_i}$. Each cospan is similar to the cospan in (1), but its legs include all morphisms $K_i^a: [k_i]_{l_i}^a \rightarrow \delta(O)$, for all $a \in Act$, instead of a single morphism for a given a . This means that, for each set of causes K_i , in $\delta(O)$ we have fresh events labelled by all possible actions.

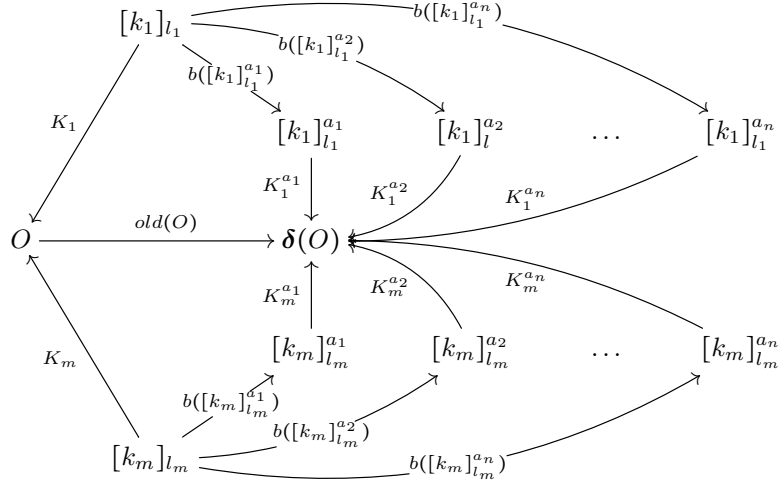


Figure 5: Colimit computing $\delta(O)$.

Notice that $\delta(O)$ and $old(O)$ do not depend on K and a . We can recover *new* maps as follows

$$new(O, K_i, a) = K_i^a \circ \top([k_i]_{l_i}^a): [1]_a \rightarrow \delta(O)$$

Given a morphism $\sigma: O \rightarrow O'$, we denote $\delta(\sigma): \delta(O) \rightarrow \delta(O')$ the corresponding morphism induced by the universal property of the above colimit. Since the colimit in Figure 5 is formed by many diagrams like the inner pushout in (1), by the universal property of pushouts there are unique maps

$$\epsilon(O, K_i, a): \delta(O, K_i, a) \rightarrow \delta(O) .$$

Then we can relate $\delta(O)$ and each $old(O, K_i, a)$

$$old(O) = \epsilon(O, K_i, a) \circ old(O, K_i, a): O \rightarrow \delta(O)$$

and see how each σ^+ “embeds” into $\delta(\sigma)$, namely

$$\begin{array}{ccc} \delta(O, K_i, a) & \xrightarrow{\epsilon(O, K_i, a)} & \delta(O) \\ \sigma^+ \downarrow & & \downarrow \delta(\sigma) \\ \delta(O', \sigma(K_i), a) & \xrightarrow{\epsilon(O', \sigma(K_i), a)} & \delta(O') \end{array}$$

The intuition is that $\delta(\sigma)$ acts as σ on old events (as all σ^+ do) and as the specific σ^+ on new ones. Since each σ^+ is an order-embedding (Lemma 6.4), also $\delta(\sigma)$ is, so $\delta(\sigma)$ is a morphism of \mathbb{O} . This means that δ defines a proper *allocation endofunctor* on \mathbb{O} .

Example 6.5. Suppose $Act = \{c, d\}$ and let O be the discrete abstract poset $\{e_a, e'_b\}$. Then $\delta(O)$

contains $new(O, \emptyset, c)$, $new(O, \emptyset, d)$, and the following pairs (we omit reflexive ones):

$$\begin{array}{ll}
e_a \leq new(O, \{e_a\}, c) & e_a \leq new(O, \{e_a\}, d) \\
e_a \leq new(O, \{e_a, e'_b\}, c) & e_a \leq new(O, \{e_a, e'_b\}, d) \\
e'_b \leq new(O, \{e'_b\}, c) & e'_b \leq new(O, \{e'_b\}, d) \\
e'_b \leq new(O, \{e_a, e'_b\}, c) & e'_b \leq new(O, \{e_a, e'_b\}, d)
\end{array}$$

Remark 6.6. Our definition of δ may not seem the best one, as it generates a new event for each possible set of causes and each label, whereas a transition only generates one of these events. However, having a functor on \mathbb{O} allows us to lift it to presheaves in a way that ensures the existence of both left and right adjoint (giving Kan extensions along δ) for the lifted functor, and then preservation of both limits and colimits, which is essential for coalgebras employing such functor. Generation of unused events is not really an issue: as we will see later, it is always possible to recover the support of a P-marking, i.e., the poset formed by events actually appearing in it.

Now we look at the category $\mathbf{Set}^{\mathbb{O}}$ of presheaves on labelled posets. Since \mathbb{O} is small it follows that $\mathbf{Set}^{\mathbb{O}}$ is locally finitely presentable and has all limits and colimits, in particular products and coproducts. The following functors are relevant for us.

Presheaf of event names. $\mathcal{E}: \mathbb{O} \rightarrow \mathbf{Set}$ maps O to the set $|O|$. Formally

$$\mathcal{E} = \sum_{a \in Act} \text{Hom}_{\mathbb{O}}([1]_a, -)$$

where $e_a \in |O|$ is represented as a morphism $[1]_a \rightarrow O$. The action of \mathcal{E} on a morphism $\sigma: O \rightarrow O'$ gives the function $\lambda e_a \in \mathcal{E}(O). \sigma \circ e_a$, which renames the event e_a according to σ .

Finite powerset. $\mathcal{P}_f: \mathbf{Set}^{\mathbb{O}} \rightarrow \mathbf{Set}^{\mathbb{O}}$, defined as $\mathcal{P}_f \circ (-)$, where \mathcal{P}_f is the finite powerset on \mathbf{Set} .

Event allocation operator. $\Delta: \mathbf{Set}^{\mathbb{O}} \rightarrow \mathbf{Set}^{\mathbb{O}}$, given by $(-) \circ \delta$. Explicitly, for $P: \mathbb{O} \rightarrow \mathbf{Set}$ and $O \in |\mathbb{O}|$, $\Delta P(O) = P(\delta(O))$. Intuitively, it generates causal markings with additional fresh events.

Presheaf of labels. $\mathcal{L}: \mathbb{O} \rightarrow \mathbf{Set}$ given by

$$\mathcal{L}(O) = Act \times \mathcal{P}_f \mathcal{E}(O)$$

For each $O \in |\mathbb{O}|$, this functor gives pairs (a, K) of an action a and a finite set of causes K , selected among events in O .

We use these operators to define our behavioral endofunctor.

Definition 6.7 (Behavioral functor). The behavioral functor $B: \mathbf{Set}^{\mathbb{O}} \rightarrow \mathbf{Set}^{\mathbb{O}}$ is

$$BP = \mathcal{P}_f(\mathcal{L} \times \Delta P) .$$

To understand this definition, consider a B -coalgebra (P, ρ) . Given $O \in |\mathbb{O}|$ and $p \in P(O)$, $\rho_O(p)$ is a finite set of triples (a, K, p') , meaning that p' is the continuation of p after observing $K \vdash a$. The continuation always belongs to $\Delta P(O)$, because every transition allocates a new event.

The category $B\text{-Coalg}$ is well-behaved: it has a final B -coalgebra, and the behavioral equivalence it induces coincides with B -bisimilarity. This is thanks to the following properties.

Proposition 6.8. *B is accessible and covers pullbacks.*

B-coalgebras can be regarded as particular LTSs whose states are elements of presheaves, i.e., pairs $O \triangleright p$.

Definition 6.9 (\mathbb{O} -ILTS). An \mathbb{O} -indexed labelled transition system (\mathbb{O} -ILTS) is a pair (P, \Longrightarrow) of a presheaf $P: \mathbb{O} \rightarrow \mathbf{Set}$ and a finitely-branching transition relation $\Longrightarrow \subseteq \int P \times \int \mathcal{L} \times \int P$ of the form:

$$O \triangleright p \xrightarrow{K \vdash a} \delta(O) \triangleright p' \quad (a, K) \in \mathcal{L}(O)$$

such that, for each morphism $\sigma: O \rightarrow O'$ in \mathbb{O} :

- (i) if $O \triangleright p \xrightarrow{l} \delta(O) \triangleright p'$ then $O' \triangleright p[\sigma] \xrightarrow{l[\sigma]} \delta(O') \triangleright p'[\delta(\sigma)]$ (transitions are *preserved* by σ);
- (ii) if $O' \triangleright p[\sigma] \xrightarrow{l'} \delta(O') \triangleright p''$ then there are l and $\delta(O) \triangleright p''$ such that $l'[\sigma] = l$, $p''[\delta(\sigma)] = p'$ and $O \triangleright p \xrightarrow{l'} \delta(O) \triangleright p''$ (transitions are *reflected* by σ);

Now, notice that labels and continuations of \mathbb{O} -ILTSs agree with those generated by *B*, and (i) and (ii) say that the transition relation behaves like a natural transformation. Therefore we have the following correspondence.

Proposition 6.10. *\mathbb{O} -ILTSs are in bijection with *B*-coalgebras.*

The natural notion of bisimulation for these transition systems is \mathbb{O} -indexed bisimulation.

Definition 6.11 (\mathbb{O} -indexed bisimulation). An \mathbb{O} -indexed bisimulation on a \mathbb{O} -ILTS (P, \Longrightarrow) is an indexed family of relations $\{R_O \subseteq P(O) \times P(O)\}_{O \in |\mathbb{O}|}$ such that, for all $(p, q) \in R_O$:

- (i) if $O \triangleright p \xrightarrow{K \vdash a} \delta(O) \triangleright p'$ then there is $\delta(O) \triangleright q'$ such that $O \triangleright q \xrightarrow{K \vdash a} \delta(O) \triangleright q'$ and $(p', q') \in R_{\delta(O)}$;
- (ii) for all $\sigma: O \rightarrow O'$, $(p, q) \in R_O$ if and only if $(p[\sigma]_P, q[\sigma]_P) \in R_{O'}$.

This definition closely resembles that of **AC**-bisimulations (Definition 3.11). We have an additional condition (ii), requiring closure under morphisms of \mathbb{O} . This is not satisfied by all **AC**-bisimulations, but it holds for the greatest one (Theorem 3.15). We have the following correspondence.

Proposition 6.12. *Let (P, ρ) be a *B*-coalgebra. Then *B*-bisimulations on (P, ρ) are in bijection with \mathbb{O} -indexed bisimulations on the induced \mathbb{O} -ILTS.*

Notice that, unlike Aczel-Mendel bisimulations, a *B*-bisimulation (namely, a Hermida-Jacobs one) needs not be the carrier of a *B*-coalgebra in order to be a bisimulation. This strong requirement is the reason why some \mathbb{O} -indexed bisimulations cannot be turned into Aczel-Mendler ones (see [24, 3.3, Anomaly]).

We now show that $\mathbf{CG}_{\mathbf{AC}}$ can be represented as an \mathbb{O} -ILTS. We form a presheaf from P-markings as follows.

Definition 6.13 (Presheaf of P-markings). The presheaf of P-markings $\mathcal{M}: \mathbb{O} \rightarrow \mathbf{Set}$ is given by

$$\mathcal{M}(O) = \{c \mid O \triangleright c \text{ is an abstract P-marking}\} \quad \mathcal{M}(\sigma: O \rightarrow O') = \lambda(O \triangleright c). O' \triangleright (c\sigma) \downarrow_{O'}$$

The action of \mathcal{M} on morphisms needs to apply the closure operator, after renaming the causal marking: this guarantees that the result is a proper P-marking. The functor \mathcal{M} has the following useful property.

Lemma 6.14. *\mathcal{M} preserves pullbacks.*

Intuitively, thanks to this property, if we take $c \in \mathcal{M}(O)$ and all subsets O' of O such that $\mathcal{M}(O')$ contains a “version” of c (typically with fewer events) then the set obtained by applying \mathcal{M} to the pullback of these subsets, i.e., to their minimal common subset, still contains a version of c . This will be essential, in the next section, to compute minimal representatives of P-markings.

We are ready to translate CG_{AC} to an \mathbb{O} -ILTS.

Definition 6.15 (Causal \mathbb{O} -ILTS_{AC}). The *Causal* \mathbb{O} -ILTS (\mathbb{O} -ILTS_{AC}) $(\mathcal{M}, \Longrightarrow)$ is the smallest one generated by the rule

$$\frac{O \triangleright c \xrightarrow{K+a} \delta(O, K, a) \triangleright c'}{O \triangleright c \xrightarrow{K+a} \delta(O) \triangleright c'[\epsilon(O, K, a)]}$$

This translation does not affect bisimilarities: two states can do the same transitions in CG_{AC} if and only if they can do the same transitions also in \mathbb{O} -ILTS_{AC}; continuations only differ for an order-embedding, but by Theorem 3.15 and Definition 6.11(ii), the \mathbb{O} -indexed bisimilarity and \sim_{AC} are closed under order-embeddings.

We call *causal coalgebra* the B -coalgebra equivalent to $(\mathcal{M}, \Longrightarrow)$. We have the following theorem, which collects the results of this section, instantiated to the causal coalgebra.

Theorem 6.16. *\mathbb{O} -indexed bisimulations on $(\mathcal{M}, \Longrightarrow)$ are equivalent to:*

- *B -bisimulations on the causal coalgebra;*
- *AC-bisimulations closed under order-embeddings.*

In particular, we have that the greatest \mathbb{O} -indexed bisimulation, B -bisimilarity on the causal coalgebra and \sim_{AC} are all equivalent, thanks to Theorem 3.15. These, by Proposition 6.8, are equivalent to the kernel of the unique morphism from the causal coalgebra to the final one.

7. From coalgebras to HD-automata

In order to give a characterization of the causal coalgebra in terms of named sets, we employ the results of [10]. Here authors define a *symmetry group over a category \mathbf{C}* to be a collection of morphisms in $\mathbf{C}[c, c]$, for any $c \in |\mathbf{C}|$, which is a group w.r.t. composition of morphisms. Then they take families of such groups as their notion of generalized named sets. A first result establishes the equivalence between these families and *coproducts of symmetrized representables*, that are functors of the form

$$\sum_{i \in I} \text{Hom}_{\mathbf{C}}(c_i, -) / \Phi_i$$

where Φ_i is a symmetry group over \mathbf{C} with domain c_i , and the quotient identifies morphisms that are obtained one from the other by precomposing elements of Φ_i . These functors, in turn, are shown to be isomorphic to *wide-pullback-preserving* presheaves on \mathbf{C} , a wide pullback being the limit of a diagram with an arbitrary number of morphisms pointing to the same object (pullbacks are a special case, with two such morphisms). The described results are summarized in the following theorem from [10].

Theorem 7.1. *Let \mathbf{C} be a category that is small, has wide pullbacks, and such that all its morphisms are monic and those in $\mathbf{C}[c, c]$ are isomorphisms, for every $c \in |\mathbf{C}|$. Then every wide-pullback-preserving $P \in |\mathbf{Set}^{\mathbf{C}}|$ is equivalent to a coproduct of symmetrized representables.*

Our category \mathbb{O} satisfies the hypothesis of this theorem: it is small and has wide pullbacks due to the existence of pullbacks. In fact, the diagram of a wide pullback in \mathbb{O} is formed by a finite number of morphisms, because a finite poset always has a finite number of ingoing poset-reflecting monomorphisms, so its limit can be computed via binary pullbacks. Moreover, \mathbb{O} has only monos, as order-embeddings are always monic, and $\mathbb{O}[O, O]$ clearly has only isomorphisms, for each $O \in |\mathbb{O}|$. Finally, our presheaf of causal markings \mathcal{M} preserves (wide) pullbacks (Lemma 6.14), so there exists an equivalent coproduct of symmetrized representables.

Theorem 7.1 indeed describes an equivalence between pullback-preserving presheaves and families, which induces one on coalgebras. We shall now investigate this point. Let $\mathbf{Set}_{\diamond}^{\mathbb{O}}$ be the full subcategory of $\mathbf{Set}^{\mathbb{O}}$ formed by pullback-preserving presheaves. We have that our behavioral endofunctor B indeed defines an endofunctor on $\mathbf{Set}_{\diamond}^{\mathbb{O}}$.

Proposition 7.2. *All the endofunctors on $\mathbf{Set}^{\mathbb{O}}$ in Definition 6.7 can be restricted to endofunctors on $\mathbf{Set}_{\diamond}^{\mathbb{O}}$.*

Let $B_{\diamond}: \mathbf{Set}_{\diamond}^{\mathbb{O}} \rightarrow \mathbf{Set}_{\diamond}^{\mathbb{O}}$ be the restricted behavioral endofunctor. The causal coalgebra is clearly a B_{\diamond} -coalgebra. Restricting to $\mathbf{Set}_{\diamond}^{\mathbb{O}}$ does not affect the final coalgebra: B -**Coalg** and B_{\diamond} -**Coalg** have the same final object and final morphisms from common objects. In fact, the terminal sequence starts from the final presheaf 1 , pointwise defined as the singleton set, which trivially preserves pullbacks, and goes through $B^n(1) = B_{\diamond}^n(1)$, for any n .

Corollary 7.3 (of Theorem 7.1). *Let \tilde{B} be the behavioral endofunctor on families defined by lifting all functors in Definition 6.7 along the equivalence. Then the category B_{\diamond} -**Coalg** is equivalent to \tilde{B} -**Coalg**.*

In particular, the equivalence relates the final B_{\diamond} -coalgebra and the final \tilde{B} -coalgebra, and their final morphisms. Moreover, since kernels are preserved by equivalence, identifications made by the final morphisms are preserved, hence behavioral equivalence is preserved too.

Now that we have proved that our categorical setting is suitable for HD-automata, we can translate the causal coalgebra to a HD-automaton. We adopt the definition of HD-automaton given in [11]: a HD-automaton is a(ny) coalgebra over a named set. We introduce a notion of named set closer to a more traditional one, but indeed equivalent to the families mentioned above. Given a set S of morphism and a morphism σ in \mathbb{O} , we write $S \circ \sigma$ for the set $\{\tau \circ \sigma \mid \tau \in S\}$ (analogously for $\sigma \circ S$).

Definition 7.4 (Category $\mathbf{Sym}(\mathbb{O})$). Let $\mathbf{Sym}(\mathbb{O})$ be the category defined as follows:

- objects Φ are subsets of $\mathbb{O}[O, O]$ that are groups w.r.t. composition in \mathbb{O} ;
- morphisms $\Phi_1 \rightarrow \Phi_2$ are sets of morphisms $\sigma \circ \Phi_1$ such that $\sigma: \text{dom}(\Phi_1) \rightarrow \text{dom}(\Phi_2)$ and $\Phi_2 \circ \sigma \subseteq \sigma \circ \Phi_1$.

Definition 7.5 (Category \mathbb{O} -**Set**). The category \mathbb{O} -**Set** is defined as follows:

- objects are \mathbb{O} -named sets, that are pairs $N = (Q_N, \mathbf{G}_N)$ of a set Q_N and a function $\mathbf{G}_N: Q_N \rightarrow |\mathbf{Sym}(\mathbb{O})|$. The local poset of $q \in Q_N$, denoted $\|q\|$, is $\text{dom}(\sigma)$, for any $\sigma \in \mathbf{G}_N(q)$.

- morphisms $f: N \rightarrow M$ are \mathbb{O} -named functions, that are pairs (h, Σ) of a function $h: Q_N \rightarrow Q_M$ and a function Σ mapping each $q \in Q_N$ to a morphism $\mathbb{G}_M(h(q)) \rightarrow \mathbb{G}_N(q)$ in $Sym(\mathbb{O})$.

In the rest of this section we give an explicit description of the \mathbb{O} -named set produced from \mathcal{M} by the equivalence. Its elements will be minimal P-markings with symmetries. We will show that the translation from P-markings to minimal ones with symmetries is achieved via categorical constructions. We need the notions of support, seed and orbit.

Definition 7.6 (Support and seed). Given $O \triangleright c$, its *support*, denoted $supp(c)$, is the wide-pullback-object of the following morphisms

$$\{\sigma: O' \rightarrow O \mid \exists O' \triangleright c' : c'[\sigma] = c\}$$

Let Σ_c be the embedding $supp(c) \hookrightarrow O$ given by the pullback. Then the *seed* of c , denoted $seed(c)$, is the unique element of $\mathcal{M}(supp(c))$ such that $seed(c)[\Sigma_c] = c$.

As shown in [10, 15], preservation of pullbacks by \mathcal{M} is essential to ensure existence and uniqueness of seeds. The seed operation achieves the first two properties of minimal P-markings (see Definition 3.16): $seed(c)$ just contains immediate causes for each token and $supp(c)$ contains all and only those causes. This is illustrated by the following example.

Example 7.7. Consider the following P-marking for the running example

$$\{e_a \leq e'_a, e''_a \leq e'''_a\} \triangleright \{\{e_a, e'_a\} \vdash s_1, \{e''_a, e'''_a\} \vdash s_2\}$$

which is reachable after firing t_1 and t_2 twice. The set of morphisms of Definition 7.6 has four elements

$$f_1, f_2: \{e_a \leq e'_a, e''_a\} \rightarrow \{e_a \leq e'_a, e''_a \leq e'''_a\} \quad f_3, f_4: \{e_a, e'_a\} \rightarrow \{e_a \leq e'_a, e''_a \leq e'''_a\}$$

$$f_1 = \begin{cases} e_a \mapsto e_a \\ e'_a \mapsto e'_a \\ e''_a \mapsto e'''_a \end{cases} \quad f_2 = \begin{cases} e_a \mapsto e''_a \\ e'_a \mapsto e'''_a \\ e''_a \mapsto e'_a \end{cases} \quad f_3 = \begin{cases} e_a \mapsto e'_a \\ e'_a \mapsto e'''_a \end{cases} \quad f_4 = \begin{cases} e_a \mapsto e'''_a \\ e'_a \mapsto e'_a \end{cases}$$

In fact, we have

$$\left. \begin{array}{l} (\{e_a \leq e'_a, e''_a\} \triangleright \{\{e_a, e'_a\} \vdash s_1, \{e''_a\} \vdash s_2\}) [f_1] \\ (\{e_a \leq e'_a, e''_a\} \triangleright \{\{e''_a\} \vdash s_1, \{e_a, e'_a\} \vdash s_2\}) [f_2] \\ (\{e_a, e'_a\} \triangleright \{\{e_a\} \vdash s_1, \{e'_a\} \vdash s_2\}) [f_3] \\ (\{e_a, e'_a\} \triangleright \{\{e'_a\} \vdash s_1, \{e_a\} \vdash s_2\}) [f_4] \end{array} \right\} = \{e_a \leq e'_a, e''_a \leq e'''_a\} \triangleright \{\{e_a, e'_a\} \vdash s_1, \{e''_a, e'''_a\} \vdash s_2\}$$

Recall that each $[f_i] = \mathcal{M}(f_i)$ is a function that, when applied to a P-marking, replaces events according to f_i and then down-closes the result w.r.t. $\{e_a \leq e'_a, e''_a \leq e'''_a\}$. It is easy to check that the pullback object of all four morphisms is $\{e_a, e'_a\}$, so the corresponding seed is

$$\{e_a, e'_a\} \triangleright \{\{e_a\} \vdash s_1, \{e'_a\} \vdash s_2\}.$$

Notice that two events have been discarded, because they are not immediate causes.

Definition 7.8 (Orbit). The *orbit* of $O \triangleright c$ is

$$\text{orb}(c) = \{c[\sigma] \mid \sigma \in \mathbb{O}[O, O]\}$$

We denote by $[c]^\circ$ a canonical choice of an element of $\text{orb}(c)$.

The *orbit* of c is the set of causal markings obtained by applying to c all functions induced by poset automorphisms. Automorphisms are isomorphisms, so taking a canonical representative for this orbit achieves the third requirement of minimal P-markings: it amounts to applying the operation $[O \triangleright c]_\pm$, i.e., choosing a representative of isomorphism classes for $O \triangleright c$.

Definition 7.9. The \mathbb{O} -named set of minimal P-markings is (M, \mathbf{G}_M) , where

$$\begin{aligned} M &= \{ \text{supp}(c) \blacktriangleright [\text{seed}(c)]^\circ \mid O \triangleright c \in \int \mathcal{M} \} \\ \mathbf{G}_M &= \lambda O \blacktriangleright c. \{ \Phi \in |\mathbf{Sym}(\mathbb{O})| \mid \text{dom}(\Phi) = O \wedge \forall \sigma \in \Phi : c[\sigma] = c \} \end{aligned}$$

The set M is produced from elements of \mathcal{M} : for each of these, we compute the seed, and then we only take the canonical representative for the seed's orbit. As explained, the final result is indeed a minimal P-marking $O \blacktriangleright c$. This P-marking is associated a symmetry by \mathbf{G}_M , namely $\Phi = \mathbf{G}_M(O \blacktriangleright c)$, so it becomes the P-marking with symmetry $O \blacktriangleright_\Phi c$.

The derivation of an HD-automaton on (M, \mathbf{G}_M) in \tilde{B} -**Coalg** from the causal coalgebra, along the equivalence, is the category-theoretic counterpart of the derivation of CG_{ICS} from CG_{AC} . The correspondence between CG_{ICS} and coalgebras over named sets is analogous to the π -calculus case, where we have set-theoretical HD-automata on one side [19] and categorical ones, namely coalgebras over named sets, on the other side. The correspondence for the π -calculus has been worked out in [9, 11], and the theory introduced therein seems robust enough to accommodate different notions of named sets such as ours. In particular, functors used to define coalgebras over named sets, such as powerset and allocation functors, should be very similar to those defining \tilde{B} .

We briefly illustrate the \tilde{B} -coalgebra for the running example. The \mathbb{O} -named set (M, \mathbf{G}_M) is as follows: M includes all P-markings in Figure 3, and \mathbf{G}_M returns the symmetry $\{id\}$ for each of them. Transitions are represented as a \mathbb{O} -named function $(h, \Sigma) : (M, \mathbf{G}_M) \rightarrow \tilde{B}(M, \mathbf{G}_M)$, where h maps each state $O \blacktriangleright_{\{id\}} c$ to its label and continuation, and $\Sigma(O \blacktriangleright_{\{id\}} c)$ encodes all history maps for outgoing transitions.

We leave a deeper investigation of the category of \mathbb{O} -named sets and of \tilde{B} -coalgebras for future work.

8. Conclusions

In this paper we have introduced an approach to derive compact operational models for causality in Petri nets. In order to do this, we have constructed a labelled semantics of Petri nets in terms of causal case graphs, and we have given a procedure to refine them in order to get minimal, possibly finite-state, representations. We have then modeled causal case graphs in a categorical setting, exploiting a nominal representation of causal relations: they are modeled as posets over event names with action labels. Our categorical treatment is simpler and more natural than the set-theoretic one, and employs standard constructs and results for nominal calculi, namely presheaf-based coalgebras and their equivalence with HD-automata. In particular, reducing the state-space and showing that this operation preserves the semantics require some technical effort in the set-theoretic version,

whereas the categorical version employs a general construction that automatically performs this reduction in a semantics-preserving way.

Our approach has a practical significance: we show how to synthesize HD-automata from Petri nets, and how to compute minimal realizations for them, in order to detect bisimilar states. As mentioned, minimization of HD-automata is possible in many cases. Even if our approach does not actually provide a way to minimize nets themselves, one can still decide bisimilarity of markings by minimizing their reachable HD-automata and matching the results.

Finally, our contribution is also methodological: we provide a further example in which the presheaf/HD-automata framework is successfully applied. We emphasize that this framework is highly parametric and can possibly be useful in many other cases.

8.1. Related work

This paper follows a line of research on coalgebraic models of causality, started in [7] by the same authors. The categorical machinery is the same in both papers, namely presheaf-based coalgebras, HD-automata, and the equivalence among them. However, this paper takes a further step towards a general categorical theory of causality. In [7], in fact, we have provided models for a particular class of causal LTSs, namely Degano-Darondeau ones. In this paper, instead, we treat Petri nets, which are much more general. For instance, unlike Degano-Darondeau LTSs, Petri nets can describe synchronizations of more than two processes.

In [7] we start from existing set-theoretic models, similar to abstract CGs, whereas the models we introduce here are novel. In both papers we represent causal dependencies as posets over events, but in [7] events are unlabelled and are canonically represented as natural numbers. Here we have labels and we take a more general approach: instead of choosing specific representatives of events, we make abstract CGs parametric in this choice. This requires more technical work and it further validates the categorical approach, where book-keeping details are abstracted away. The categorical environment in this paper is more elaborate than [7], due to labelling. In particular, event generation is more complex, and is studied in greater detail. Another difference is that here we give conditions under which the model with only immediate causes is finite, whereas in [7] decidability is not treated.

A first version of HD-automata for Petri nets, called *causal automata*, has been introduced in [18]. However, their construction is purely set-theoretical and does not include symmetries, so the existence of a minimal model is not guaranteed. This version of HD-automata is similar to what we call immediate causes CG (without symmetries). HD-automata with symmetries were developed for the π -calculus in [21, 19], and a general categorical treatment was provided in [11]. In all these cases nominal structures associated to states are just a sets of (event) names, whereas we have posets, which are more adequate to represent causal dependencies.

We can cite [8] for the introduction of transitions systems for causality whose states are elements of presheaves, intended to model the causal semantics of the π -calculus as defined in [6]. However, the index of a state is a set of names, without any information about events and causal relations. The advantage of our index category is that it allows reducing the state-space in an automatic way, exploiting a standard categorical construction. This cannot be done in the framework of [8]. Finally, an HD-automaton for causality has been described in [11], but it is derived as a direct translation of causal automata and its states do not take into account causal relations.

Other related works are [26, 28], where event structures have been characterized as (contravariant) presheaves on posets. While the meaning of presheaves is similar, the context is different: we consider the more concrete realm of coalgebras and nominal automata. A more precise correspondence with such models should be worked out.

8.2. Future work

Logics for causality have been recently studied in [5]. As future work, we would like to understand whether they can be captured in our coalgebraic setting. Another open research question is how to obtain coalgebraic models for other notions of causal bisimulation, such as hereditary history preserving bisimulation.

References

- [1] Jiří Adámek. Introduction to coalgebra. *Theory and Applications of Categories*, 14(8):157–199, 2005.
- [2] Jiří Adámek, Filippo Bonchi, Mathias Hülsbusch, Barbara König, Stefan Milius, and Alexandra Silva. A coalgebraic perspective on minimization and determinization. In *FoSSaCS*, pages 58–73, 2012.
- [3] Jiří Adámek and Jiří Rosický. *Locally Presentable and Accessible Categories*. Cambridge University Press, 1994.
- [4] Paolo Baldan and Alberto Carraro. Non-interference by unfolding. In *Application and Theory of Petri Nets and Concurrency - 35th International Conference, PETRI NETS 2014, Tunis, Tunisia, June 23-27, 2014. Proceedings*, pages 190–209, 2014.
- [5] Paolo Baldan and Silvia Crafa. A logic for true concurrency. *J. ACM*, 61(4):24, 2014.
- [6] Michele Boreale and Davide Sangiorgi. A fully abstract semantics for causality in the π -calculus. *Acta Inf.*, 35(5):353–400, 1998.
- [7] Roberto Bruni, Ugo Montanari, and Matteo Sammartino. Revisiting causality, coalgebraically. *Acta Inf.*, pages 1–29, 2014. Preprint available at <http://www.di.unipi.it/~sammarti/publications/ACTA2014.pdf>.
- [8] Gian Luca Cattani and Peter Sewell. Models for name-passing processes: interleaving and causal. *Inf. Comput.*, 190(2):136–178, 2004.
- [9] Vincenzo Ciancia. *Accessible Functors and Final Coalgebras for Named Sets*. PhD thesis, University of Pisa, 2008.
- [10] Vincenzo Ciancia, Alexander Kurz, and Ugo Montanari. Families of symmetries as efficient models of resource binding. *Electr. Notes Theor. Comput. Sci.*, 264(2):63–81, 2010.
- [11] Vincenzo Ciancia and Ugo Montanari. Symmetries, local names and dynamic (de)-allocation of names. *Inf. Comput.*, 208(12):1349 – 1367, 2010.
- [12] Gian Luigi Ferrari, Ugo Montanari, and Emilio Tuosto. Coalgebraic minimization of HD-automata for the π -calculus using polymorphic types. *Theor. Comput. Sci.*, 331(2-3):325–365, 2005.
- [13] Marcelo P. Fiore and Daniele Turi. Semantics of name and value passing. In *LICS*, pages 93–104, 2001.

- [14] Sibylle B. Fröschle and Thomas T. Hildebrandt. On plain and hereditary history-preserving bisimulation. In *MFCS*, pages 354–365, 1999.
- [15] Fabio Gadducci, Marino Miculan, and Ugo Montanari. About permutation algebras, (pre)sheaves and named sets. *Higher-Order and Symbolic Computation*, 19(2-3):283–304, 2006.
- [16] André Joyal, Mogens Nielsen, and Glynn Winskel. Bisimulation from open maps. *Inf. Comput.*, 127(2):164–185, 1996.
- [17] Paris C. Kanellakis and Scott A. Smolka. Ccs expressions, finite state processes, and three problems of equivalence. *Inf. Comput.*, 86(1):43–68, 1990.
- [18] Ugo Montanari and Marco Pistore. Minimal transition systems for history-preserving bisimulation. In *STACS*, pages 413–425, 1997.
- [19] Ugo Montanari and Marco Pistore. Structured coalgebras and minimal HD-automata for the π -calculus. *Theor. Comput. Sci.*, 340(3):539–576, 2005.
- [20] Mogens Nielsen, Gordon D. Plotkin, and Glynn Winskel. Petri nets, event structures and domains, part I. *Theor. Comput. Sci.*, 13:85–108, 1981.
- [21] Marco Pistore. *History Dependent Automata*. PhD thesis, University of Pisa, 1999.
- [22] Jan J. M. M. Rutten. Universal coalgebra: a theory of systems. *Theor. Comput. Sci.*, 249(1):3–80, 2000.
- [23] Davide Sangiorgi and David Walker. *π -Calculus: A Theory of Mobile Processes*. Cambridge University Press, New York, NY, USA, 2001.
- [24] Sam Staton. Name-passing process calculi: operational models and structural operational semantics. Technical Report 688, University of Cambridge, 2007.
- [25] Sam Staton. Relating coalgebraic notions of bisimulation. *Logical Methods in Computer Science*, 7(1), 2011.
- [26] Sam Staton and Glynn Winskel. On the expressivity of symmetry in event structures. In *LICS*, pages 392–401, 2010.
- [27] Boris A. Trakhtenbrot and Alexander Moshe Rabinovich. Behavior structures and nets of processes. *Fundamenta Informaticae*, 11(4):357–403, 1988.
- [28] Glynn Winskel. Event structures as presheaves — two representation theorems. In *CONCUR*, pages 541–556, 1999.
- [29] James Worrell. Terminal sequences for accessible endofunctors. *Electr. Notes Theor. Comput. Sci.*, 19:24–38, 1999.

A. Proofs

We first introduce some technical lemmata. Then we give proofs for the claims in the paper.

A.1. Additional lemmata

Lemma A.1. *Let O_1, O_2 be finite Act-labeled posets and let $\sigma: O_1 \rightarrow O_2$ be an order-embedding. Then:*

- (i) $O_1 \triangleright c \xrightarrow{K \vdash e_a} \delta(O_1, K, e_a) \triangleright c'$ implies $O_2 \triangleright (c\sigma) \downarrow_{O_2} \xrightarrow{\sigma(K) \vdash e'_a} O'_2 \triangleright (c'\sigma[e'_a/e_a]) \downarrow_{O'_2}$, for any $e' \notin X_{O_2}$, with $O'_2 = \delta(O_2, \sigma(K), e'_a)$;
- (ii) $O_2 \triangleright c \xrightarrow{K \vdash e_a} \delta(O_2, K, e_a) \triangleright c'$ implies $O_1 \triangleright c'' \xrightarrow{K' \vdash e'_a} \delta(O_1, K', e'_a) \triangleright c'''$, with $c''\sigma = c$, $\sigma(K') = K$ and $c'''\sigma[e'_a/e_a] = c'$, for any $e' \notin X_{O_1}$.

Proof. We prove item (i), the other one is analogous. Suppose $O_1 \triangleright c \xrightarrow{K \vdash e_a} \delta(O_1, K, e_a) \triangleright c'$ is derived from the rule of Definition 3.3 as follows

$$\frac{t \in T \quad |c_1| = \bullet t \quad a = l(t) \quad e \notin X_{O_1} \quad K = \max_{O_1} \mathcal{K}(c_1)}{O_1 \triangleright c_1 \cup c_2 \xrightarrow{K \vdash e_a} \delta(O_1, K, e_a) \triangleright (\mathcal{K}(c_1) \cup \{e_a\} \vdash t^\bullet) \cup c_2}$$

where $c = c_1 \cup c_2$ and $c' = (\mathcal{K}(c_1) \cup \{e_a\} \vdash t^\bullet) \cup c_2$. Clearly we have $(c\sigma) \downarrow_{O_2} = (c_1\sigma) \downarrow_{O_2} \cup (c_2\sigma) \downarrow_{O_2}$, with $|(c_1\sigma) \downarrow_{O_2}| = |c_1|$, because σ only affects events, not tokens. Moreover, it can be easily verified that $\max_{O_2} \mathcal{K}((c_1\sigma) \downarrow_{O_2}) = \sigma(\max_{O_1} \mathcal{K}(c_1)) = \sigma(K)$. In fact, causes of $(c_1\sigma) \downarrow_{O_2}$ are: those of $c_1\sigma$, related exactly as their counterimages, due to σ preserving and reflecting order; additional causes, smaller than those of $c_1\sigma$, added by the closure. Therefore we can again apply the rule as follows

$$\frac{t \in T \quad |(c_1\sigma) \downarrow_{O_2}| = \bullet t \quad a = l(t) \quad e' \notin X_{O_2} \quad \sigma(K) = \max_{O_2} \mathcal{K}((c_1\sigma) \downarrow_{O_2})}{O_2 \triangleright (c_1\sigma) \downarrow_{O_2} \cup (c_2\sigma) \downarrow_{O_2} \xrightarrow{\sigma(K) \vdash e'_a} O'_2 \triangleright (\mathcal{K}((c_1\sigma) \downarrow_{O_2}) \cup \{e'_a\} \vdash t^\bullet) \cup (c_2\sigma) \downarrow_{O_2}}$$

where $O'_2 = \delta(O_2, \sigma(K), e'_a)$. Now, observe that, by definition of δ , we have

$$\mathcal{K}((c_1\sigma) \downarrow_{O_2}) \subseteq \mathcal{K}((c_1\sigma) \downarrow_{O'_2}) \quad \{e'_a\} \downarrow_{O'_2} = \mathcal{K}((c_1\sigma) \downarrow_{O'_2}) \cup \{e'_a\}$$

which implies

$$\begin{aligned} \mathcal{K}((c_1\sigma) \downarrow_{O_2}) \cup \{e'_a\} \vdash t^\bullet &= (\mathcal{K}(c_1\sigma) \cup \{e'_a\}) \downarrow_{O'_2} \vdash t^\bullet \\ &= (\mathcal{K}(c_1) \cup \{e_a\}) \sigma[e'_a/e_a] \downarrow_{O'_2} \vdash t^\bullet \\ &= (\mathcal{K}(c_1) \cup \{e_a\} \vdash t^\bullet) \sigma[e'_a/e_a] \downarrow_{O'_2} \end{aligned}$$

From this equation, and from $(c_2\sigma) \downarrow_{O_2} = (c_2\sigma[e'_a/e_a]) \downarrow_{O'_2}$, because $e_a \notin \mathcal{K}(c_2)$, it follows that the continuation derived from the above rule has the required shape. \square

Lemma A.2. *Let $\sigma: O \rightarrow O'$ be an isomorphism. Then $O \triangleright c_1 \sim_C O \triangleright c_1$ implies $O' \triangleright c_1\sigma \sim_C O' \triangleright c_2\sigma$.*

Proof. We will prove that the following relation is a C-bisimulation

$$R_{O'} = \{(O' \triangleright c_1\sigma, O' \triangleright c_2\sigma) \mid O \triangleright c_1 \sim_C O \triangleright c_2, \sigma: O \rightarrow O' \text{ is an isomorphism}\}$$

Take $(O' \triangleright c_1\sigma, O' \triangleright c_2\sigma) \in R_{O'}$ and

$$O' \triangleright c_1\sigma \xrightarrow{K' \vdash e'_a} \delta(O', K', e'_a) \triangleright c'_1$$

We have to find a simulating transition of $O' \triangleright c_2\sigma$. Let $e \notin X_O$. We can apply Lemma 3.9, using the isomorphism $\sigma^{-1}[e_a/e'_a]$, and get

$$O \triangleright c_1 \xrightarrow{\sigma^{-1}(K) \vdash e_a} \delta(O, \sigma^{-1}(K), e_a) \triangleright c'_1 \sigma^{-1}[e_a/e'_a]$$

Since $O \triangleright c_1 \sim_{\mathcal{C}} O \triangleright c_2$, there is a simulating transition

$$O \triangleright c_2 \xrightarrow{\sigma^{-1}(K) \vdash e_a} \delta(O, \sigma^{-1}(K), e_a) \triangleright c'_2 .$$

Applying again Lemma 3.9 with $\sigma[e'_a/e_a]$ to this transition, we get

$$O' \triangleright c_2\sigma \xrightarrow{K' \vdash e'_a} \delta(O', K', e'_a) \triangleright c'_2 \sigma[e'_a/e_a] .$$

This is the required simulating transition. In fact, since

$$\delta(O, \sigma^{-1}(K), e_a) \triangleright c'_1 \sigma^{-1}[e_a/e'_a] \sim_{\mathcal{C}} \delta(O, \sigma^{-1}(K), e_a) \triangleright c'_2$$

and $\sigma[e'_a/e_a]$ is an isomorphism, by definition of $R_{O'}$ we have

$$(\delta(O', K', e'_a) \triangleright c'_1 , \delta(O', K', e'_a) \triangleright c'_2 \sigma[e'_a/e_a]) \in R_{O'} .$$

□

Lemma A.3. *Let $O \triangleright c_1$ and $O \triangleright c_2$ be abstract P-markings. Then $O \triangleright c_1 \sim_{\mathcal{C}} O \triangleright c_2$ if and only if $O \triangleright c_1 \sim_{\text{AC}} O \triangleright c_2$.*

Proof. We show the left-to-right implication, the other one is analogous. We prove that the following relation is an AC-bisimulation

$$R_O = \{(O \triangleright c_1, O \triangleright c_2) \mid O \triangleright c_1 \sim_{\mathcal{C}} O \triangleright c_2\}$$

Take $(O \triangleright c_1, O \triangleright c_2) \in R_O$ and suppose

$$O \triangleright c_1 \xrightarrow{K \vdash a} \delta(O, K, a) \triangleright c'_1$$

then we must find a simulating transition of $O \triangleright c_2$. By Definition 3.8, the above transition can be derived from

$$O \triangleright c_1 \xrightarrow{K \vdash e_a} \delta(O, K, e_a) \triangleright c''_1$$

with $c''_1 \text{old}(O, K, e_a)[\text{new}(O, K, e_a)/e_a] = c'_1$. Since $O \triangleright c_1 \sim_{\mathcal{C}} O \triangleright c_2$ by hypothesis, this transition can be simulated by

$$O \triangleright c_2 \xrightarrow{K \vdash e_a} \delta(O, K, e_a) \triangleright c''_2 .$$

Applying again Definition 3.8, we get the required transition

$$O \triangleright c_2 \xrightarrow{K \vdash a} \delta(O, K, a) \triangleright c''_2 (\text{old}(O, K, e_a)[\text{new}(O, K, e_a)/e_a]).$$

In fact, from $\delta(O, K, e_a) \triangleright c''_1 \sim_{\mathcal{C}} \delta(O, K, e_a) \triangleright c''_2$, using Lemma A.2 with the isomorphism $\text{old}(O, K, e_a)[\text{new}(O, K, e_a)/e_a]$, we get

$$\delta(O, K, a) \triangleright c'_1 \sim_{\mathcal{C}} \delta(O, K, a) \triangleright c''_2 (\text{old}(O, K, e_a)[\text{new}(O, K, e_a)/e_a])$$

and we can conclude that these P-markings are related by $R_{\delta(O, K, a)}$, by its definition. □

Lemma A.4. Let $O_2 \xleftarrow{\sigma_2} O \xrightarrow{\sigma_1} O_1$ be a span in \mathbb{O} and let

$$\begin{array}{ccc} O & \xrightarrow{\sigma_1} & O_1 \\ \sigma_2 \downarrow & & \downarrow p_1 \\ O_2 & \xrightarrow{p_2} & O_3 \end{array} \quad \sqcap$$

be its pushout in \mathbf{P} . Then it is also a pushout in \mathbf{O} , with

$$l_{O_3}(x) = \begin{cases} l_{O_1}(y) & x = p_1(y) \\ l_{O_2}(y) & x = p_2(y) \end{cases}$$

Proof. In [7, Lemma 8] we have proved that pushouts in \mathbf{P} are computed as in **Graph**, plus transitive closure of the pushout object. We will use this fact to prove our claim.

First of all, we check that l_{O_3} is well-defined. We only have to verify that its definition is correct for $x = p_1(y_1) = p_2(y_2)$. If $p_1(y_1) = p_2(y_2)$ then y_1 and y_2 are images via σ_1 and σ_2 of the same element of O , by definition of pushout in **Graph**. Since σ_1 and σ_2 preserve labels, we must have $l_{O_1}(y_1) = l_{O_2}(y_2)$, so $l_{O_3}(x)$ is well-defined on x .

Preservation of labels by p_1 and p_2 follows immediately from the definition of l_{O_3} .

Now we prove that the square is indeed a pushout in \mathbf{O} . Consider the following situation:

$$\begin{array}{ccc} O & \xrightarrow{\sigma_1} & O_1 \\ \sigma_2 \downarrow & & \downarrow p_1 \\ O_2 & \xrightarrow{p_2} & O_3 \end{array} \quad \sqcap \quad \begin{array}{c} \downarrow q_1 \\ O_4 \end{array}$$

$\downarrow q_2$

$\dots m \dots$

We have to check that, when q_1 and q_2 preserve labels, also the unique mediating morphism m , as computed in \mathbf{P} , does. We prove it by contradiction. Suppose m does not preserve labels, then there exists $x \in X_{O_3}$ such that $l_{O_4}(m(x)) \neq l_{O_3}(x)$. Suppose x is image of $y \in X_{O_1}$ via p_1 (the case $y \in X_{O_2}$ and $x = p_2(y)$ is analogous). Then we have

$$\begin{aligned} l_{O_1}(y) &= l_{O_3}(x) && \text{(by } p_1 \text{ preserving labels)} \\ &\neq l_{O_4}(m(x)) && \text{(by hypothesis)} \\ &= l_{O_4}(q_1(y)) && \text{(by } q_1 = m \circ p_1) \end{aligned}$$

which implies that q_1 does not preserve labels, a contradiction. □

A.2. Main proofs

Proof of Lemma 3.9. It is just a corollary of Lemma A.1. □

Proof of Proposition 3.13. We prove (i), the other point is similar. Suppose

$$O \triangleright c \xrightarrow{K \vdash a} \delta(O, K, a) \triangleright c'$$

Then, by Definition 3.8, this transition can be derived from

$$O \triangleright c \xrightarrow{K \vdash e_a} \delta(O, K, e_a) \triangleright c''$$

with $c' = c'' \text{old}(O, K, e_a)[\text{new}(O, K, e_a)/e_a]$, for any $e \notin X_O$. Suppose $e \notin X_{O'}$. By Lemma A.1(i), we have

$$O' \triangleright c\sigma \xrightarrow{\sigma(K) \vdash e_a} \delta(O', \sigma(K), e_a) \triangleright (c''\sigma[e_a/e_a]) \downarrow_{\delta(O', \sigma(K), e_a)}$$

from which, using Definition 3.8, we get

$$O' \triangleright c\sigma \xrightarrow{\sigma(K) \vdash a} \delta(O', \sigma(K), a) \triangleright (c''\sigma[e_a/e_a]) \downarrow_{\delta(O', \sigma(K), e_a)} \omega$$

where $\omega = \text{old}(O', \sigma(K), e_a)[\text{new}(O', \sigma(K), e_a)/e_a]$. We have to prove that the continuation of this transition has the required form.

It is immediate to verify that, for any isomorphism $\sigma: O \rightarrow O'$ and causal marking c such that $\mathcal{K}(c) \subseteq |O|$, we have

$$(c\sigma) \downarrow_{O'} = c \downarrow_O \sigma$$

which, for $\sigma = \omega$, implies

$$(c''\sigma[e_a/e_a]) \downarrow_{\delta(O', \sigma(K), e_a)} \omega = (c''\sigma[e_a/e_a] \omega) \downarrow_{\delta(O', \sigma(K), e_a)} \cdot \quad (\text{A.1})$$

Now, observe that, by the definition of σ^+ we have

$$\sigma[e_a/e_a] \omega = \text{old}(O, K, e_a)[\text{new}(O, K, e_a)/e_a] \sigma^+$$

therefore (A.1) is equal to

$$(c'' \text{old}(O, K, e_a)[\text{new}(O, K, e_a)/e_a] \sigma^+) \downarrow_{\delta(O', \sigma(K), a)} = (c'\sigma^+) \downarrow_{\delta(O', \sigma(K), a)}$$

as required. \square

Proof of Theorem 3.12. Both implications can be proved by combining Lemma A.3 and Lemma A.2. \square

Proof of Theorem 3.22. This is proved as [7, Theorem 2], where specific choices for abstract posets and *old* and *new* maps are made in order to accommodate Darondeau-Degano LTSs. The proof is exactly the same, where each specific operation is replaced by its general version described in this paper. \square

Proof of Proposition 3.18. Take $c \in [n_0]$. Then its tokens have been created by at most $|c|$ transitions. Since we only take immediate causes, i.e., events generated when those transitions were fired, each $O \blacktriangleright c$ is such that $|O|$ contains at most $|c|$ events. O can be any poset on those events but, since posets of minimal P-markings must be abstract, there are finitely-many such posets. \square

Proof of Lemma 4.4.

- (i) Immediately from the fact that any path from $\emptyset \triangleright \emptyset \vdash m_0$ to $O_c \triangleright c$ builds O_c and c incrementally, adding one event for each transition.
- (ii) Suppose there are two parallel transitions from $O \triangleright c$ to $O' \triangleright c'$, with labels a and b . Then $O' = \delta(O, K, e_a) = \delta(O, K', e'_b)$, which can only happen when $K = K'$ and $e_a = e'_b$, i.e., when the two transitions coincide.

Suppose there is a directed cycle starting and ending at $O \triangleright c$. Each transition in the cycle would add a new event to O , so the final state would be $O' \triangleright c$, with O' a strict superposet of O , a contradiction.

□

Proof of Theorem 4.6.

- (i) Consider a transition $c_1 \xrightarrow{a}_{B_c} c'_1$ and suppose the corresponding transition in CG_c^r is

$$O_{c_1} \triangleright c_1 \xrightarrow{K \vdash e_a} \delta(O_{c_1}, K, e_a) \triangleright c'_1$$

Now, observe that there is a trivial embedding of O_{c_1} into O . In fact, causes of c_1 are down-closed w.r.t. both posets, so O_{c_1} must be a prefix of O . Then, using Lemma A.1(i) and the embedding $O_{c_1} \hookrightarrow O$ on the above transition, we get

$$O \triangleright c_1 \xrightarrow{K \vdash e'_a} \delta(O, K, e'_a) \triangleright c'_1[e'_a/e_a]$$

for any $e' \notin X_O$. By the hypothesis $O \triangleright c_1 \sim_c O \triangleright c_2$, this transition can be simulated by

$$O \triangleright c_2 \xrightarrow{K \vdash e'_a} \delta(O, K, e'_a) \triangleright c'_2$$

with $\delta(O, K, e'_a) \triangleright c'_1[e'_a/e_a] \sim_c \delta(O, K, e'_a) \triangleright c'_2$. Using Lemma A.1(ii) on the embedding of O_{c_2} into O , and noting that $e' \notin X_{O_{c_2}}$, we recover a transition

$$O_{c_2} \triangleright c_2 \xrightarrow{K \vdash e'_a} \delta(O_{c_2}, K, e'_a) \triangleright c'_2$$

and from this, using the rule in Proposition 4.5, we get $c_2 \xrightarrow{a}_{B_c} c'_2$. In order to show that this transition simulates $c_1 \xrightarrow{a}_{B_c} c'_1$, we have to find an isomorphism $\sigma': O_{c'_1} \rightarrow O_{c'_2}$ such that the following diagram commutes

$$\begin{array}{ccc} O_{c_1} & \xrightarrow{\phi_{c_1, c'_1}^c} & O_{c'_1} \\ \sigma \downarrow & & \downarrow \sigma' \\ O_{c_2} & \xrightarrow{\phi_{c_2, c'_2}^c} & O_{c'_2} \end{array}$$

We can define $\sigma'(x)$ as $\sigma(x)$ if $x \in |O_{c_1}|$ and as e'_a if $x = e_a$.

(ii) We want to prove that the following relation is an AC-bisimulation

$$R_{O_{c_2}} = \{(O_{c_2} \triangleright c_1 \sigma, O \triangleright c_2) \mid c_1 \sim_{bs}^\sigma c_2\}$$

Suppose $c_1 \sim_{bs}^\sigma c_2$ and

$$O_{c_2} \triangleright c_1 \sigma \xrightarrow{K \vdash e_a} \delta(O_{c_2}, K, e_a) \triangleright c'_1. \quad (\text{A.2})$$

We have to find a simulating transition of $O_{c_2} \triangleright c_2$. Applying Lemma 3.9 to the last transition, with isomorphism σ^{-1} , we get

$$O_{c_1} \triangleright c_1 \xrightarrow{\sigma^{-1}(K) \vdash e'_a} \delta(O_{c_1}, \sigma^{-1}(K), e'_a) \triangleright c''_1$$

where $c''_1 = c'_1 \sigma^{-1}[e'_a/e_a]$, for any $e' \notin X_{O_{c_1}}$. This transition corresponds, via Proposition 4.5, to the following transition in B_c

$$c_1 \xrightarrow{a}_{B_c} c''_1$$

which, by the hypothesis $c_1 \sim_{bs}^\sigma c_2$, can be simulated by

$$c_2 \xrightarrow{a}_{B_c} c'_2 \quad (\text{A.3})$$

with $c'_1 \sim_{bs}^{\sigma'} c''_1$ such that

$$\phi_{c_2, c'_2}^c \circ \sigma = \sigma' \circ \phi_{c_1, c''_1}^c \quad (\text{A.4})$$

Now, suppose for simplicity $\{e_a\} = |O_{c'_2}| \setminus |O_{c_2}|$ (the general case where $|O_{c'_2}| \setminus |O_{c_2}|$ contains any event fresh w.r.t. O_{c_2} requires minor changes). By definition of ϕ_{c_2, c'_2}^c and ϕ_{c_1, c''_1}^c , and by (A.4), σ' should act as σ on O_{c_1} , so $\sigma' = \sigma[e_a/e'_a]$. Moreover, since σ' is an isomorphism, we have that the maximal causes of e'_a , namely $\sigma^{-1}(K)$, are mapped by σ' to the maximal causes of e''_a , which then are $\sigma'(\sigma^{-1}(K)) = \sigma(\sigma^{-1}(K)) = K$, where the first equation follows from $e'_a \notin \sigma^{-1}(K)$. Therefore $O_{c'_2} = \delta(O_{c_2}, K, e_a)$ and (A.3) is derived, using Proposition 4.5, from

$$O_{c_2} \triangleright c_2 \xrightarrow{K \vdash e_a} \delta(O_{c_2}, K, e_a) \triangleright c'_2$$

This transition is the required one simulating (A.2). In fact, $c''_1 \sim_{bs}^{\sigma'} c'_2$ implies

$$(\delta(O_{c_2}, K, e_a) \triangleright c''_1 \sigma', \delta(O_{c_2}, K, e_a) \triangleright c'_2) \in R_{\delta(O_{c_2}, K, e_a)}$$

by definition of R , and for the first P-marking we have $c''_1 \sigma' = c'_1 \sigma[e_a/e'_a] = (c'_1 \sigma^{-1}[e'_a/e_a]) \sigma[e_a/e_a] = c'_1$, which is the causal marking in the continuation of (A.2). \square

Proof of Proposition 6.3. Smallness follows from skeletality. In [7] we have proved that pullbacks in \mathbf{P}_m are computed as the category **Graph** of graphs and their homomorphisms. It can be easily verified that, given a cospan $O_1 \xrightarrow{f} O_3 \xleftarrow{g} O_2$ in \mathbb{O} , we can forget labels and compute the pullback as in **Graph**. In fact, the pullback poset O has an element y for each pair of elements $x_1 \in X_{O_1}$ and $x_2 \in X_{O_2}$ such that $f(x_1) = g(x_2)$. But then, since f and g preserve labels, we must have $l_{O_1}(x_1) = l_{O_2}(x_2) = a$, so $l_O(y) = a$ and the pullback maps preserve labels. It is easy to check that pullback mediating morphisms preserve labels, as they must commute with morphisms with such property. \square

Proof of Lemma 6.4. In ([7, Lemma 8]) we have proved that pushouts of order-embeddings in \mathbf{P} are commuting squares in \mathbf{P}_m . Therefore we can compute the two pushouts of (1) in \mathbf{P} , take the corresponding commuting squares in \mathbf{P}_m and then use Lemma A.4 to get labeling functions for their bottom-right corners. Diagrams in \mathbf{P}_m made of label preserving functions are also diagrams in \mathbb{O} .

Finally, the fact that σ^+ reflects orders follows from its definition. \square

Proof of Proposition 6.8. B is obtained by composition and product of accessible functors: \mathcal{P}_f is known to be accessible; \mathcal{L} is accessible, because it can be regarded as a constant endofunctor on $\mathbf{Set}^{\mathbb{O}}$; Δ is accessible, because it has a right adjoint, namely the functor computing right Kan extensions along δ .

In order to show that B covers pullbacks, we will show that it has the form $\mathcal{P}_f \circ B'$, with B' a pullback preserving endofunctor on $\mathbf{Set}^{\mathbb{O}}$. The thesis will follow from \mathcal{P}_f covering pullbacks (see [25]). Δ has a left adjoint, namely the functor computing left Kan extensions along δ , then it preserves pullbacks; \mathcal{L} can be seen as a constant, hence pullback-preserving, endofunctor on $\mathbf{Set}^{\mathbb{O}}$. B' is the product of these two functors, so it preserves pullbacks. \square

Proof of Proposition 6.12. Requirement Definition 6.11(ii) corresponds to the fact that a B -bisimulation R on (P, ρ) is a functor and its projections are natural transformations, so we have $(p, q)[\sigma]_R = (p[\sigma]_P, q[\sigma]_P)$, for any morphism σ in \mathbb{O} . Requirement (i) corresponds to the fact that RO is “almost” an ordinary bisimulation, because computing $\overline{BR}(O)$ essentially amounts to computing $\overline{B}_{flts}(RO)$ (see section 5.2) for each $O \in |\mathbb{O}|$, as images in $\mathbf{Set}^{\mathbb{O}}$ are computed pointwise in \mathbf{Set} , with the difference that continuations are not in RO , but in $R(\delta O)$. \square \square

Proof of Lemma 6.14. We have to prove that if the square on the left is a pullback then so is the outer square on the right.

$$\begin{array}{ccc}
 O & \xrightarrow{p_1} & O_1 \\
 \downarrow p_2 & \lrcorner & \downarrow \sigma_1 \\
 O_2 & \xrightarrow{\sigma_2} & O_3
 \end{array}
 \qquad
 \begin{array}{ccc}
 \mathcal{M}(O) & \xrightarrow{[p_1]} & \mathcal{M}(O_1) \\
 \downarrow [p_2] & \begin{array}{c} \text{---} \mu \text{---} \\ \text{---} P \text{---} \\ \text{---} \pi_1 \end{array} & \downarrow [\sigma_1] \\
 \mathcal{M}(O_2) & \xrightarrow{[\sigma_2]} & \mathcal{M}(O_3)
 \end{array}$$

In the right diagram, let P be the pullback in \mathbf{Set} of $[\sigma_1]$ and $[\sigma_2]$, namely

$$P = \{(c_1, c_2) \mid c_1[\sigma_1] = c_2[\sigma_2]\}$$

We will show that that the mediating morphism μ is an isomorphism, which implies that $\mathcal{M}(O)$ is a pullback object.

Take $(c_1, c_2) \in P$ and $c = c_1[\sigma_1] = c_2[\sigma_2]$. Then these causal markings must be of the form

$$c_1 = \{K_1 \vdash s_1, \dots, K_n \vdash s_n\} \quad c_2 = \{H_1 \vdash s_1, \dots, H_n \vdash s_n\} \quad c = \{L_1 \vdash s_1, \dots, L_n \vdash s_n\}$$

because $[\sigma_1]$ and $[\sigma_2]$ do not affect tokens. Moreover, we must have

$$L_i = \sigma_1(K_i) \downarrow_{O_1} = \sigma_2(H_i) \downarrow_{O_2} \quad (i = 1, \dots, n)$$

by definition of the action of \mathcal{M} on morphisms, and in particular

$$\max_{O_1} \sigma_1(K_i) = \max_{O_2} \sigma_2(H_i) = \max_{O_3} L_i$$

because K_i, H_i and L_i are down-closed sets, so they coincide with the closure of their maxima. It is easy to check that order-preserving and reflecting morphisms preserve maxima, so we have

$$\sigma_1(\max_{O_1} K_i) = \max_{O_3} \sigma_1(K_i) = \max_{O_2} \sigma_2(H_i) = \sigma_2(\max_{O_2} H_i).$$

Therefore, by definition of pullback in \mathbb{O} (computed as in **Graph**), there are $J_i \subseteq |O|$ such that

$$p_1(J_i) = \max_{O_1} K_i \quad p_2(J_i) = \max_{O_2} H_i \quad (\text{A.5})$$

and we can define the following causal marking in $\mathcal{M}(O)$

$$c' = \{\hat{J}_1 \vdash s_1, \dots, \hat{J}_n \vdash s_n\}$$

where $\hat{J}_i = J_i \downarrow_O$.

Now, observe that $c'[p_1] = c_1$ and $c'[p_2] = c_2$, because (A.5) implies $p_1(\hat{J}_i) \downarrow_{O_1} = K_i$ and $p_2(\hat{J}_i) \downarrow_{O_2} = H_i$. Therefore letting $\mu(c') = (c_1, c_2)$ makes the whole right diagram commute. So far we have proved that μ is surjective. For injectivity, suppose there is another $c'' \in \mathcal{M}(O)$ such that $\mu(c'') = (c_1, c_2)$. Since $c''[p_1] = c_1$ and $c''[p_2] = c_2$, c'' is again of the form $\{M_1 \vdash s_1, \dots, M_n \vdash s_n\}$, with $p_1(M_i) \downarrow_{O_1} = K_i$. Since also $K_i = p_1(\hat{J}_i) \downarrow_{O_1}$, M_i and \hat{J}_i must have the same set X of maxima. But then we have $M_i = X \downarrow_O = \hat{J}_i$, so $c'' = c'$. □

Proof of Theorem 6.16. The first item is just an instance of Proposition 6.12.

For the second item, we shall show that R is an AC-bisimulation closed under order-embeddings if and only if it is a \mathbb{O} -indexed bisimulation:

\implies : take $(O \triangleright c, O \triangleright \tilde{c}) \in R_O$ and suppose

$$O \triangleright c \xrightarrow{K \vdash a} \delta(O) \triangleright c'. \quad (\text{A.6})$$

Then, by Definition 6.15, there is

$$O \triangleright c \xrightarrow{K \vdash a} \delta(O, K, a) \triangleright c''$$

such that $c' = c''[\epsilon(O, K, a)]$. Since R is a AC-bisimulation, there is

$$O \triangleright \tilde{c} \xrightarrow{K \vdash a} \delta(O, K, a) \triangleright \tilde{c}'$$

such that $(\delta(O, K, a) \triangleright c'', \delta(O, K, a) \triangleright \tilde{c}') \in R_{\delta(O, K, a)}$. Again by Definition 6.15, from the last transition we get

$$O \triangleright \tilde{c} \xrightarrow{K \vdash a} \delta(O) \triangleright \tilde{c}'[\epsilon(O, K, a)].$$

This is a simulating transition for (A.6), because $(\delta(O, K, a) \triangleright c'', \delta(O, K, a) \triangleright \tilde{c}') \in R_{\delta(O, K, a)}$ implies $(\delta(O) \triangleright c', \delta(O) \triangleright \tilde{c}'[\epsilon(O, K, a)]) \in R_{\delta(O)}$, by closure of R under order-embeddings.

\Leftarrow : analogous to the previous point. Closure under order-embeddings of R follows from Definition 6.11(ii).

□

Proof of Proposition 7.2. Analogous to the proof of [7, Proposition 8].

□