

# Categorical specification and implementation of Replicated Data Types

Fabio Gadducci

*Dipartimento di Informatica, Università di Pisa*

Hernán Melgratti

*ICC – Universidad de Buenos Aires – CONICET, Argentina*

Christian Roldán

*IMDEA Software Institute*

Matteo Sammartino

*Royal Holloway University of London, University College London*

---

## Abstract

Replicated Data Types (RDTs) have been introduced as an abstraction for dealing with weakly consistent data stores, which may (temporarily) expose multiple, inconsistent views of their state. In the literature, RDTs are usually presented in set-theoretical terms, and only recently different specification flavours have been proposed. This paper offers a categorical presentation for the specification and implementation of RDTs. This paves the way for a method that allows distilling an operational semantics from a specification, which is then exploited to define a notion of implementation correctness via simulation.

*Keywords:* Replicated data types, Specification, Operational semantics, Functorial characterisation, Implementation correctness

---

## 1. Introduction

Modern distributed applications rely on data replication to achieve low latency and high availability even in the presence of failures at the expense of consistency, i.e., applications should tolerate and deal with temporarily inconsistent

---

\*Research partially supported by the MIUR PRIN 2017FTXR7S “IT-MaTTeRS”, by the EU H2020 RISE programme under the Marie Skłodowska-Curie grant agreement 778233, by the UBACyT projects 20020170100544BA and 20020170100086BA.

(partial) views of the state. One of the prominent abstractions for programming and reasoning about replicated states are *Replicated Data Types* (RDTs). Roughly, an RDT is an abstract data type whose behaviour embodies replication, i.e., its states describe partial views of the global state and the behaviour of operations is described in terms of the partial views of the state.

Different specification approaches for RDTs have been proposed in the literature [5, 6, 7, 8, 9, 11, 13, 15, 20, 22]. Despite stylistic differences, they abstractly represent the state of a system in terms of two relations defined on executed operations: *visibility*, which explains the partial view of the state over which each operation is executed, and *arbitration*, which totally orders operations and is used for resolving conflicting effects of concurrent operations. Leveraging this abstract view, in this paper we move from a set-theoretic presentation of RDTs to a categorical one, which offers a unifying perspective on the specification and implementation of RDTs. This allows distilling an operational semantics from a specification via a standard categorical construction and, ultimately, leads to a notion of implementation correctness defined as a standard simulation relation.

We start off with an illustrative example of RDTs, and then proceed to outline the contributions of this paper.

### 1.1. Illustrative example

Consider an RDT *Register* that represents a memory cell whose contents can be updated and read. Following the approach in [11], the RDT *Register* is specified by a function that maps visibility relations into sets of arbitrations: Here we call such function  $\mathcal{S}_{lwwR}$ . Figure 1a illustrates the definition of  $\mathcal{S}_{lwwR}$  for the case in which the visibility relation involves two concurrent writes and a read. Events are depicted by their labels  $\langle \text{operation}, \text{result} \rangle$  where  $\text{wr}(k)$  stands for an operation that writes the value  $k$  and  $\text{rd}$  stands for a read. The two writes are unrelated (i.e., they are not visible to each other), while the read operation sees both writes. The return value of the read operation is 2, which coincides with one of the visible written values.

According to fig. 1a,  $\mathcal{S}_{lwwR}$  maps the visibility graph into a set containing those arbitrations (i.e., total orders over the three events in the visibility relation) in which  $\text{wr}(1)$  precedes  $\text{wr}(2)$ . Arbitrations may not reflect the causal ordering of events; in fact, the last two arbitrations in the right-hand-side of the equation in Fig. 1a place the operation that read the value 2 before the operation that writes that same value.

We remark that arbitrations do not necessarily account for real-time orderings of events: They are instead possible ways in which events can be *logically* ordered to explain a given visibility. For instance, the excluded arbitrations in the image of  $\mathcal{S}_{lwwR}$  are the total orders in which  $\text{wr}(2)$  precedes  $\text{wr}(1)$ , i.e., the specification bans the behaviour in which a read operation returns a value that is different from the last written one. An extreme situation is the case in which the specification maps a visibility relation into an empty set of arbitrations, which means that events cannot be logically ordered to explain such visibility. For instance, the equation in Fig. 1b assigns an empty set of arbitrations to a visibility relation in which the read operation returns a value that is different

$$\mathcal{S}_{lwwR} \left( \begin{array}{cc} \langle \mathbf{wr}(1), ok \rangle & \langle \mathbf{wr}(2), ok \rangle \\ & \swarrow \quad \searrow \\ & \langle \mathbf{rd}, 2 \rangle \end{array} \right) = \left\{ \begin{array}{ccc} \langle \mathbf{wr}(1), ok \rangle & \langle \mathbf{wr}(1), ok \rangle & \langle \mathbf{rd}, 2 \rangle \\ | & | & | \\ \langle \mathbf{wr}(2), ok \rangle, & \langle \mathbf{rd}, 2 \rangle & , \langle \mathbf{wr}(1), ok \rangle \\ | & | & | \\ \langle \mathbf{rd}, 2 \rangle & \langle \mathbf{wr}(2), ok \rangle & \langle \mathbf{wr}(2), ok \rangle \end{array} \right\}$$

(a) Visibility relation with admissible arbitrations

$$\mathcal{S}_{lwwR} \left( \begin{array}{c} \langle \mathbf{wr}(1), ok \rangle \\ \downarrow \\ \langle \mathbf{rd}, 0 \rangle \end{array} \right) = \emptyset$$

(b) Non admissible arbitrations

Figure 1: A register specification

from the unique visible written value (i.e., it returns 0 instead of 1). In this way, the specification bans the behaviour in which a read operation returns a value that does not match a previously written value. As shown in [11], this style of specification can be considered equivalent to (and it is actually more general than) the model for the operational description of RDTs proposed in [7]. We refer the reader to [9] for a formal comparison of the two different approaches.

## 1.2. Outline of the paper

We now outline the structure of the paper and its main contributions.

*Section 2.* We start by recalling functional specifications of RDTs, and we provide several examples. We then restrict our attention to *coherent* specifications, which satisfy a suitable condition on arbitrations. We recall two set-theoretic operational models for coherent specifications: the *one-replica Labelled Transition System (LTS)*, modelling the behaviour of a single replica, and the *multi-replica LTS*, modelling the behaviour of multiple replica evolving concurrently.

*Section 3.* In this section we present our categorical model of coherent specifications. Given a set of labels  $\mathcal{L}$  for RDT operations, we model visibilities as  $\mathcal{L}$ -labelled direct acyclic graphs, and we define a suitable category  $\mathbf{PIDag}(\mathcal{L})$  for these graphs. Similarly, we form a category  $\mathbf{SPath}(\mathcal{L})$ , where objects are sets of  $\mathcal{L}$ -labelled paths (i.e., total orders) modelling arbitrations. As we shall see, particular care is required in defining morphisms for these categories. Our formulation ensures that certain limits and colimits exist, and that they capture important operations on arbitrations and visibilities. Once the basic categories are in place, we are able to generalise coherent specifications to a class of functors  $\mathbf{PIDag}(\mathcal{L}) \rightarrow \mathbf{SPath}(\mathcal{L})$ , dubbed *coherent functors*. The main result of this section is a sound (Theorem 3.32) and complete (Theorem 3.33) functorial characterisation of coherent specifications, which implies a one-to-one correspondence (Theorem 3.40) between coherent functors and specifications. Furthermore, we show that these results seamlessly extend to specific classes of specifications, namely *saturated* and *topological* ones.

*Section 4.* RDTs are usually implemented by several replicas that keep their own local state and propagate changes asynchronously. In this section we show that several well-known implementation of common RDTs can be defined in terms of monoidal LTSs, that is, LTSs whose states form a monoid and whose transition relation is compositional with respect to the state combination. The advantage of our formulation is a notion of bisimulation for RDT implementations that is closer to a standard high-order bisimulation (see, e.g., [16]).

*Section 5.* In this section we exploit the structure of coherent functors to provide a systematic way of recovering an operational semantics from a specification. This is achieved via a standard categorical construction, the so-called *category of elements* of the functor associated with a specification. In particular, we recover both the one-replica (Proposition 5.3) and the multi-replica LTSs (Proposition 5.6). Furthermore, we recover a *contextual* LTS, where labels provide the contextual information that explains how a local computation is embedded into a global context. Implementation correctness is thus straightforwardly stated in terms of a simulation between such LTSs and the one distilled from the specification itself. This characterisation allows for reframing previous ad-hoc formulations of implementation correctness by applying well-known notions from concurrency theory, thus paving the way for the application of more standard techniques in the analysis of RDTs.

*Section 6.* Finally, in this section we obtain a functorial characterisation of implementations. Implementation functors are defined as functors from a category  $\mathbf{IR}(\mathcal{R})$ , where arrows represent sequences of actions of a set  $\mathcal{R}$  of replica, to a category  $\mathbf{P}(\mathbf{Mon})$  representing non-deterministic computations of those replica. The set of states is modelled in the latter category as a monoid, which provides us with a one-to-one representation of the monoidal LTSs for implementations (Theorem 6.6). The category of elements construction again allows us to derive a suitable implementation LTS from implementation functors.

*Differences from conference paper.* This paper is an extended and revised version of the conference papers [12, 10]. Besides providing clarifications, proofs, and additional examples, in this presentation we give an alternative presentation for implementations, and consequently of their functorial characterisation. Technically, this is achieved by handling operations that mutate the state of an RDT from those that query it differently. Such *mutators* usually preserve the monoidal structure of states, hence the behaviour of a replica (as far as such mutators are concerned) can be described in terms of an LTS that preserves the monoidal structure of the states. This observation allows us to opt for a simpler functorial characterisation of computations in terms of power-monoids instead of the analogous power-domain construction over pre-ordered monoids introduced in [10]. The behaviour of *queries* is captured instead as barbs, i.e., as predicates accounting for the possible observations of states. Implementation equivalence translates to a barbed (higher-order) bisimulation [16], and thus also implementation correctness now relies on barbs.

## 2. Preliminaries on RDTs

In this section we briefly recall the functional model for RDTs specification and implementation introduced in [11, 9].

### 2.1. Notation

Let  $\mathcal{L}$  be a finite set of *labels*. A *labelled graph* is a triple  $\langle \mathcal{E}, \prec, \lambda \rangle$ , where  $\mathcal{E}$  is the set of vertices (they actually stand for “events”, hence the notation),  $\prec \subseteq \mathcal{E} \times \mathcal{E}$  is the (directed) connectivity relation, i.e.,  $e \prec e'$  means that there is an edge from  $e$  to  $e'$ , and  $\lambda: \mathcal{E} \rightarrow \mathcal{L}$  is a labelling function, assigning a label to each vertex. A graph is *acyclic* if the transitive closure of  $\prec$  is a strict partial order. We write  $\mathcal{E}_{\mathbb{G}}$ ,  $\prec_{\mathbb{G}}$  and  $\lambda_{\mathbb{G}}$  for the corresponding component of a specific graph  $\mathbb{G}$ . Given two labelled graphs  $\mathbb{G}$  and  $\mathbb{G}'$ , an isomorphism  $\varphi: \mathbb{G} \cong \mathbb{G}'$  between them is a bijective function  $\varphi: \mathcal{E}_{\mathbb{G}} \rightarrow \mathcal{E}_{\mathbb{G}'}$  that preserves graph structure and labels, i.e., such that  $e \prec_{\mathbb{G}} e'$  implies  $\varphi(e) \prec_{\mathbb{G}'} \varphi(e')$ , and  $\lambda_{\mathbb{G}'}(\varphi(e)) = \lambda_{\mathbb{G}}(e)$ , for all  $e, e' \in \mathcal{E}_{\mathbb{G}}$ . A *path*  $\langle \mathcal{E}, \leq, \lambda \rangle$  is a graph where  $\leq$  is a total order.

Given a graph  $\mathbb{G} = \langle \mathcal{E}, \prec, \lambda \rangle$  and a subset  $\mathcal{E}' \subseteq \mathcal{E}$ , we denote by  $\mathbb{G}|_{\mathcal{E}'}$  the obvious restriction (and the same applies to a path  $\mathbb{P}$ ), noting that  $\mathbb{G}|_{\emptyset} = \epsilon$ , with  $\epsilon$  the empty graph. Given  $\ell \in \mathcal{L}$  and a fresh event  $\top$ , we denote by  $\mathbb{G}_{\mathcal{E}'}^{\ell} = \langle \mathcal{E}_{\top}, \prec \cup (\mathcal{E}' \times \{\top\}), \lambda[\top \rightarrow \ell] \rangle$  the *extension* of  $\mathbb{G}$  over  $\mathcal{E}'$  with  $\ell$ , writing  $\mathbb{G}^{\ell}$  when  $\mathcal{E}' = \mathcal{E}$ .

We denote with  $\mathbb{G}(\mathcal{L})$  and  $\mathbb{P}(\mathcal{L})$  the collections of (finite) graphs and (finite) paths, respectively, labelled by  $\mathcal{L}$ . Also, when the set of labels  $\mathcal{L}$  is chosen, we let  $\mathbb{G}(\mathcal{E}, \lambda)$  and  $\mathbb{P}(\mathcal{E}, \lambda)$  be the collections of graphs and paths, respectively, whose set of vertices is  $\mathcal{E}$  and which are labelled by  $\lambda: \mathcal{E} \rightarrow \mathcal{L}$ .

### 2.2. Functional specification of replicated data types

We first recall the notion of (functional) specification. Hereafter we fix a set of labels  $\mathcal{L}$ .

**Definition 2.1** (Specifications). *A specification  $\mathcal{S}$  is an isomorphism-preserving function  $\mathcal{S}: \mathbb{G}(\mathcal{L}) \rightarrow 2^{\mathbb{P}(\mathcal{L})}$  such that  $\mathcal{S}(\epsilon) = \{\epsilon\}$  and  $\forall \mathbb{G}. \mathcal{S}(\mathbb{G}) \in 2^{\mathbb{P}(\mathcal{E}_{\mathbb{G}}, \lambda_{\mathbb{G}})}$ .*

A specification  $\mathcal{S}$  maps a graph (interpreted as the visibility relation of a RDT) to a set of paths (that is, the admissible arbitrations of the events). Indeed, each  $\mathbb{P} \in \mathcal{S}(\mathbb{G})$  is a path over  $\mathcal{E}_{\mathbb{G}}$ , hence a total order of the events in  $\mathbb{G}$ . In specifications we only care about the relative order of events and their labels, hence we require that isomorphic graphs are mapped to sets of paths that are element-wise isomorphic. More precisely, any isomorphism  $\varphi: \mathbb{G} \cong \mathbb{G}'$  of labelled graphs lifts to the corresponding sets of paths, meaning that  $\mathbb{P}' \in \mathcal{S}(\mathbb{G}')$  if and only if there exists  $\mathbb{P} \in \mathcal{S}(\mathbb{G})$  such that  $\varphi: \mathbb{P} \cong \mathbb{P}'$ , for all  $\mathbb{P}' \in \mathcal{S}(\mathbb{G}')$ .

We illustrate the approach by specifying a few well-known examples of RDTs.

### 2.2.1. Counters

A data type Counter maintains an integer value and provides operations for retrieving and modifying it. We consider the following two variants: increment-only counters (*GCounter*) and positive/negative counters (*PNCCounter*). The former account for non-negative counters, which provide the operation `inc` for increments, and the query operation `rd` for retrieving the current value of the counter. The latter may instead contain negative values and provide also the operation `dec` for decrementing. Hereafter, we assume that every counter is initialised in 0. Their specification is given below.

**Example 2.2** (*GCounter*). *The specification of the RDT GCounter relies on the following set of labels*

$$\mathcal{L}_{\mathcal{S}_{GC}} = \{\langle \text{inc}, ok \rangle\} \cup (\{\text{rd}\} \times \mathbb{N})$$

Essentially, a label stands for an executed operation (identified by the first component of the pair) and its observed result (given as the second component of the pair). By the definition of  $\mathcal{L}_{\mathcal{S}_{GC}}$  above, the execution of operation `inc` always returns `ok` (i.e., it succeeds), while `rd` returns a natural number. The specification of *GCounter* is given by the function  $\mathcal{S}_{GC}$  defined as follows: given a visibility graph  $\mathbf{G}$  and  $\mathbf{P} \in \mathbb{P}(\mathcal{E}_{\mathbf{G}}, \lambda_{\mathbf{G}})$

$$\mathbf{P} \in \mathcal{S}_{GC}(\mathbf{G}) \text{ iff } \forall e \in \mathcal{E}_{\mathbf{G}}. \lambda(e) = \langle \text{rd}, k \rangle \Rightarrow k = \#P$$

where  $P = \{e' \mid e' \prec_{\mathbf{G}} e \wedge \lambda(e') = \langle \text{inc}, ok \rangle\}$ . Hereafter, when defining specifications, we will keep the assumption  $\mathbf{P} \in \mathbb{P}(\mathcal{E}_{\mathbf{G}}, \lambda_{\mathbf{G}})$  implicit.

A visibility graph  $\mathbf{G}$  has admissible arbitrations (i.e.,  $\mathcal{S}_{GC}(\mathbf{G}) \neq \emptyset$ ) only if every event  $e$  in  $\mathbf{G}$  whose label corresponds to a read operation with return value  $k$  (i.e.,  $\langle \text{rd}, k \rangle$ ) observes (according to  $\mathbf{G}$ ) exactly  $k$  increment operations.

We illustrate two cases for the definition of  $\mathcal{S}_{GC}$  in Fig. 2. While the configuration in Fig. 2a has admissible arbitrations, the one in Fig. 2b has not, because the unique event labelled by `rd` returns 0 even if it is preceded by an observed increment. In other words, an execution is not allowed to generate such a visibility graph.

We remark that  $\mathcal{S}_{GC}$  does not impose any constraint on the ordering  $<_{\mathbf{P}}$ . In fact, a path  $\mathbf{P} \in \mathcal{S}_{GC}(\mathbf{G})$  does not need to reflect the ordering of  $\mathbf{G}$  as, for instance, the rightmost path in the set of Fig. 2a.

**Example 2.3** (*PNCCounter*). *A PNCCounter extends the definition of a GCounter by allowing decrement operations and negative values, which is reflected by the considered set of labels*

$$\mathcal{L}_{\mathcal{S}_{PNC}} = \{\langle \text{inc}, ok \rangle, \langle \text{dec}, ok \rangle\} \cup (\{\text{rd}\} \times \mathbb{Z})$$

The corresponding specification  $\mathcal{S}_{PNC}(\mathbf{G})$  is defined such that the following holds

$$\mathbf{P} \in \mathcal{S}_{PNC}(\mathbf{G}) \text{ iff } \left\{ \begin{array}{l} \forall e \in \mathcal{E}_{\mathbf{G}}. \lambda(e) = \langle \text{rd}, k \rangle \Rightarrow k = P - N \text{ where} \\ P = \#\{e' \mid e' \prec_{\mathbf{G}} e \wedge \lambda(e') = \langle \text{inc}, ok \rangle\} \\ N = \#\{e' \mid e' \prec_{\mathbf{G}} e \wedge \lambda(e') = \langle \text{dec}, ok \rangle\} \end{array} \right.$$

$$\begin{aligned}
\mathcal{S}_{GC} \left( \begin{array}{c} \langle \text{inc}, ok \rangle \\ \downarrow \\ \langle \text{rd}, 1 \rangle \end{array} \right) &= \left\{ \begin{array}{cc} \langle \text{inc}, ok \rangle & \langle \text{rd}, 1 \rangle \\ | & | \\ \langle \text{rd}, 1 \rangle & \langle \text{inc}, ok \rangle \end{array} \right\} \\
&\text{(a)} \\
\mathcal{S}_{GC} \left( \begin{array}{c} \langle \text{inc}, ok \rangle \\ \downarrow \\ \langle \text{rd}, 0 \rangle \end{array} \right) &= \emptyset \\
&\text{(b)}
\end{aligned}$$

Figure 2: Counter specification.

A graph  $\mathbf{G}$  has admissible arbitrations when each read returns a value  $k$  that corresponds to the difference between the number of increments and decrements observed by that read.

### 2.2.2. Registers

A data type Register maintains a value and provides two operations, one for updating it and another one for retrieving the current value. We provide specification for two alternative semantics, namely, a *Multi-value register* and a *Last-write wins* register.

**Example 2.4** (*Multi-value register*). A common abstraction of a memory cell in a replicated system is given by a Multi-value register. Differently from a traditional register, a multi-value one may contain several values when it is updated concurrently. Hence, we can fix the following set of labels

$$\mathcal{L}_{mvR} = \{\langle \mathbf{wr}(k), ok \rangle \mid k \in \mathcal{V}\} \cup (\{\mathbf{rd}\} \times 2^{\mathcal{V}})$$

where  $\langle \mathbf{wr}(k), ok \rangle$  stands for an operation that writes the integer  $k$  and  $\langle \mathbf{rd}, S \rangle$  for a read that retrieves the (possibly empty) set of values  $S$  stored in the register. The return value of every write operation is *ok* since they always succeed.

The specification is given by  $\mathcal{S}_{mvR} : \mathbb{G}(\mathcal{L}_{mvR}) \rightarrow 2^{\mathbb{P}(\mathcal{L}_{mvR})}$  defined as follows

$$\mathbf{P} \in \mathcal{S}_{mvR}(\mathbf{G}) \text{ iff } \left\{ \begin{array}{l} \forall e \in \mathcal{E}_{\mathbf{G}}. \\ \lambda(e) = \langle \mathbf{rd}, S \rangle \Rightarrow S = \{k \mid \exists e' \prec_{\mathbf{G}} e. \lambda(e') = \langle \mathbf{wr}(k), ok \rangle \wedge \\ (\forall e'' \prec_{\mathbf{G}} e, k'. e' \prec_{\mathbf{G}} e'' \Rightarrow \lambda(e'') \neq \langle \mathbf{wr}(k'), ok \rangle)\} \end{array} \right.$$

The condition on the right requires that any event  $e$  in  $\mathbf{G}$  associated with a read (i.e., labelled by  $\langle \mathbf{rd}, S \rangle$ ) returns a set  $S$  that contains all values written by maximal (according to  $\prec_{\mathbf{G}}$ ) concurrent updates seen by it. If this is the case, all arbitrations are admissible, i.e.,  $\mathbf{P} \in \mathcal{S}_{mvR}(\mathbf{G})$  for all  $\mathbf{P}$ , otherwise  $\mathcal{S}_{mvR}(\mathbf{G}) = \emptyset$ .

An instance of  $\mathcal{S}_{mvR}$  is shown in Fig. 3, where  $\mathbf{G}$  consists of three writes:  $\mathbf{wr}(2)$  overwrites  $\mathbf{wr}(1)$  and both are concurrent with  $\mathbf{wr}(3)$ . Additionally, there




Figure 3: Specification of a *Multi-value register*

are two reads: One observes all writes (right-most at the bottom), the other does not see  $\mathbf{wr}(3)$  (left-most one). Both reads return the set of values written by the maximal observed events: None of them returns 1 because it has been overwritten by 2. A graph is mapped by  $\mathcal{S}_{mvR}$  to  $\emptyset$  if it describes an inconsistent configuration, e.g., if the return value of one read in Fig. 3 were changed to  $\{1\}$ .

According to  $\mathcal{S}_{mvR}$ , events can be arbitrated in any order, allowing read events to happen before observed writes (as in the second and third path in Fig. 3). This is a common approach in the specification of RDTs [7] because it allows permissive strategies for implementation (we refer to [9] for details). If needed, a specification can exclude arbitrations (as illustrated in Ex. 2.5).

**Example 2.5** (*Last-write wins Register*). An alternative to the Multi-value Register is the Last-write wins Register, in which every read returns the last written value according to arbitration. We take the following set of labels

$$\mathcal{L}_{lwwR} = \{\langle \mathbf{wr}(k), ok \rangle, \langle \mathbf{rd}, k \rangle \mid k \in \mathcal{V}\} \cup \{\langle \mathbf{rd}, \perp \rangle\}$$

where  $\perp$  is the initial value of a register. Its specification  $\mathcal{S}_{lwwR}$  is given by

$$\mathbf{P} \in \mathcal{S}_{lwwR}(\mathbf{G}) \text{ iff } \begin{cases} \forall \mathbf{e} \in \mathcal{E}_{\mathbf{G}}. \\ \lambda(\mathbf{e}) = \langle \mathbf{rd}, \perp \rangle \Rightarrow \forall \mathbf{e}' \prec_{\mathbf{G}} \mathbf{e}, k'. \lambda(\mathbf{e}') \neq \langle \mathbf{wr}(k'), ok \rangle \wedge \\ \lambda(\mathbf{e}) = \langle \mathbf{rd}, k \rangle \Rightarrow \exists \mathbf{e}' \prec_{\mathbf{G}} \mathbf{e}. \lambda(\mathbf{e}') = \langle \mathbf{wr}(k), ok \rangle \wedge \\ (\forall \mathbf{e}'' \prec_{\mathbf{G}} \mathbf{e}, k'. \mathbf{e}' \prec_{\mathbf{P}} \mathbf{e}'' \Rightarrow \lambda(\mathbf{e}'') \neq \langle \mathbf{wr}(k'), ok \rangle) \end{cases}$$

According to  $\mathcal{S}_{lwwR}$ , a read returns  $\perp$  when it does not observe any write. On the contrary, a read  $\mathbf{e}$  returns a natural number  $k$  when it observes some event  $\mathbf{e}'$  that writes  $k$ . In such case, the arbitration  $\mathbf{P}$  must order  $\mathbf{e}'$  as the maximal event (accordingly to  $\prec_{\mathbf{P}}$ ) among all write operations seen by  $\mathbf{e}$ . In this way, the specification constrains the allowed arbitrations of a graph.

### 2.3. Sets

A data type for sets provides the operations for adding, removing and examining the elements within the set. Different alternatives have been proposed in the literature for resolving conflicts in the presence of concurrent additions and removals; we illustrate two possible alternatives below.



**Example 2.6** (*Observed-remove set*). An Observed-remove set (Or-Set) is an abstraction for sets in which conflicting concurrent additions and removals are resolved by making the former to win over the latter. Since it is an abstraction for sets, it provides the operations to add and remove elements, and also to lookup its content. Hence, we take the following set of labels

$$\mathcal{L}_{Or-Set} = \{\langle \text{add}(k), \text{ok} \rangle, \langle \text{rem}(k), \text{ok} \rangle \mid k \in \mathcal{V}\} \cup (\{\text{lookup}\} \times 2^{\mathcal{V}})$$

The specification of the Or-Set is as follows

$$P \in \mathcal{S}_{Or-Set}(\mathbb{G}) \text{ iff } \begin{cases} \forall e \in \mathcal{E}_{\mathbb{G}}. \\ \lambda(e) = \langle \text{lookup}, S \rangle \Rightarrow \\ S = \{k \mid \exists e' \prec_{\mathbb{G}} e. \lambda(e') = \langle \text{add}(k), \text{ok} \rangle \wedge \\ (\forall e'' \prec_{\mathbb{G}} e. e' \prec_{\mathbb{G}} e'' \Rightarrow \lambda(e'') \neq \langle \text{rem}(k), \text{ok} \rangle)\} \end{cases}$$

The value  $S$  returned by a read event  $e$  contains all those elements that have been added but not subsequently removed according to the visibility of  $e$ .

**Example 2.7** (2P-Set). A Two-Phase set (2P-Set) is an abstraction in which elements can be added and removed; however, after being removed, an element cannot be re-added to the set. For its specification, we fix a set of labels identical to the OR-Set, namely

$$\mathcal{L}_{2P-Set} = \{\langle \text{add}(k), \text{ok} \rangle, \langle \text{rem}(k), \text{ok} \rangle \mid k \in \mathcal{V}\} \cup (\{\text{lookup}\} \times 2^{\mathcal{V}})$$

Its specification of then defined as follows

$$P \in \mathcal{S}_{2P-Set}(\mathbb{G}) \text{ iff } \begin{cases} \forall e \in \mathcal{E}_{\mathbb{G}}. \\ \lambda(e) = \langle \text{lookup}, S \rangle \Rightarrow \\ S = \{k \mid \exists e' \prec_{\mathbb{G}} e. \lambda(e') = \langle \text{add}(k), \text{ok} \rangle \wedge \\ \forall e'' \prec_{\mathbb{G}} e. \lambda(e'') \neq \langle \text{rem}(k), \text{ok} \rangle\} \end{cases}$$

The value  $S$  returned by a read event  $e$  contains all those elements that have been added but not removed according to the visibility of  $e$  (regardless of whether the addition took place before or after the removal).

#### 2.4. Coherent Specifications

We now restrict our attention to *coherent* specifications, which suffice for the standard specification of RDTS [9] and are amenable to a categorical characterisation, as illustrated in the next section. Coherence expresses that admissible arbitrations of a visibility graph are obtained by composing the admissible arbitrations corresponding to smaller visibilities. Its formal definition relies on an auxiliary operation for composing sets of paths. We say that the paths of a set  $\mathcal{X} = \{P_i\}_{i \in I}$  are *compatible* if we have  $\lambda_j(e) = \lambda_k(e)$  for all  $e \in \mathcal{E}_j \cap \mathcal{E}_k$ .

**Definition 2.8** (Product). The product of a finite set  $\mathcal{X}$  of compatible paths is

$$\bigotimes \mathcal{X} = \{P \mid P \text{ is a path over } \bigcup_i \mathcal{E}_i \text{ and } P|_{\mathcal{E}_i} \in \mathcal{X}\}$$

The product of paths is analogous to the synchronous product of transition systems: Common elements are identified and the remaining ones can be freely interleaved, as long as the original orders are respected. Clearly, the result is  $\emptyset$  if so is  $\mathcal{X}$ , while if the set contains only the empty path  $\epsilon$  the result is  $\{\epsilon\}$ .

A set of sets of paths  $\mathcal{X}_1, \mathcal{X}_2, \dots$  is compatible if  $\bigcup_i \mathcal{X}_i$  is so, and we can define  $\bigotimes_i \mathcal{X}_i$  as  $\bigotimes \bigcup_i \mathcal{X}_i$ . But first, for an element  $e$  of a graph  $\mathbf{G}$ , we denote by  $[e]$  the set of elements of  $\mathbf{G}$  that precede  $e$ , namely,  $\{e' \mid e' \prec_{\mathbf{G}}^* e\}$ .

**Definition 2.9** ((Past-)Coherent Specification). *Let  $\mathcal{S}$  be a specification. We say that  $\mathcal{S}$  is past-coherent (briefly, coherent) if  $\forall \mathbf{G} \neq \epsilon. \mathcal{S}(\mathbf{G}) = \bigotimes_{e \in \mathcal{E}_{\mathbf{G}}} \mathcal{S}(\mathbf{G}|_{[e]})$ .*

In a coherent specification  $\mathcal{S}$  the arbitrations of a configuration  $\mathbf{G}$  (i.e., the set of paths  $\mathcal{S}(\mathbf{G})$ ) are the composition of the arbitrations of its sub-graphs  $\mathbf{G}|_{[e]}$ . It can be shown that the specifications in Section 2.2 are coherent.


Coherent specifications can be given an operational interpretation representing admissible computations. We start by defining a *one-replica Labelled Transition System (LTS)*, modelling the behaviour of a single replica.

**Definition 2.10** (One-replica LTS). *Given a specification  $\mathcal{S}$ , its one-replica LTS is defined as follows*


- states are pairs  $\langle \mathbf{G}, \mathbf{P} \rangle$  such that  $\mathbf{P} \in \mathcal{S}(\mathbf{G})$ ;
- transitions are of the form  $\langle \mathbf{G}, \mathbf{P} \rangle \xrightarrow{\ell} \langle \mathbf{G}^\ell, \mathbf{P}' \rangle$ , where  $\mathbf{P}'|_{\mathcal{E}_{\mathbf{G}}} = \mathbf{P}$ .

The behaviour of a single replica that implements a specification  $\mathcal{S}$  can be defined in terms of the evolution of the relations  $\mathbf{G}$  and  $\mathbf{P}$ , which respectively account for the visibility and arbitration of the events known by that particular replica. Only pairs  $\langle \mathbf{G}, \mathbf{P} \rangle$  in which  $\mathbf{P}$  is an admissible arbitration of  $\mathbf{G}$ , i.e., such that  $\mathbf{P} \in \mathcal{S}(\mathbf{G})$ , are considered. Each transition  $\langle \mathbf{G}, \mathbf{P} \rangle \xrightarrow{\ell} \langle \mathbf{G}^\ell, \mathbf{P}' \rangle$  represents the execution of an operation  $\ell$  over the state  $\langle \mathbf{G}, \mathbf{P} \rangle$ . Such execution extends the visibility relation  $\mathbf{G}$  with a top event labelled by  $\ell$  (i.e.,  $\mathbf{G}^\ell$ ), which accounts for the fact that the operation observes all the events known by the replica. Since  $\langle \mathbf{G}^\ell, \mathbf{P}' \rangle$  should be a state of the LTS,  $\mathbf{P}'$  is a total order of the events in  $\mathbf{G}^\ell$ . The condition  $\mathbf{P}'|_{\mathcal{E}_{\mathbf{G}}} = \mathbf{P}$  ensures that  $\mathbf{P}'$  preserves the arbitration  $\mathbf{P}$ , i.e.,  $\mathbf{P}'$  is obtained by inserting the fresh event labelled by  $\ell$  in any position within  $\mathbf{P}$ .

**Example 2.11.** *Consider the RDT GCounter specified in 2.2.1. Suppose we have a replica that is aware of the execution of three increment operations, two of which have been performed locally, and another one has been (concurrently) executed over a different replica. An abstract representation of the state of that replica is given by the visibility relation  $\mathbf{G}$  depicted below and a total order  $\mathbf{P}$  over those events (any total order would work in this case because arbitration is uninfluential for GCounter).*



Note that  $\mathbf{G}$  abstracts away from the identity of the replicas over which the operations have been executed and only keeps track of their (visibility) dependencies: Each pair of unrelated events corresponds to operations concurrently executed over different replicas. A state of a single replica is associated with a visibility graph accounting for operations executed over different replicas; in this way, we abstractly represent the propagation of events that occurs in concrete implementations. In this scenario, the pair  $\langle \mathbf{G}, \mathbf{P} \rangle$  can be regarded as the state of a replica that has locally performed two increments and then has received an increment propagated from other replica. Interestingly,  $\langle \mathbf{G}, \mathbf{P} \rangle$  also describes several other possibilities; e.g., a replica that has performed one increment and then has received information about two (sequential) increments from a different replica; or a third replica that has received the updates propagated by other two replicas. In any case, the behaviour of the replica is just determined by  $\langle \mathbf{G}, \mathbf{P} \rangle$ . The transitions from  $\langle \mathbf{G}, \mathbf{P} \rangle$  have the following shape  $\langle \mathbf{G}, \mathbf{P} \rangle \xrightarrow{\ell} \langle \mathbf{G}^\ell, \mathbf{P}' \rangle$  where  $\ell \in \mathcal{L}_{\mathcal{S}_{GC}} = \{\langle \text{inc}, ok \rangle\} \cup (\{\text{rd}\} \times \mathbb{N})$  and  $\mathbf{G}^\ell$  is defined as follows



If  $\ell = \langle \text{inc}, ok \rangle$ , then it suffices to take  $\mathbf{P}'$  as one of the possible extensions of  $\mathbf{P}$  (where the new event can be added in any position in  $\mathbf{P}$ ). If  $\ell = \langle \text{rd}, k \rangle$ , then we have two cases. If  $k = 3$ , then the situation is analogous to previous case. On the contrary, if  $k \neq 3$ , we first note that  $\mathcal{S}_{GC}(\mathbf{G}^{\langle \text{rd}, k \rangle}) = \emptyset$  because  $\mathcal{S}_{GC}$  imposes the return value of every read operation to coincide with the number of observed reads. Hence, no transition from  $\langle \mathbf{G}, \mathbf{P} \rangle$  can be labelled by  $\langle \text{rd}, k \rangle$  if  $k \neq 3$  (i.e., no read operation can return a value different from 3 when the replica knows three increments).

A *multi-replica* LTS models multiple replica evolving concurrently. Suppose we have two replica, and the current state for each is  $\langle \mathbf{G}_i, \mathbf{P}_i \rangle$ . We assume that  $\mathbf{G}_1$  and  $\mathbf{G}_2$  are *compatible*, i.e., shared events have the same labels and predecessors in the two graphs (that is,  $\lambda_{\mathbf{G}_1}(e) = \lambda_{\mathbf{G}_2}(e)$  and  $\{e' \mid e' \prec_{\mathbf{G}_1}^* e\} = \{e' \mid e' \prec_{\mathbf{G}_2}^* e\}$  for all  $e \in \mathcal{E}_{\mathbf{G}_1} \cap \mathcal{E}_{\mathbf{G}_2}$ ). We denote by  $\mathbf{G}_1 \sqcup \mathbf{G}_2$  the union of compatible graphs. A


replicated state is of the form  $\langle \mathbf{G}_1 \sqcup \mathbf{G}_2, \mathbf{P} \rangle$ , where  $\mathbf{P} \in \mathbf{P}_1 \otimes \mathbf{P}_2$ , i.e.,  $\mathbf{P}$  is obtained by “synchronising” the individual arbitrations.

**Definition 2.12** (Multi-replica LTS). *Given a specification  $\mathcal{S}$ , its multi-replica LTS is derived from the one-replica LTS by adding the following inference rule*

$$\text{(COMP)} \frac{\langle \mathbf{G}_1, \mathbf{P}|_{\mathcal{E}_{\mathbf{G}_1}} \rangle \xrightarrow{\ell} \langle \mathbf{G}'_1, \mathbf{P}'_1 \rangle \quad \mathbf{P}' \in \mathbf{P} \otimes \mathbf{P}'_1 \quad (\mathcal{E}_{\mathbf{G}'_1} \setminus \mathcal{E}_{\mathbf{G}_1}) \cap \mathcal{E}_{\mathbf{G}_2} = \emptyset}{\langle \mathbf{G}_1 \sqcup \mathbf{G}_2, \mathbf{P} \rangle \xrightarrow{\ell} \langle \mathbf{G}'_1 \sqcup \mathbf{G}_2, \mathbf{P}' \rangle}$$

The rule says that transitions from a replicated state  $\langle \mathbf{G}_1 \sqcup \mathbf{G}_2, \mathbf{P} \rangle$  are obtained from those of a single component  $\langle \mathbf{G}_1, \mathbf{P}|_{\mathcal{E}_{\mathbf{G}_1}} \rangle$  (note that the projection recovers the local arbitration). After the transition, the replicated state evolves to the composition of the part that has not changed (namely  $\mathbf{G}_2$ ) and the new local state  $\langle \mathbf{G}'_1, \mathbf{P}'_1 \rangle$ . In short, global computations arise from computations of single replicas. The last condition in the premise of (COMP) ensures that events generated during computation are globally unique, i.e., new events in  $\mathbf{G}'_1$  should not be already present in  $\mathbf{G}_2$ .

**Example 2.13.** *Consider again the RDT GCounter and a scenario similar to the one in Example 2.11, in which a replica  $r_1$  has performed two sequential increments and  $r_2$  one, but none of them has propagated its local information yet. Then, the state of the system can be represented as the composition of the states of the two replicas. Since  $\mathbf{P}$  is uninteresting for GCounter, we will just focus on the visibility graph  $\mathbf{G}$ , which is depicted as follows*




Note that  $\mathbf{G}$  combines the visibility relations corresponding to each of the replicas (which are depicted by dotted boxes).


As discussed in Example 2.11, the one-replica LTS describes the behaviour of each operation by considering all events in the state, i.e., each fresh event dominates all other events in the visibility relation. In our scenario, this means that a read operation over  $\mathbf{G}$  can only return 3 (according to the one-replica LTS). However, a read operation executed over  $r_1$ , which is aware of two increments, should return 2; and analogously, a read over  $r_2$  should return 1.

The multi-replica LTS accounts for the execution of operations that have a partial view of the state of the system. This is achieved via the rule (COMP), which accounts for decompositions of the global state. For instance, (COMP) allows us to derive a transition describing the evolution of the system when a read operation is performed over the replica  $r_1$ . In fact, by the one-replica LTS,

we can derive a transition from  $\langle G_1, P_1 \rangle$  to  $\langle G'_1, P'_1 \rangle$ , which extends  $G_1$  with a new event labelled by  $\langle rd, 2 \rangle$ . Then, (COMP) establishes that  $G = G_1 \sqcup G_2$  can evolve to  $G'$  below



Similarly, a read operation over  $r_2$  would extend the visibility with a new event labelled by  $\langle rd, 1 \rangle$  that dominates the events in  $G_2$  (i.e., a read returns 1 if performed over  $r_2$ ).



As highlighted in Example 2.11, we do not explicitly represent replicas. Consequently, rule (COMP) accounts for the behaviour of operations performed over any partial view of  $G$ .

For instance, consider the decomposition  $G = \emptyset \sqcup G$ . The rule (COMP) can be applied to such decomposition to account for the behaviour of an operation performed over the empty state  $\emptyset$ . Intuitively, this is the state of a replica that has not performed any operation and has not received any information from the other replicas, so a read operation will return 0. Analogously, if we consider  $G = G \sqcup G$ , (COMP) will account for operations executed over a replica that has complete view of the replicated state. Considering all possible decompositions will yield transitions in the multi-replica LTS accounting for the fact that a read operation over  $G$  may return any of the values in  $\{0, 1, 2, 3\}$ .

### 3. Categorical model of specifications

In this section we first introduce the key ingredients of our categorical approach, and then we present our categorical model of coherent specifications. We assume familiarity with basic concepts of category theory, such as functors, natural transformations, limits and colimits.

### 3.1. Categories of relations

We first introduce a category of binary relations that will serve as the basic category for modelling both visibilities and arbitrations.

**Definition 3.1** (Binary Relations). *Given a finite set  $\mathbf{E}$ , a (binary) relation  $\rho$  over  $\mathbf{E}$ , written  $\langle \mathbf{E}, \rho \rangle$ , is a subset  $\rho \subseteq \mathbf{E} \times \mathbf{E}$ . A subset  $\mathbf{E}' \subseteq \mathbf{E}$  is downward closed with respect to  $\rho$  if  $(e, e') \in \rho$  implies  $e \in \mathbf{E}'$ , for all  $e' \in \mathbf{E}'$ .*

We write  $\emptyset$  for the empty relation and  $e \rho e'$  to mean  $(e, e') \in \rho$ . And we write  $[e]_\rho$  for the smallest downward closed set with respect to  $\rho$  including  $e \in \mathbf{E}$ , omitting the subscript  $\rho$  if clear from the context: It coincides with  $\{e' \mid e' \rho^* e\}$ .

We now look at morphisms between relations. Besides ordinary relation-preserving morphisms, we need a special type of morphism, called *past-reflecting (pr-)morphism*, which intuitively adds no dependencies in the past of an event. As we shall see, pr-morphisms play a key role in defining the category of visibilities.

**Definition 3.2** ((Binary Relation) Morphisms). *A (binary relation) morphism  $\mathbf{f} : \langle \mathbf{E}, \rho \rangle \rightarrow \langle \mathbf{T}, \gamma \rangle$  is a function  $\mathbf{f} : \mathbf{E} \rightarrow \mathbf{T}$  such that  $e \rho e'$  implies  $\mathbf{f}(e) \gamma \mathbf{f}(e')$  for all  $e, e' \in \mathbf{E}$ . A morphism  $\mathbf{f} : \langle \mathbf{E}, \rho \rangle \rightarrow \langle \mathbf{T}, \gamma \rangle$  is past-reflecting (shortly, pr-morphism) if  $\mathbf{t} \gamma \mathbf{f}(e)$  implies that there is  $e' \in \mathbf{E}$  such that  $e' \rho e$  and  $\mathbf{t} = \mathbf{f}(e')$  for all  $e \in \mathbf{E}$  and  $\mathbf{t} \in \mathbf{T}$ .*

Past-reflecting morphisms are known under various names in different contexts, e.g. as bounded morphisms in modal logic (see, e.g., [3, Pag. 60]). We now give a useful characterisation result.

**Lemma 3.3** (Characterising pr-morphisms). *Let  $\mathbf{f} : \langle \mathbf{E}, \rho \rangle \rightarrow \langle \mathbf{T}, \gamma \rangle$  be a morphism. If*

1.  $\mathbf{f}(e) \gamma \mathbf{f}(e')$  implies  $e \rho e'$ , and
2.  $\bigcup_{e \in \mathbf{E}} \mathbf{f}(e)$  is downward closed,

*then it is a pr-morphism. If  $\mathbf{f}$  is injective, then the converse holds.*

*Proof.* For  $\Rightarrow$ ), let us take  $e \in \mathbf{E}$  and  $\mathbf{t} \in \mathbf{T}$ . If  $\mathbf{t} \gamma \mathbf{f}(e)$ , then there exists  $e' \in \mathbf{E}$  such that  $\mathbf{t} = \mathbf{f}(e')$  because of (2). By (1),  $\mathbf{f}(e') \gamma \mathbf{f}(e)$  implies  $e' \rho e$ .

For  $\Leftarrow$ ), by the definition of pr-morphism  $\mathbf{f}(e) \gamma \mathbf{f}(e')$  implies  $\exists \bar{e} \in \mathbf{E}. \bar{e} \rho e' \wedge \mathbf{f}(e) = \mathbf{f}(\bar{e})$ . Since  $\mathbf{f}$  is injective,  $\bar{e} = e$  and hence  $e \rho e'$ . So, let  $\mathcal{T} = \bigcup_{e \in \mathbf{E}} \mathbf{f}(e)$ . We want to show that

$$\forall \mathbf{t} \in \mathbf{T}, \mathbf{t}' \in \mathcal{T}. \mathbf{t} \gamma \mathbf{t}' \text{ implies } \mathbf{t} \in \mathcal{T}$$

The proof follows by contradiction. Assume that  $\exists \mathbf{t} \in \mathbf{T}, \mathbf{t}' \in \mathcal{T}. \mathbf{t} \gamma \mathbf{t}' \wedge \mathbf{t} \notin \mathcal{T}$ . By definition of  $\mathcal{T}$ ,  $\exists e \in \mathbf{E}$  such that  $\mathbf{f}(e) = \mathbf{t}'$ . Since  $\mathbf{f}$  is a pr-morphism, then

$$\mathbf{t} \gamma \mathbf{f}(e) \text{ implies } \exists e' \in \mathbf{E}. e' \rho e \wedge \mathbf{t} = \mathbf{f}(e')$$

Therefore  $\mathbf{t} = \mathbf{f}(e') \in \mathcal{T}$ , which contradicts the assumption  $\mathbf{t} \notin \mathcal{T}$ .  $\square$

Both classes of morphisms are closed under composition: **Bin** denotes the category of relations and their morphisms and **PBin** the sub-category of pr-morphisms. The category **Bin** has finite limits and finite colimits, which are computed point-wise as in **Set**. The structure is lifted by and large to **PBin**: Finite colimits and binary pullbacks in **PBin** are computed as in **Bin**. Yet there is no terminal object, as clearly not all relations admit a past-reflecting morphism into the singleton.

Monos in **Bin** are just morphisms whose underlying function is injective, and similarly in **PBin**, so that the inclusion functor preserves (and reflects) them. Thus, we can lift the following lemma from the category of sets and functions.

**Lemma 3.4** (Monos under pushouts). *Pushouts in **Bin** (and thus in **PBin**) preserve monos.*

Given a set of labels  $\mathcal{L}$ , the category of labelled relations is **Bin**( $\mathcal{L}$ ).

**Definition 3.5** (Category of labelled relations). *The category **Bin**( $\mathcal{L}$ ) is defined as the comma category  $\mathbf{U}_r \downarrow \mathcal{L}$ , where  $\mathbf{U}_r : \mathbf{Bin} \rightarrow \mathbf{Set}$  is the forgetful functor, mapping a binary relation to its underlying set of elements. Explicitly, an object in **Bin**( $\mathcal{L}$ ) is a triple  $(\mathbf{E}, \rho, \lambda)$  for a labeling function  $\lambda : \mathbf{E} \rightarrow \mathcal{L}$ . A label-preserving morphism  $(\mathbf{E}, \rho, \lambda) \rightarrow (\mathbf{E}', \rho', \lambda')$  is a morphism  $\mathbf{f} : (\mathbf{E}, \rho) \rightarrow (\mathbf{E}', \rho')$  such that  $\lambda(\mathbf{s}) = \lambda'(\mathbf{f}(\mathbf{s}))$  for all  $\mathbf{s} \in \mathbf{E}$ .*

The category **PBin**( $\mathcal{L}$ ) is defined analogously, with the requirement that the morphisms are past-reflecting. In both categories, finite colimits and binary pullbacks always exist and are essentially computed as in **Bin**.

### 3.2. Categories of Graphs and Paths

We now move to introduce specific sub-categories that are going to be used for both the syntax and the semantics of specifications.

**Definition 3.6** (**PDag**). ***PDag** is the full sub-category of **PBin** whose objects are acyclic graphs (relations whose transitive closures is a strict partial order).*

The full sub-category of **Bin** whose objects are acyclic graphs is not suited for our purposes, since it does not admit pushouts, not even along monos. Pushouts are important for our work, as they will allow us to generalise composition of visibilities. The sub-category of **Bin** with pr-morphisms is much more well-behaved, still remaining computationally simple. In fact, it enjoys the following property. Recall that a functor *reflects* (co)limits if every diagram mapped to a (co)limit by the functor is a (co)limit itself.

**Proposition 3.7** (Properties of **PDag**). *The inclusion functor **PDag**  $\rightarrow$  **PBin** reflects finite colimits and binary pullbacks.*

In other terms, finite colimits and binary pullbacks in **PDag** are computed as in **PBin**, hence as in **Bin**.

We now move to consider the category of paths.




Figure 4: Some labelled graphs

**Definition 3.8 (Path).** *Path* is the full sub-category of **Bin** whose objects are paths (i.e., total orders).

Note that defining **Path** as only containing pr-morphisms would be too restrictive, since there exists a pr-morphism between two paths if and only if one path is a prefix of the other. Nevertheless, as for **PDag**, finite colimits in **Path** are computed as in **Bin**.

**Proposition 3.9** (Properties of **Path**). *The inclusion functor  $\mathbf{Path} \rightarrow \mathbf{Bin}$  reflects finite colimits.*

As for relations, we consider suitable comma categories in order to capture labelled paths and graphs. In particular, we use the forgetful functors  $\mathbf{U}_{\text{rp}} : \mathbf{Path} \rightarrow \mathbf{Set}$  and  $\mathbf{U}_{\text{pd}} : \mathbf{PDag} \rightarrow \mathbf{Set}$ : For a set of labels  $\mathcal{L}$  we denote  $\mathbf{PDag}(\mathcal{L}) = \mathbf{U}_{\text{rp}} \downarrow \mathcal{L}$  and  $\mathbf{Path}(\mathcal{L}) = \mathbf{U}_{\text{pd}} \downarrow \mathcal{L}$ . Once more, finite colimits and binary pullbacks always exist and are essentially computed as in **Bin**.

**Example 3.10.** *We illustrate some labelled graphs in Fig. 4 and remark that  $P_1$  is the only path in the figure. Note that  $G_1$  is not a path because the relation is not transitive. There is an obvious label-preserving morphism  $f_1 : G_1 \rightarrow P_1$ , but this is not a pr-morphism because the edge from  $\langle \mathbf{wr}(1), ok \rangle$  to  $\langle \mathbf{rd}, \{2\} \rangle$  in  $P_1$  is not matched in  $G_1$ . On the contrary, there is no morphism from  $P_1$  to  $G_1$ . Note that the label-preserving morphisms  $f_2 : P_1 \rightarrow G_2$  and  $f_3 : G_2 \rightarrow G_3$  are pr-morphisms (and consequently,  $f_2; f_3 : P_1 \rightarrow G_3$  is so).*

### 3.3. Category of visibilities

We now study the category of visibility relations. We first introduce an operation that will be handy for our categorical characterisation. We say that a graph  $G$  is *rooted* if there exists a (necessarily unique) event  $e \in \mathcal{E}_G$  such that  $G = G|_{[e]}$ .

**Definition 3.11** (Extension). *Let  $f : G_1 \rightarrow G_2$  be a mono pr-morphism. It is an extension along  $\ell$  (shortly,  $\ell$ -extension) if, for a fresh event  $\top$ ,  $\mathcal{E}_{G_2} = \mathcal{E}_{G_1} \cup \{\top\}$ ,  $f : \mathcal{E}_{G_1} \rightarrow \mathcal{E}_{G_1} \cup \{\top\}$  is the associated injection, and  $\lambda_{G_2}(\top) = \ell$ .*

Intuitively,  $G_2$  is obtained by adding to the visibility relation  $G_1$  a new event “seeing” some events in  $G_1$ . We say that an extension  $f : G_1 \rightarrow G_2$  is a *root extension* if  $G_2$  is rooted.



**Example 3.12.** Noting that  $\mathbf{f}_2$  and  $\mathbf{f}_3$  in Ex. 3.10 are mono, the former is an extension along  $\langle \mathbf{wr}(3), \text{ok} \rangle$  and the latter is an extension along  $\langle \mathbf{rd}, \{2, 3\} \rangle$ . Since  $\mathbf{G}_3$  is rooted,  $\mathbf{f}_3$  is a rooted extension; however,  $\mathbf{f}_2$  is not rooted because  $\mathbf{G}_2$  is not. Note that  $\mathbf{f}_2; \mathbf{f}_3$  is a pr-morphism but not an extension because the target  $\mathbf{G}_3$  adds two new events to the source  $\mathbf{P}_1$ .

**Proposition 3.13.** Rooted graphs form a family of separators of  $\mathbf{PDag}(\mathcal{L})$ .

*Proof.* We need to show that for any pair of pr-morphisms  $\mathbf{f}_1, \mathbf{f}_2 : \mathbf{G}_1 \rightarrow \mathbf{G}_2$  such that  $\mathbf{f}_1 \neq \mathbf{f}_2$  there is a rooted graph  $\mathbf{G}$  and a morphism  $\mathbf{f} : \mathbf{G} \rightarrow \mathbf{G}_1$  such that  $\mathbf{f}; \mathbf{f}_1 \neq \mathbf{f}; \mathbf{f}_2$ . Given  $e \in \mathcal{E}_{\mathbf{G}_1}$  such that  $\mathbf{f}_1(e) \neq \mathbf{f}_2(e)$ , it suffices to consider the pr-morphism  $\mathbf{f} : \mathbf{G}_1|_{[e]} \rightarrow \mathbf{G}_1$ .  $\square$

We now further curb the arrows in  $\mathbf{PDag}(\mathcal{L})$  to *monic* ones. Intuitively, we are only interested in what happens if we add further events to visibility relations. We thus consider the sub-category  $\mathbf{PIDag}(\mathcal{L})$  of acyclic graphs and monic pr-morphisms. Note that the chosen morphism  $\mathbf{f}$  in the proof of Proposition 3.13 is monic, since a morphism in  $\mathbf{PDag}(\mathcal{L})$  is monic if and only if the underlying function is injective.

We can then show that rooted graphs are also a family of generators for the sub-category  $\mathbf{PIDag}(\mathcal{L})$ . Before giving the result, we observe that, as is the case for the category of sets and injective functions,  $\mathbf{PIDag}(\mathcal{L})$  lacks pushouts. However, the following simple technical result, analogous to Lemma 3.4, allows us to compute those pushouts in the larger category  $\mathbf{PDag}(\mathcal{L})$

**Lemma 3.14** (Monos under pushouts, 2). *Pushouts in  $\mathbf{PDag}(\mathcal{L})$  preserve monos.*

We can now state the announced important characterisation of  $\mathbf{PIDag}(\mathcal{L})$ .

**Proposition 3.15.**  $\mathbf{PIDag}(\mathcal{L})$  is the smallest sub-category of  $\mathbf{PDag}(\mathcal{L})$  containing all root extension morphisms and closed under finite colimits.

*Proof.* First, note that, since pushouts in  $\mathbf{PDag}(\mathcal{L})$  preserve monos, the smallest sub-category of  $\mathbf{PDag}(\mathcal{L})$  containing all root extensions and closed under finite colimits is surely a sub-category of  $\mathbf{PIDag}(\mathcal{L})$ . So, given a monic pr-morphism  $\mathbf{f} : \mathbf{G}_1 \rightarrow \mathbf{G}_2$ , we need to prove that it can be generated from root extension morphisms via colimits. We proceed by induction on the cardinality of  $\mathcal{E}_{\mathbf{G}_2}$ .

If the cardinality is 0, then  $\mathbf{f}$  must be the identity of the empty graph. Otherwise, consider  $\mathbf{G}_2$  and assume that it is rooted with root  $e$ . Now, if  $e \in \text{img}(\mathbf{f})$ , since the image of a pr-morphism is downward closed, it turns out that  $\mathbf{f}$  is the identity of  $\mathbf{G}_2$ . If it is not in the image, then  $\mathbf{f}$  can be decomposed as  $\mathbf{G}_1 \rightarrow (\mathbf{G}_2 \setminus e) \rightarrow \mathbf{G}_2$ : the left-most is given by induction, while the right-most is a root extension morphism.

Without loss of generality, let us assume now that  $\mathbf{G}_2$  has two distinct roots, namely  $e_1$  and  $e_2$ , and that the image of  $\mathbf{f}$  is contained in  $\mathbf{G}_2|_{[e_1]}$ . Now,  $\mathbf{f}$  can be decomposed as  $\mathbf{G}_1 \rightarrow \mathbf{G}_2|_{[e_1]} \rightarrow \mathbf{G}_2$ : the left-most is given by induction, while the right-most is obtained via the pushout of the span  $\mathbf{G}_2|_{[e_1]} \cap \mathbf{G}_2|_{[e_2]} \rightarrow \mathbf{G}_2|_{[e_1]}$ .  $\square$

### 3.4. Category of arbitrations

We know from Section 2.4 that the objects of our category of arbitrations must be compatible sets of paths. We now turn our attention to morphisms. We first look at two constructions, called *saturation* and *retraction*, which act on arbitrations by extending or restricting them, respectively, in a maximal way, according to a function between sets of events. These constructions will be used to define a notion of morphisms of arbitrations, and to describe limits and colimits in the ensuing category.

**Definition 3.16** (Path saturation). *Let  $P$  be a path and  $f : (\mathcal{E}_P, \lambda_P) \rightarrow (\mathcal{E}, \lambda)$  a label-preserving function. The saturation of  $P$  along  $f$  is defined as*

$$\text{sat}(P, f) = \{Q \mid Q \in \mathbb{P}(\mathcal{E}, \lambda) \text{ and } f \text{ induces a morphism } f : P \rightarrow Q \text{ in } \mathbf{Path}(\mathcal{L})\}$$

*Saturation is generalised to sets of paths  $\mathcal{X} \subseteq \mathbb{P}(\mathcal{E}, \lambda)$  as  $\bigcup_{P \in \mathcal{X}} \text{sat}(P, f)$ .*

Note that, should  $f$  not be injective, it could be that  $\text{sat}(P, f) = \emptyset$ .

**Example 3.17.** *Consider the injective, label-preserving function  $f$  from set  $\{\langle \text{wr}(1), \text{ok} \rangle, \langle \text{wr}(2), \text{ok} \rangle\}$  to set  $\{\langle \text{wr}(1), \text{ok} \rangle, \langle \text{wr}(2), \text{ok} \rangle, \langle \text{rd}, 2 \rangle\}$ . Then, we have*

$$\text{sat} \left( \begin{array}{c} \langle \text{wr}(1), \text{ok} \rangle \\ | \\ \langle \text{wr}(2), \text{ok} \rangle \end{array}, f \right) = \left\{ \begin{array}{ccc} \langle \text{wr}(1), \text{ok} \rangle & \langle \text{wr}(1), \text{ok} \rangle & \langle \text{rd}, 2 \rangle \\ | & | & | \\ \langle \text{wr}(2), \text{ok} \rangle, & \langle \text{rd}, 2 \rangle & \langle \text{wr}(1), \text{ok} \rangle \\ | & | & | \\ \langle \text{rd}, 2 \rangle & \langle \text{wr}(2), \text{ok} \rangle & \langle \text{wr}(2), \text{ok} \rangle \end{array} \right\}$$

*Intuitively, saturation adds  $\langle \text{rd}, 2 \rangle$  – and in general those events that are not in the image of  $f$  – to the original path in all possible ways, preserving the order of the original events.*

**Definition 3.18** (Path retraction). *Let  $Q$  be a path and  $f : \mathcal{E} \rightarrow \mathcal{E}_Q$  a function. The retraction of  $Q$  along  $f$  is defined as*

$$\text{ret}(Q, f) = \{P \mid P \in \mathbb{P}(\mathcal{E}, \lambda) \text{ and } f \text{ induces a morphism } f : P \rightarrow Q \text{ in } \mathbf{Path}(\mathcal{L})\}$$

*The notion of retraction is extended to sets of paths  $\mathcal{X} \subseteq \mathbb{P}(\mathcal{E}, \lambda)$  as  $\bigcup_{Q \in \mathcal{X}} \text{ret}(Q, f)$ .*

Note that  $\lambda$  is fully characterised as the restriction of  $\lambda_Q$  along the mapping. Should  $f$  be injective,  $\text{ret}(Q, f)$  would be a singleton, and if  $f$  is an inclusion, then  $\text{ret}(Q, f) = Q|_{\mathcal{E}}$ .

We may now start considering the relationship between the two notions.

**Lemma 3.19.** *Let  $\mathcal{X}_1 \subseteq \mathbb{P}(\mathcal{E}_1, \lambda_1)$  be a set of paths and  $f : (\mathcal{E}_1, \lambda_1) \rightarrow (\mathcal{E}_2, \lambda_2)$  a label-preserving function. Then  $\mathcal{X}_1 \subseteq \text{ret}(\text{sat}(\mathcal{X}_1, f), f)$ . If  $f$  is injective, then the equality holds.*

**Lemma 3.20.** *Let  $\mathcal{X}_2 \subseteq \mathbb{P}(\mathcal{E}_2, \lambda_2)$  be a set of paths and  $\mathbf{f} : \mathcal{E}_1 \rightarrow \mathcal{E}_2$  a function. Then  $\mathcal{X}_2 \subseteq \mathbf{sat}(\mathbf{ret}(\mathcal{X}_2, \mathbf{f}), \mathbf{f})$ .*

Here, injectiveness does not ensure that the equality holds. We say that an injective function  $\mathbf{f}$  is *saturated* with respect to  $\mathcal{X}_2$  if the equality does hold.

**Example 3.21.** *Consider the set of paths  $\mathcal{X}_1$  and  $\mathcal{X}_2$  and the pr-morphism  $\mathbf{f}$  below*

$$\mathcal{X}_1 = \left\{ \begin{array}{c} \langle \mathbf{wr}(1), ok \rangle \\ | \\ \langle \mathbf{wr}(2), ok \rangle \end{array} \right\} \quad \mathcal{X}_2 = \left\{ \begin{array}{c} \langle \mathbf{wr}(1), ok \rangle \\ | \\ \langle \mathbf{wr}(2), ok \rangle \\ | \\ \langle \mathbf{rd}, 2 \rangle \end{array} \right\} \quad \mathbf{f} : \begin{array}{ccc} \langle \mathbf{wr}(1), ok \rangle & & \langle \mathbf{wr}(1), ok \rangle \\ | & & | \\ \langle \mathbf{wr}(2), ok \rangle & \rightarrow & \langle \mathbf{wr}(2), ok \rangle \\ & & | \\ & & \langle \mathbf{rd}, 2 \rangle \end{array}$$

the underlying function  $\mathbf{f}$  (defined in Example 3.17) is not saturated with respect to  $\mathcal{X}_2$  because

$$\left\{ \begin{array}{c} \langle \mathbf{wr}(1), ok \rangle \\ | \\ \langle \mathbf{wr}(2), ok \rangle \\ | \\ \langle \mathbf{rd}, 2 \rangle \end{array} \right\} \neq \mathbf{sat}(\mathbf{ret}\left(\left\{ \begin{array}{c} \langle \mathbf{wr}(1), ok \rangle \\ | \\ \langle \mathbf{wr}(2), ok \rangle \\ | \\ \langle \mathbf{rd}, 2 \rangle \end{array} \right\}, \mathbf{f}\right), \mathbf{f}) = \mathbf{sat}\left(\left\{ \begin{array}{c} \langle \mathbf{wr}(1), ok \rangle \\ | \\ \langle \mathbf{wr}(2), ok \rangle \end{array} \right\}, \mathbf{f}\right)$$

We can now exploit saturation to get a simple definition of our category of arbitrations.

**Definition 3.22** (ps-morphism). *Let  $\mathcal{X}_1 \subseteq \mathbb{P}(\mathcal{E}_1, \lambda_1)$  and  $\mathcal{X}_2 \subseteq \mathbb{P}(\mathcal{E}_2, \lambda_2)$  be sets of paths. A path-set morphism (shortly, ps-morphism)  $\mathbf{f} : \mathcal{X}_1 \rightarrow \mathcal{X}_2$  is a label-preserving function  $\mathbf{f} : (\mathcal{E}_1, \lambda_1) \rightarrow (\mathcal{E}_2, \lambda_2)$  such that  $\mathcal{X}_2 \subseteq \mathbf{sat}(\mathcal{X}_1, \mathbf{f})$ .*

Intuitively, there is a ps-morphism from the set of paths  $\mathcal{X}_1$  to the set of paths  $\mathcal{X}_2$  if any path in  $\mathcal{X}_2$  can be obtained by adding events to some path in  $\mathcal{X}_1$ . This notion captures the idea that arbitrations of larger visibilities are obtained as extensions of smaller visibilities. We say that mono ps-morphism  $\mathbf{f} : \mathcal{X}_1 \rightarrow \mathcal{X}_2$  is a *ps-extension* (along  $\ell$ ) if its underlying function between events is of the form  $\mathbf{f} : \mathcal{E}_{\mathcal{X}_1} \rightarrow \mathcal{E}_{\mathcal{X}_1} \cup \{\top\}$ , with  $\ell$  labelling the fresh event  $\top$ .

**Example 3.23.** *Consider the following three sets and the function  $\mathbf{f}$  from Example 3.17*

$$\mathcal{X}_1 = \left\{ \begin{array}{c} \langle \mathbf{wr}(1), ok \rangle \\ | \\ \langle \mathbf{wr}(2), ok \rangle \end{array} \right\} \quad \mathcal{X}_2 = \left\{ \begin{array}{cc} \langle \mathbf{wr}(1), ok \rangle & \langle \mathbf{wr}(1), ok \rangle \\ | & | \\ \langle \mathbf{wr}(2), ok \rangle, & \langle \mathbf{rd}, 2 \rangle \\ | & | \\ \langle \mathbf{rd}, 2 \rangle & \langle \mathbf{wr}(2), ok \rangle \end{array} \right\} \quad \mathcal{X}_3 = \left\{ \begin{array}{cc} \langle \mathbf{wr}(1), ok \rangle & \langle \mathbf{wr}(2), ok \rangle \\ | & | \\ \langle \mathbf{wr}(2), ok \rangle, & \langle \mathbf{rd}, 2 \rangle \\ | & | \\ \langle \mathbf{rd}, 2 \rangle & \langle \mathbf{wr}(1), ok \rangle \end{array} \right\}$$

Now,  $\mathbf{f}$  induces a ps-morphism  $\mathbf{f} : \mathcal{X}_1 \rightarrow \mathcal{X}_2$  because  $\mathcal{X}_2 \subseteq \mathbf{sat}(\mathcal{X}_1, \mathbf{f})$  (the latter is shown in Example 3.17). On the contrary, there is no ps-morphism from  $\mathcal{X}_1$  to  $\mathcal{X}_3$ : the rightmost path of  $\mathcal{X}_3$  cannot be obtained by extending a path of  $\mathcal{X}_1$  with an event labelled by  $\langle \mathbf{rd}, 2 \rangle$ . Note that the ps-morphism induced by  $\mathbf{f}$  is a ps-extension along  $\langle \mathbf{rd}, 2 \rangle$ .

**Definition 3.24** (Sets of Paths Category). We define  $\mathbf{SPath}(\mathcal{L})$  as the category whose objects are sets of paths labelled over  $\mathcal{L}$  and arrows are ps-morphisms.

**Proposition 3.25** (Properties of  $\mathbf{SPath}$ ). The category  $\mathbf{SPath}(\mathcal{L})$  has finite colimits along monos and binary pullbacks.

*Proof.* (Strict) initial object. The (unique) initial object is  $\langle \emptyset, \{\epsilon\}, \emptyset \rangle$ , with  $\epsilon \in \mathbb{P}(\emptyset, \emptyset)$  the empty path. Let  $\mathcal{X} \subseteq \mathbb{P}(\mathcal{E}, \lambda)$  and  $! : \emptyset \rightarrow \mathcal{E}$  the unique function from the empty set to  $\mathcal{E}$ . We have the (unique) function  $! : (\emptyset, \emptyset) \rightarrow (\mathcal{E}, \lambda)$  such that  $\mathcal{X} \subseteq \mathbf{sat}(\{\epsilon\}, !) = \mathbb{P}(\mathcal{E}, \lambda)$ .

*Binary Pushouts.* Let  $\mathcal{X}, \mathcal{X}_1$ , and  $\mathcal{X}_2$  be sets of paths and  $\mathbf{f}_i : \mathcal{X} \rightarrow \mathcal{X}_i$  ps-morphisms. Consider the underlying functions  $\mathbf{f}_i : \mathcal{E} \rightarrow \mathcal{E}_i$  and their pushout  $\mathbf{f}'_i : \mathcal{E}_i \rightarrow \mathcal{E}_1 +_{\mathcal{E}} \mathcal{E}_2$  in the category of sets and functions: it induces a pushout  $\mathbf{f}'_i : \mathcal{X}_i \rightarrow \mathbf{sat}(\mathcal{X}_1, \mathbf{f}'_1) \cap \mathbf{sat}(\mathcal{X}_2, \mathbf{f}'_2)$  in  $\mathbf{SPath}(\mathcal{L})$ .

*Binary Pullbacks.* Let  $\mathcal{X}, \mathcal{X}_1$ , and  $\mathcal{X}_2$  be sets of paths and  $\mathbf{f}_i : \mathcal{X}_i \rightarrow \mathcal{X}$  ps-morphisms. Consider the underlying functions  $\mathbf{f}_i : \mathcal{E}_i \rightarrow \mathcal{E}$  and their pullback  $\mathbf{f}'_i : \mathcal{E}_1 \times_{\mathcal{E}} \mathcal{E}_2 \rightarrow \mathcal{E}$  in the category of sets: It induces a pullback  $\mathbf{f}'_i : \mathbf{ret}(\mathcal{X}_1, \mathbf{f}'_1) \cup \mathbf{ret}(\mathcal{X}_2, \mathbf{f}'_2) \rightarrow \mathcal{X}_i$  in  $\mathbf{SPath}(\mathcal{L})$ .  $\square$

The above characterisation of pushouts is enabled by the fact that we considered injective functions. To help intuition, we now instantiate that characterisation to suitable inclusions. We obtain that pushouts of inclusions precisely amount to products of sets of paths, as defined in Definition 2.8, which is a crucial step towards characterising coherent specifications.

**Lemma 3.26.** Let  $\mathbf{f}_i : \mathcal{X} \rightarrow \mathcal{X}_i$  be ps-morphisms such that the underlying functions  $\mathbf{f}_i : \mathcal{E} \rightarrow \mathcal{E}_i$  are inclusions and  $\mathcal{E} = \mathcal{E}_1 \cap \mathcal{E}_2$ . Then their pushout is given by  $\mathbf{f}'_i : \mathcal{X}_i \rightarrow \mathcal{X}_1 \otimes \mathcal{X}_2$ .

*Proof.* By definition  $\mathcal{X}_1 \otimes \mathcal{X}_2 = \{\mathbf{P} \mid \mathbf{P} \text{ is a path over } \bigcup_i \mathcal{E}_i \text{ and } \mathbf{P}|_{\mathcal{E}_i} \in \mathcal{X}_i\}$ . Note also that  $\mathbf{sat}(\mathcal{X}_i, \mathbf{f}'_i) = \bigcup_{\mathbf{Q} \in \mathcal{X}_i} \{\mathbf{P} \mid \mathbf{P} \in \mathbb{P}(\bigcup_i \mathcal{E}_i, \bigcup_i \lambda_i) \text{ and } \mathbf{f}'_i \text{ induces a path morphism } \mathbf{f}'_i : \mathbf{P} \rightarrow \mathbf{Q}\}$ . Since  $\mathbf{f}'_i$  is an inclusion, the latter condition equals to  $\mathbf{P}|_{\mathcal{E}_i} = \mathbf{Q}$ , thus the property holds.  $\square$

**Example 3.27.** Consider the following ps-morphisms

$$\left\{ \begin{array}{c} \langle \mathbf{wr}(1), ok \rangle \\ | \\ \langle \mathbf{wr}(2), ok \rangle \\ | \\ \langle \mathbf{rd}, 2 \rangle \end{array} \right\} \leftarrow \left\{ \begin{array}{c} \langle \mathbf{wr}(1), ok \rangle \\ | \\ \langle \mathbf{wr}(2), ok \rangle \end{array} \right\} \rightarrow \left\{ \begin{array}{cc} \langle \mathbf{wr}(1), ok \rangle & \langle \mathbf{wr}(2), ok \rangle \\ | & | \\ \langle \mathbf{wr}(2), ok \rangle, & \langle \mathbf{wr}(1), ok \rangle \\ | & | \\ \langle \mathbf{rd}, 1 \rangle & \langle \mathbf{rd}, 1 \rangle \end{array} \right\}$$

then, the pushout is given by the following ps-morphisms

$$\left\{ \begin{array}{c} \langle \mathbf{wr}(1), ok \rangle \\ | \\ \langle \mathbf{wr}(2), ok \rangle \\ | \\ \langle \mathbf{rd}, 2 \rangle \end{array} \right\} \rightarrow \left\{ \begin{array}{cc} \langle \mathbf{wr}(1), ok \rangle & \langle \mathbf{wr}(1), ok \rangle \\ | & | \\ \langle \mathbf{wr}(2), ok \rangle & \langle \mathbf{wr}(2), ok \rangle \\ | & | \\ \langle \mathbf{rd}, 1 \rangle & \langle \mathbf{rd}, 2 \rangle \\ | & | \\ \langle \mathbf{rd}, 2 \rangle & \langle \mathbf{rd}, 1 \rangle \end{array} \right\} \leftarrow \left\{ \begin{array}{cc} \langle \mathbf{wr}(1), ok \rangle & \langle \mathbf{wr}(2), ok \rangle \\ | & | \\ \langle \mathbf{wr}(2), ok \rangle, & \langle \mathbf{wr}(1), ok \rangle \\ | & | \\ \langle \mathbf{rd}, 1 \rangle & \langle \mathbf{rd}, 1 \rangle \end{array} \right\}$$

An analogous property holds for pullbacks. Let  $\mathbf{f}_i : \mathcal{X}_i \rightarrow \mathcal{X}$  be ps-morphisms such that the underlying functions are inclusions: The pullback is given as  $\mathbf{f}'_i : \bigcup_i \mathcal{X}_i|_{\mathcal{E}_1 \cap \mathcal{E}_2} \rightarrow \mathcal{X}_i$ . In particular, the square below is both a pullback and a pushout

$$\begin{array}{ccc} \bigcup_i \mathcal{X}_i|_{\mathcal{E}_1 \cap \mathcal{E}_2} & \hookrightarrow & \mathcal{X}_1 \\ \downarrow & & \downarrow \\ \mathcal{X}_2 & \hookrightarrow & \mathcal{X}_1 \otimes \mathcal{X}_2 \end{array}$$

### 3.5. Coherent Functors

It is now time for moving towards our categorical characterisation of coherent specifications. We first provide a simple technical result.

**Lemma 3.28.** *Let  $\mathcal{S}$  be a coherent specification and  $\mathcal{E} \subseteq \mathcal{E}_{\mathbf{G}}$ . If  $\mathcal{E}$  is downward closed, then  $\mathcal{S}(\mathbf{G})|_{\mathcal{E}} \subseteq \mathcal{S}(\mathbf{G}|_{\mathcal{E}})$ .*

*Proof.* Let  $\mathcal{E}$  be downward closed, and note that this amounts to requiring  $\mathcal{E} = \bigcup_{e \in \mathcal{E}} [e]$ , hence for all  $e \in \mathcal{E}$  we have that  $(\mathbf{G}|_{\mathcal{E}})|_{[e]} = \mathbf{G}|_{[e]}$ . By the latter and by coherence we have  $\mathcal{S}(\mathbf{G})|_{\mathcal{E}} = \left( \bigotimes_{e \in \mathcal{E}_g} \mathcal{S}(\mathbf{G}|_{[e]}) \right)|_{\mathcal{E}}$  and  $\mathcal{S}(\mathbf{G}|_{\mathcal{E}}) = \bigotimes_{e \in \mathcal{E}} \mathcal{S}(\mathbf{G}|_{[e]})$ . Note that  $\left( \bigotimes_{e \in \mathcal{E}_g} \mathcal{S}(\mathbf{G}|_{[e]}) \right)|_{\mathcal{E}} \subseteq \bigotimes_{e \in \mathcal{E}} \mathcal{S}(\mathbf{G}|_{[e]})$  because a path in the former can always be restricted to a suitable path with fewer events on the latter (the converse in general does not hold).  $\square$

We now introduce our functorial models of specifications, consisting of functors preserving relevant categorical structure. In the following we say that a functor  $\mathbb{F} : \mathbf{PIDag}(\mathcal{L}) \rightarrow \mathbf{SPath}(\mathcal{L})$  *weakly preserves* finite colimits if any commuting diagram in  $\mathbf{PIDag}(\mathcal{L})$  that is a colimit (via the inclusion functor) in  $\mathbf{PDag}(\mathcal{L})$  is mapped by  $\mathbb{F}$  to a colimit in  $\mathbf{SPath}(\mathcal{L})$ .

**Definition 3.29.** *A functor  $\mathbb{F} : \mathbf{PIDag}(\mathcal{L}) \rightarrow \mathbf{SPath}(\mathcal{L})$  is coherent if it maps root  $\ell$ -extensions to  $\ell$ -extensions and weakly preserves finite colimits.*

The rest of this section is dedicated to proving that coherent functors are a sound and complete characterisation of coherent specifications. More specifically, we shall prove that coherent specifications induce coherent functors (soundness) and vice versa (completeness). To this purpose, we need the following two technical lemmas. The first one spells out two important properties of coherent functors: They preserve monos and map graphs to sets of paths of the correct type. The second one says that, for each graph  $\mathbf{G}$ , a coherent functor gives a set of paths which is the product of paths for sub-graphs of  $\mathbf{G}$ , as required by coherence (Definition 2.9).

**Lemma 3.30.** *Let  $\mathbb{F}$  be a coherent functor. Then*

- (i)  $\mathbb{F}$  preserves monos;

(ii)  $\mathbb{F}(\mathbf{G})$  is a set of paths over  $(\mathcal{E}_{\mathbf{G}}, \lambda_{\mathbf{G}})$ , for all  $\mathbf{G}$  in  $\mathbf{PIDag}(\mathcal{L})$ .

*Proof.* Recall that morphisms of  $\mathbf{PIDag}(\mathcal{L})$  are generated from root  $\ell$ -extensions via finite colimits (Proposition 3.15). One can show that  $\mathbb{F}$  maps  $\mathbf{f}: \mathbf{G}_1 \hookrightarrow \mathbf{G}_2$  to a monic ps-morphism by induction on the number of vertices of  $\mathbf{G}_2$ , using the fact that  $\mathbb{F}$  maps root  $\ell$ -extensions to  $\ell$ -extensions, which are monic, and weakly preserves finite colimits, hence it maps pushouts of monos to pushouts of monos. Point (ii) can be shown via an analogous argument, noting that  $\mathbb{F}$  preserves the label of the additional event when mapping root  $\ell$ -extensions to  $\ell$ -extensions.  $\square$

**Lemma 3.31.** *Let  $\mathbb{F}$  be a coherent functor and  $\mathbf{f}_i: \mathbf{G} \hookrightarrow \mathbf{G}_i$  ( $i = 1, 2$ ) mono ps-morphisms such that  $\mathcal{E}_{\mathbf{G}} = \mathcal{E}_{\mathbf{G}_1} \cap \mathcal{E}_{\mathbf{G}_2}$ . Then there exists a pushout in  $\mathbf{SPath}(\mathcal{L})$*

$$\begin{array}{ccc} \mathbb{F}(\mathbf{G}) & \xleftarrow{\mathbb{F}(\mathbf{f}_1)} & \mathbb{F}(\mathbf{G}_1) \\ \mathbb{F}(\mathbf{f}_2) \downarrow & & \downarrow \\ \mathbb{F}(\mathbf{G}_2) & \hookrightarrow & \mathbb{F}(\mathbf{G}_1) \otimes \mathbb{F}(\mathbf{G}_2) \end{array}$$

*Proof.*  $\mathbb{F}$  preserves monos, by Lemma 3.30, so  $\mathbb{F}(\mathbf{f}_i): \mathbb{F}(\mathbf{G}) \hookrightarrow \mathbb{F}(\mathbf{G}_i)$  are monos and by Lemma 3.26 their pushout is the diagram above.  $\square$

We now state our soundness result.

**Theorem 3.32** (Soundness). *A coherent specification  $\mathcal{S}$  induces a coherent functor  $\mathbb{M}(\mathcal{S}): \mathbf{PIDag}(\mathcal{L}) \rightarrow \mathbf{SPath}(\mathcal{L})$ .*

*Proof.* For  $\mathbf{G}$  we define  $\mathbb{M}(\mathcal{S})(\mathbf{G})$  as  $\mathcal{S}(\mathbf{G})$  and for  $\mathbf{f}: \mathbf{G} \rightarrow \mathbf{G}'$  we define  $\mathbb{M}(\mathcal{S})(\mathbf{f})$  as the ps-morphism with underlying injective function  $\mathbf{f}: (\mathcal{E}_{\mathbf{G}}, \lambda_{\mathbf{G}}) \hookrightarrow (\mathcal{E}_{\mathbf{G}'}, \lambda_{\mathbf{G}'})$ .

We first show that  $\mathbf{f}$  is a valid ps-morphism from  $\mathcal{S}(\mathbf{G})$  into  $\mathcal{S}(\mathbf{G}')$ , i.e.,  $\mathcal{S}(\mathbf{G}') \subseteq \mathbf{sat}(\mathcal{S}(\mathbf{G}), \mathbf{f})$  and, since we are considering specifications preserving isomorphisms, we can restrict our attention to the case where  $\mathbf{f}$  is an inclusion. Since  $\mathbf{f}$  is a pr-morphism,  $\bigcup_{e \in \mathcal{E}_{\mathbf{G}}} \mathbf{f}(e)$  is downward-closed in  $\mathbf{G}'$  and thus by Lemma 3.28 we have  $\mathcal{S}(\mathbf{G}')|_{\mathcal{E}_{\mathbf{G}}} \subseteq \mathcal{S}(\mathbf{G}'|_{\mathcal{E}_{\mathbf{G}}}) = \mathcal{S}(\mathbf{G})$ , the latter equality given by coherence. Now, consider a path  $\mathbf{P} \in \mathcal{S}(\mathbf{G}')$ . Since  $\mathbf{P}|_{\mathcal{E}_{\mathbf{G}}} \in \mathcal{S}(\mathbf{G})$ , we have  $\mathbf{P} \in \mathbf{sat}(\mathcal{S}(\mathbf{G}), \mathbf{f})$ , because saturation adds missing events – namely those in  $\mathcal{E}_{\mathbf{G}'} \setminus \mathcal{E}_{\mathbf{G}}$  – to  $\mathbf{P}|_{\mathcal{E}_{\mathbf{G}}}$  in all possible ways. Therefore we can conclude  $\mathcal{S}(\mathbf{G}') \subseteq \mathbf{sat}(\mathcal{S}(\mathbf{G}), \mathbf{f})$ . Note that, if  $\mathbf{f}'$  is a root  $\ell$ -extension,  $\mathbb{M}(\mathcal{S})(\mathbf{f})$  must be an  $\ell$ -extension, as it has the same underlying function between events.

Finally, we need to show that  $\mathbb{M}(\mathcal{S})$  weakly preserves colimits. Preservation of the initial object is easy, since  $\mathcal{S}(\epsilon) = \{\epsilon\}$  by definition (see Definition 2.1) and the right hand side is the initial object in  $\mathbf{SPath}(\mathcal{L})$ . As for pushouts: Since  $\mathcal{S}$  is coherent, preservation boils down to Lemma 3.26.  $\square$

We now give our completeness result.

**Theorem 3.33.** *Any coherent functor  $\mathbb{F}: \mathbf{PIDag}(\mathcal{L}) \rightarrow \mathbf{SPath}(\mathcal{L})$  induces a coherent specification  $\mathbb{S}(\mathbb{F})$ .*

*Proof.* The claim follows from Lemma 3.31 by setting  $\mathbb{S}(\mathbb{F})(\mathbf{G}) = \mathbb{F}(\mathbf{G})$ .  $\square$

### 3.6. More completeness

We can sharpen the completeness result by requiring functors to preserve specific properties for suitable arrows of  $\mathbf{PIDag}(\mathcal{L})$ . The candidates are root extension morphisms, given the properties shown in Section 3.3. In order to define the functors, we also need to consider a subset of the arrows of  $\mathbf{SPath}(\mathcal{L})$ .

**Definition 3.34** (Saturated specifications). *Let  $\mathcal{S}$  be a specification. It is saturated if for all graphs  $\mathbb{G}$ , the function  $\mathbf{f} : \mathcal{E}_{\mathbb{G}} \rightarrow \mathcal{E}_{\mathbb{G}^\ell}$  is saturated with respect to  $\mathcal{S}(\mathbb{G}^\ell)$  (see Lemma 3.20), that is*

$$\forall \mathbb{G}, \mathcal{E}, \ell. \mathcal{S}(\mathbb{G}^\ell) = \mathbf{sat}(\mathbf{ret}(\mathcal{S}(\mathbb{G}^\ell), \mathbf{f}), \mathbf{f})$$

A *saturation ps-morphism* (along  $\ell$ ) is a ps-extension  $\mathbf{f} : \mathcal{X}_1 \rightarrow \mathcal{X}_2$  along  $\ell$  whose underlying function  $\mathbf{f} : \mathcal{E} \rightarrow \mathcal{E} \cup \{\top\}$  is saturated. We can now prove an instance of Theorem 3.33 concerning saturated specifications.

**Proposition 3.35.** *Let  $\mathbb{F} : \mathbf{PIDag}(\mathcal{L}) \rightarrow \mathbf{SPath}(\mathcal{L})$  be a coherent functor. If it maps root extension morphisms to saturation ps-morphisms (along the same labels), then the induced specification  $\mathbb{S}(\mathbb{F})$  is saturated and coherent.*

*Proof.* We know by Theorem 3.33 that  $\mathbb{S}(\mathbb{F})$  is coherent. It remains to be shown that  $\mathbb{S}(\mathbb{F})$  is saturated, that is  $\mathbb{F}(\mathbb{G}^\ell) = \mathbf{sat}(\mathbf{ret}(\mathbb{F}(\mathbb{G}^\ell), \mathbf{f}), \mathbf{f})$ . If  $\mathbb{G}^\ell$  is rooted, this follows from  $\mathbb{F}$  mapping root extensions to saturation ps-morphisms. Otherwise, by coherence,  $\mathbb{F}(\mathbb{G}^\ell)$  can be decomposed into the product  $\bigotimes_{e \in (\mathcal{E}_{\mathbb{G}})_{\top}} \mathbb{F}(\mathbb{G}^\ell|_{[e]})$ . For each component of the product we have a root extension  $\mathbb{G}^\ell|_{[e]} \setminus e \rightarrow \mathbb{G}^\ell|_{[e]}$ , which is mapped by  $\mathbb{F}$  to a saturation ps-morphism, therefore we have  $\mathbb{F}(\mathbb{G}^\ell|_{[e]}) = \mathbf{sat}(\mathbf{ret}(\mathbb{F}(\mathbb{G}^\ell|_{[e]}), \mathbf{f}_e), \mathbf{f}_e)$ , where  $\mathbf{f}_e$  is the underlying function between events of the root extension. Saturation of  $\mathbb{F}(\mathbb{G}^\ell)$  follows by computing the product of these sets of paths.  $\square$

The need of finding a suitable image for root extension morphisms allows for alternative choices. To this end, we introduce a different subset of the arrows of  $\mathbf{SPath}(\mathcal{L})$ .

**Definition 3.36** (Path extension/prefixing). *Let  $\mathbb{P}$  be a path and  $\mathbf{f} : (\mathcal{E}_{\mathbb{P}}, \lambda_{\mathbb{P}}) \rightarrow (\mathcal{E}, \lambda)$  a function preserving labels. The extension of  $\mathbb{P}$  along  $\mathbf{f}$  is defined as*

$$\mathbf{ext}(\mathbb{P}, \mathbf{f}) = \{\mathbb{Q} \mid \mathbb{Q} \in \mathbb{P}(\mathcal{E}, \lambda) \text{ and } \mathbf{f} \text{ induces a pr-morphism } \mathbf{f} : \mathbb{P} \rightarrow \mathbb{Q} \text{ in } \mathbf{Path}(\mathcal{L})\}$$

*Similarly, let  $\mathbb{Q}$  be a path and  $\mathbf{f} : \mathcal{E} \rightarrow \mathcal{E}_{\mathbb{Q}}$  a function preserving labels. The prefixing of  $\mathbb{Q}$  along  $\mathbf{f}$  is defined as*

$$\mathbf{pre}(\mathbb{Q}, \mathbf{f}) = \{\mathbb{P} \mid \mathbb{P} \in \mathbb{P}(\mathcal{E}, \lambda) \text{ and } \mathbf{f} \text{ induces a pr-morphism } \mathbf{f} : \mathbb{P} \rightarrow \mathbb{Q} \text{ in } \mathbf{Path}(\mathcal{L})\}$$

Both definitions immediately extend to sets of paths. Should  $\mathbf{f}$  be injective,  $\mathbf{pre}(\mathbb{Q}, \mathbf{f})$  would be a singleton, and if  $\mathbf{f}$  is an inclusion, then  $\mathbf{pre}(\mathbb{Q}, \mathbf{f}) = \mathbb{Q}|_{\mathcal{E}}$ , for the latter a prefix of  $\mathbb{Q}$ . Also, note that similarly  $\mathbb{P}$  has to be a prefix for all the paths in  $\mathbf{ext}(\mathbb{P}, \mathbf{f})$ .

**Example 3.37.** A topological specification  $\mathcal{S}_{topR}$  for a **Register** can be defined as  $\mathcal{S}_{lwwR}$  in example 2.2 with the additional requirement that paths are topological orderings of visibilities

$$P \in \mathcal{S}_{topR}(\mathbf{G}) \text{ iff } P \in \mathcal{S}_{lwwR}(\mathbf{G}) \text{ and } \prec_{\mathbf{G}} \subseteq \leq_P$$

In this way,  $\mathcal{S}_{topR}(\mathbf{G})$  excludes e.g. the two right-most arbitrations of the equation in Figure 1a.

**Definition 3.38** (Topological specifications). Let  $\mathcal{S}$  be a specification. It is topological if

$$\forall \mathbf{G}, \mathcal{E}, \ell. \mathcal{S}(\mathbf{G}_{\mathcal{E}}^{\ell}) = \text{ext}(\text{pre}(\mathcal{S}(\mathbf{G}_{\mathcal{E}}^{\ell}), \mathbf{f}), \mathbf{f})$$

A topological ps-morphism (along  $\ell$ ) is a ps-extension  $\mathbf{f} : \mathcal{X}_1 \rightarrow \mathcal{X}_2$  with underlying function  $\mathbf{f} : \mathcal{E} \rightarrow \mathcal{E} \cup \{\top\}$  such that  $\mathcal{X}_2 = \text{ext}(\text{pre}(\mathcal{S}(\mathcal{X}_2), \mathbf{f}), \mathbf{f})$ .

The name is directly reminiscent of what are called topological RDTs in [14, 8], and in fact it similarly guarantees that arbitrations preserve the visibility order. We can thus prove another instance of Theorem 3.33, now concerning topological specifications.

**Proposition 3.39.** Let  $\mathbb{F} : \mathbf{PIDag}(\mathcal{L}) \rightarrow \mathbf{SPath}(\mathcal{L})$  be a coherent functor. If it maps root extension morphisms into topological ps-morphisms (along the same labels), then the induced specification  $\mathbb{S}(\mathbb{F})$  is coherent and topological.

### 3.7. Interchangeability of Functors and Coherent Specifications

The connection between the constructions of Theorem 3.32 and Theorem 3.33 is quite tight, and in fact induces a one-to-one correspondence between functors and coherent specifications.

**Theorem 3.40.** Let  $\mathcal{S}$  be a coherent specification. Then  $\mathbb{S}(\mathbb{M}(\mathcal{S})) = \mathcal{S}$ . Conversely, let  $\mathbb{F} : \mathbf{PIDag}(\mathcal{L}) \rightarrow \mathbf{SPath}(\mathcal{L})$  be a coherent functor. Then  $\mathbb{M}(\mathbb{S}(\mathbb{F})) \simeq \mathbb{F}$ .

*Proof.* We first show that  $\mathbb{M}(\mathbb{S}(\mathbb{F})) \simeq \mathbb{F}$ . For notational convenience, we denote  $\mathbb{M}(\mathbb{S}(\mathbb{F}))$  by  $\mathbb{M}'$ . We will show the existence of a natural isomorphism  $\varphi : \mathbb{M}' \Rightarrow \mathbb{F}$ . By definition, we have  $\mathbb{M}'(\mathbf{G}) = \mathbb{S}(\mathbb{F})(\mathbf{G}) = \mathbb{F}(\mathbf{G})$ , therefore we can define  $\varphi_{\mathbf{G}} = \text{Id}_{\mathbb{F}(\mathbf{G})}$ . We need to prove that it is natural, which in this case amounts to showing  $\mathbb{M}'(\mathbf{f}) = \mathbb{F}(\mathbf{f})$  for  $\mathbf{f} : \mathbf{G} \rightarrow \mathbf{G}'$  in  $\mathbf{PIDag}(\mathcal{L})$ . This follows from  $\mathbb{M}'(\mathbf{f})$  and  $\mathbb{F}(\mathbf{f})$  having the same underlying function between events, namely the inclusion  $(\mathcal{E}_{\mathbf{G}}, \lambda_{\mathbf{G}}) \rightarrow (\mathcal{E}_{\mathbf{G}'}, \lambda_{\mathbf{G}'})$ .

Now we show that  $\mathbb{S}(\mathbb{M}(\mathcal{S})) = \mathcal{S}$  for any coherent specification  $\mathcal{S}$ . This follows directly from the definition of  $\mathbb{M}$  and  $\mathbb{S}$ . In fact,  $\mathbb{S}(\mathbb{M}(\mathcal{S}))(\mathbf{G}) = \mathbb{M}(\mathcal{S})(\mathbf{G}) = \mathcal{S}(\mathbf{G})$ .  $\square$

The one-to-one correspondence can be lifted to the specific classes of saturated and topological coherent specifications and to the functors of Proposition 3.35 and Proposition 3.39, respectively. However, what is most relevant is the fact the interchangeability allows one to leverage the categorical machinery of the functor category for providing operators on specifications.






Figure 5: Execution

**Remark 3.41.** Besides coherence, one of the keys of the previous correspondence is the (quite reasonable) choice of specifications that preserve isomorphisms. In general terms, whenever one needs to consider the relationship between different specifications, it is necessary to take into account how the underlying sets of events are related. This is quite easy to accomplish if we move to the functorial presentation. For example, we can say that a specification  $\mathcal{S}_1$  refines a specification  $\mathcal{S}_2$  if  $\mathcal{S}_1(\mathbf{G}) \subseteq \mathcal{S}_2(\mathbf{G})$  for all graphs  $\mathbf{G}$ . However, this is a very concrete characterisation: It would be more general to check for the existence of a ps-morphism  $\mathcal{S}_2(\mathbf{G}_2) \rightarrow \mathcal{S}_1(\mathbf{G}_1)$  whose underlying function  $\mathbf{f} : \mathcal{E}_{\mathbf{G}_2} \rightarrow \mathcal{E}_{\mathbf{G}_1}$  is a bijection, in order to abstract from the identities of the events. In this case, a further constraint would be that  $\mathbf{f}$  is preserved along the image of the morphisms in  $\mathbf{PIDag}(\mathcal{L})$ . These requirements boil down to the existence of a natural transformation  $\mathbb{M}(\mathcal{S}_2) \rightarrow \mathbb{M}(\mathcal{S}_1)$ .

#### 4. On the implementations of Replicated Data Types

An RDT is implemented on top of a set of replicas, which serve requests from clients according to their local state and communicate asynchronously their local changes. Fig. 5 illustrates a scenario involving two replicas, namely  $r_1$  and  $r_2$ , that implement a *multi-value Register* (as specified in Example 2.4). A horizontal line corresponds to a replica and shows the relative order (from left to right) in which events occur in that replica. The depicted scenario shows a concrete execution that generates the visibility graph in Fig. 3. The two writes on  $r_1$  are totally ordered (events **1** and **3**), and consequently, 2 overwrites 1. The remaining write takes place on  $r_2$  (event **2**) and is unknown to  $r_1$  until  $r_2$  propagates its changes. Hence, the first read on  $r_1$  (event **5**) returns 2, which is the last written value in  $r_1$ . Replicas communicate their local changes by using primitives  $\text{snd}$  and  $\text{rcv}$ . Event **6** in  $r_1$  denotes the synchronisation of its local state with the state of  $r_2$ , i.e.,  $r_1$  becomes aware of the written value 3 (depicted by the dashed line between events **4** and **6**). Since writes in  $r_1$  and  $r_2$  are concurrent, the last read on  $r_1$  returns the set of maximal concurrent updates, i.e.,  $\{2, 3\}$ .

A crucial aspect in the implementation of RDTs concerns the information exchanged through  $\text{snd}$  and  $\text{rcv}$ . Under the *state-based* approach, replicas communicate their own local states [7] while they only communicate operations (or their effects) under the *operation-based* approach [20]. Hereafter, we will focus on state-based implementations.

#### 4.1. Implementations as RLTSs

Following the approach in [7], we describe an implementation of a data type in terms of the behaviour of a replica and we assume that all the replicas of an implementation behave the same.

We define the behaviour of a replica implementing a specification as a labelled transition system, which we call *Replica LTS* (RLTS). We write  $\Sigma$  for the set of the possible states  $\sigma, \sigma_0, \dots$  of a replica. We assume that  $\Sigma$  forms a *commutative monoid*, namely that it can be equipped with an operation  $\oplus$  describing how one replica updates its state upon reception of the local state of another replica. We deviate from previous proposals, such as [7, 9], and treat operations that may change the state of a replica (i.e., *mutators*) differently from those that can only retrieve information from the state (i.e., *queries*). The foremost reason of this distinction is behavioural: The mutators are total, meaning that they can originate from any state; furthermore, they respect the monoidal structure of states. By defining RLTSs over mutators, we will be able to capture the compositional core of RDT implementations. Queries are recovered in our notion of bisimulation for RLTSs; this will be akin to higher-order bisimulation (see, e.g., [16]), where queries are modelled as *barbs*, i.e., predicates on states. Hereafter, we use  $\mathcal{M}$  and  $\mathcal{B}$  to respectively denote the sets of labels associated with mutators and queries.

**Definition 4.1** (Replica LTS). *A replica LTS (RLTS in short) on a set of actions  $\mathcal{M}$  is a quadruple  $\mathcal{Q} = \langle \Sigma, \oplus, 1, \rightsquigarrow \rangle$  such that  $\langle \Sigma, \oplus, 1 \rangle$  is a (commutative) monoid and  $\rightsquigarrow \subseteq \Sigma \times (\mathcal{M} \cup \Sigma) \times \Sigma$  is a transition relation satisfying*

$$\text{(RCV)} \quad \frac{}{\sigma \overset{\sigma'}{\rightsquigarrow} \sigma \oplus \sigma'}$$

*It is action-deterministic if the relation  $\rightsquigarrow_{\mathcal{M}} = \rightsquigarrow \cap (\Sigma \times \mathcal{M} \times \Sigma)$  is a function  $(\Sigma \times \mathcal{M}) \rightsquigarrow \Sigma$ , and it is state-deterministic if the relation  $\rightsquigarrow_{\Sigma} = \rightsquigarrow \cap (\Sigma \times \Sigma \times \Sigma)$  is a function  $(\Sigma \times \Sigma) \rightsquigarrow \Sigma$ .*

According to the rule (RCV), when receiving a state  $\sigma'$  from other replica, the local state  $\sigma$  is updated by combining the information in  $\sigma$  and  $\sigma'$ . Thus, the concrete LTS defining the behaviour of a replica that implements an RDT is obtained by instantiating the definitions of state and state combination.

**Remark 4.2.** *Note that for state-deterministic RLTSs the relation  $\rightsquigarrow$  coincides with the closure of  $\rightsquigarrow_{\mathcal{M}}$  with respect to the (RCV) rule.*

Indeed, all of our examples will be state-deterministic, albeit not action-deterministic. Notably, they will preserve the monoidal operator, i.e., the transition relation will be compositional with respect to state combination.

**Definition 4.3** (monoidal RLTS). *A monoidal RLTS  $\mathcal{Q} = \langle \Sigma, \oplus, 1, \rightsquigarrow \rangle$  is a state-deterministic RLTS such that the transition relation satisfies*

$$1 \text{ if } \sigma_1 \overset{m}{\rightsquigarrow} \sigma'_1 \text{ and } \sigma_2 \overset{m}{\rightsquigarrow} \sigma'_2 \text{ then } \sigma_1 \oplus \sigma_2 \overset{m}{\rightsquigarrow} \sigma'_1 \oplus \sigma'_2$$

2 if  $\sigma_1 \oplus \sigma_2 \xrightarrow{m} \sigma$  then  $\exists \sigma'_1, \sigma'_2. \sigma_1 \xrightarrow{m} \sigma'_1, \sigma_2 \xrightarrow{m} \sigma'_2$ , and  $\sigma = \sigma'_1 \oplus \sigma'_2$

**Remark 4.4.** In other terms, in our RLTS we say for a given state how it may a) evolve towards a different state and b) be combined with other states. This is thus more general than the lattice of states (thus with union as state composition) proposed in [20]. First of all, we implicitly consider a pre-order  $\rightsquigarrow$  on states instead of a partial order, and furthermore we do not require the state composition operator  $\otimes$  to be induced by the order relation on states. This weakening results in an algebraic structure that allows for modelling a large family of RDTs, as shown in Section 4.3.

We can now move to denote the category of our RLTSs.

**Definition 4.5** (category of RLTSs). *The category **RLTS** has RLTSs as objects and transition- and monoid-preserving functions as arrows, that is,  $f : \mathcal{Q}_1 \rightarrow \mathcal{Q}_2$  is a monoid homomorphism  $f : \langle \Sigma_1, \oplus_1 \rangle \rightarrow \langle \Sigma_2, \oplus_2 \rangle$  such that  $f(\rightsquigarrow_1) \subseteq \rightsquigarrow_2$ .*

**Remark 4.6.** Note that if  $\mathcal{Q}_2$  is state-deterministic, then any transition-preserving function  $f : \Sigma_1 \rightarrow \Sigma_2$  is bound to be a monoid homomorphism.

We finally define an implementation over a family  $\mathcal{R}$  of replicas.

**Definition 4.7** (RLTS implementation). *An RLTS implementation over  $\mathcal{R}$  is an  $\mathcal{R}$ -indexed family of isomorphic RLTSs.*

#### 4.2. Bisimulation for RLTSs

We now move to provide a notion of bisimulation for RLTSs, which is akin to the one for higher-order calculi: beside the symmetric simulation game, it exploits the notion of *barbs*, that is, predicates on states. Introduced in [19] and widely used for the definition of equivalences in process calculus, barbs intuitively express the satisfiability of an observable property by a state. More specifically, we will write  $\sigma \downarrow b$  to say that a state  $\sigma$  verifies a property  $b$ .

**Definition 4.8** (Action and implementation bisimulation). *Let  $\mathcal{Q}$  be an RLTS over a set of actions  $\mathcal{M}$ ,  $\mathcal{B}$  a set of barbs, and  $\downarrow \subseteq \Sigma \times \mathcal{B}$  a relation. Then an action bisimulation is a symmetric relation  $\mathcal{R} \subseteq \mathcal{Q} \times \mathcal{Q}$  between states in  $\mathcal{Q}$  such that if  $(\sigma_1, \sigma_2) \in \mathcal{R}$  then*

1 if  $\sigma_1 \downarrow b$  then  $\sigma_2 \downarrow b$

2 if  $\sigma_1 \xrightarrow{m} \sigma'_1$  then  $\exists \sigma'_2. \sigma_2 \xrightarrow{m} \sigma'_2$

*It is an implementation bisimulation if additionally*

3 if  $\sigma_1 \xrightarrow{\sigma'_1} \sigma''_1$  then  $\exists \sigma'_2, \sigma''_2$  such that  $\sigma_2 \xrightarrow{\sigma'_2} \sigma''_2$  and  $\{(\sigma'_1, \sigma'_2), (\sigma''_1, \sigma''_2)\} \subseteq \mathcal{R}$ .

We write  $\sim_{\mathcal{Q}}^{\mathcal{B}}$  for the largest action bisimulation, and  $\approx_{\mathcal{Q}}^{\mathcal{B}}$  for the largest implementation bisimulation.

Clearly,  $\approx_{\mathcal{Q}}^{\mathcal{B}} \subseteq \sim_{\mathcal{Q}}^{\mathcal{B}}$ , that is, if  $\sigma_1 \approx_{\mathcal{Q}}^{\mathcal{B}} \sigma_2$  then  $\sigma_1 \sim_{\mathcal{Q}}^{\mathcal{B}} \sigma_2$  for any two states  $\sigma_1, \sigma_2$ . However, under certain conditions the two notions coincide.

**Proposition 4.9** (Actions as implementations). *Let  $\mathcal{Q}$  be a state-deterministic RLTS and  $\mathcal{B}$  a set of barbs. Then  $\sim_{\mathcal{Q}}^{\mathcal{B}}$  and  $\approx_{\mathcal{Q}}^{\mathcal{B}}$  coincide if and only if*

$$\forall \sigma_1, \sigma_2. \langle \sigma_1, \sigma_2 \rangle \in \sim_{\mathcal{Q}}^{\mathcal{B}} \implies \forall \sigma'_1. \exists \sigma'_2. \{ \langle \sigma'_1, \sigma'_2 \rangle, \langle \sigma_1 \otimes \sigma'_1, \sigma_2 \otimes \sigma'_2 \rangle \} \subseteq \sim_{\mathcal{Q}}^{\mathcal{B}}$$

*Proof.* For the left-to-right part, assume  $\sim_{\mathcal{Q}}^{\mathcal{B}} = \approx_{\mathcal{Q}}^{\mathcal{B}}$  and consider  $\langle \sigma_1, \sigma_2 \rangle \in \sim_{\mathcal{Q}}^{\mathcal{B}}$ . By state-determinism, for all  $\sigma'_1$  there is a transition

$$\sigma_1 \xrightarrow{\sigma'_1} \sigma_1 \otimes \sigma'_1.$$

By  $\sim_{\mathcal{Q}}^{\mathcal{B}} = \approx_{\mathcal{Q}}^{\mathcal{B}}$  and the definition of implementation bisimulation, this transition

can be simulated by  $\sigma_2 \xrightarrow{\sigma'_2} \sigma_1 \otimes \sigma'_2$ , for some  $\sigma'_2$ , with

$$\{ \langle \sigma'_1, \sigma'_2 \rangle, \langle \sigma_1 \otimes \sigma'_1, \sigma_2 \otimes \sigma'_2 \rangle \} \subseteq \sim_{\mathcal{Q}}^{\mathcal{B}}$$

as required.

Vice versa, suppose the implication holds; we need to prove  $\sim_{\mathcal{Q}}^{\mathcal{B}} \subseteq \approx_{\mathcal{Q}}^{\mathcal{B}}$ . Consider  $\langle \sigma_1, \sigma_2 \rangle \in \sim_{\mathcal{Q}}^{\mathcal{B}}$ , and suppose  $\sigma_1 \xrightarrow{\sigma'_1} \sigma_1 \otimes \sigma'_1$ . Then by the implication we get the implementation bisimulation condition.  $\square$

Note that the property required in Proposition 4.9 for the correspondence between the two equivalences is weaker than the closure with respect to the monoidal operator, which can be stated as

$$\forall \sigma_1, \sigma_2. \langle \sigma_1, \sigma_2 \rangle \in \sim_{\mathcal{Q}}^{\mathcal{B}} \implies \forall \sigma. \langle \sigma_1 \otimes \sigma, \sigma_2 \otimes \sigma \rangle \in \sim_{\mathcal{Q}}^{\mathcal{B}}$$

### 4.3. Some key examples

We now present instantiations corresponding to well-known implementations of the RDTs specified in Section 2.2. All the associated RLTSs are state-deterministic, hence we will drop the adjectives. All the RLTSs will be proved to be monoidal, which validates the generality of our model.

**Example 4.10** (Implementation of the *GCounter*). *The RLTS corresponding to a replica  $r$  that implements the RDT *GCounter* specified in Example 2.2 is defined as  $\mathcal{Q}^r = \langle \Sigma, \oplus, 1, \rightsquigarrow^r \rangle$  where*

- $\Sigma = \mathbb{N}^{\mathcal{R}}$ , i.e., states are mappings  $v$  where  $v(r)$  is the number of increments performed over  $r$ .
- The operator  $\oplus$  is defined as the point-wise maximum of  $v$  and  $v'$ ,  $v \oplus v' = \max\{v, v'\}$  where  $\max\{v, v'\}(s) = \max\{v(s), v'(s)\}$  for all  $s \in \mathcal{R}$ .
- $1$  is the constant map to 0.

- The set of mutators is the singleton  $\mathcal{M} = \{\langle \mathbf{inc}, \mathbf{ok} \rangle\}$ . Then,  $\rightsquigarrow^r$  is defined as the least relation satisfying rule  $(\text{rcv})$  and the following one

$$\frac{}{v \overset{\langle \mathbf{inc}, \mathbf{ok} \rangle}{\rightsquigarrow^r} v[\mathbf{r} \mapsto v(\mathbf{r}) + 1]} (\text{inc})$$

which describes the effects of performing an increment over the replica  $\mathbf{r}$ :  
The entry associated with the replica  $\mathbf{r}$  in the map  $v$  is incremented.

It straightforwardly follows from the definitions above that  $\langle \Sigma, \oplus, 1 \rangle$  is a commutative monoid. We now inspect the conditions in Definition 4.3

- 1 if  $v_i \overset{\langle \mathbf{inc}, \mathbf{ok} \rangle}{\rightsquigarrow^r} v'_i$  then  $v'_i = v_i[\mathbf{r} \mapsto v_i(\mathbf{r}) + 1]$  for  $i = 1, 2$ . Then,

$$\begin{aligned} v'_1 \oplus v'_2 &= (v_1 \oplus v_2)[\mathbf{r} \mapsto \max\{v_1(\mathbf{r}) + 1, v_2(\mathbf{r}) + 1\}] \\ &= (v_1 \oplus v_2)[\mathbf{r} \mapsto \max\{v_1(\mathbf{r}), v_2(\mathbf{r})\} + 1] \\ &= (v_1 \oplus v_2)[\mathbf{r} \mapsto (v_1 \oplus v_2)(\mathbf{r}) + 1] \end{aligned}$$

Hence,  $v_1 \oplus v_2 \overset{\langle \mathbf{inc}, \mathbf{ok} \rangle}{\rightsquigarrow^r} v'_1 \oplus v'_2$ .

- 2 if  $v_1 \oplus v_2 \overset{\langle \mathbf{inc}, \mathbf{ok} \rangle}{\rightsquigarrow^r} v$  then  $v = (v_1 \oplus v_2)[\mathbf{r} \mapsto (v_1 \oplus v_2)(\mathbf{r}) + 1]$ . Moreover,  $v_i \overset{\langle \mathbf{inc}, \mathbf{ok} \rangle}{\rightsquigarrow^r} v_i[\mathbf{r} \mapsto v_i(\mathbf{r}) + 1]$  for  $i = 1, 2$ . Hence, the condition follows by taking  $v'_i = v_i[\mathbf{r} \mapsto v_i(\mathbf{r}) + 1]$ .

Consequently,  $\mathcal{Q}^r = \langle \Sigma, \oplus, 1, \rightsquigarrow^r \rangle$  is a monoidal RLTS.

The set of barbs is  $\mathcal{B} = \{\langle \mathbf{rd}, k \rangle \mid k \in \mathbb{N}\}$ , and  $\downarrow$  is given by

$$v \downarrow \langle \mathbf{rd}, k \rangle \iff k = \sum_{\mathbf{s} \in \mathcal{R}} v(\mathbf{s})$$

Note that  $v \downarrow \langle \mathbf{rd}, k \rangle$  holds only if the return value  $k$  in the barb corresponds to the value stored in  $v$  (i.e.,  $k = \sum_{\mathbf{s} \in \mathcal{R}} v(\mathbf{s})$ ).

**Example 4.11** (*GCounter bisimulation*). Consider a GCounter implementation with two replicas, so that states can be simply modelled as pairs of natural numbers  $\langle n, m \rangle$ . Also, let us consider the RLTS for the first replica, so that

$$\langle n, m \rangle \overset{\langle \mathbf{inc}, \mathbf{ok} \rangle}{\rightsquigarrow^1} \langle n + 1, m \rangle.$$

Clearly we have that  $\langle n, m \rangle \sim_{\mathcal{Q}}^{\mathcal{B}} \langle x, y \rangle$  if and only if  $n + m = x + y$ , since  $\langle n, m \rangle \downarrow \langle \mathbf{rd}, k \rangle \iff n + m = k$ . However, this is not an implementation bisimulation. In fact, let us consider the states  $\langle 9, 2 \rangle$  and  $\langle 6, 5 \rangle$ : If the former receives the state  $\langle 0, 3 \rangle$ , there is no state that can be suitably received from  $\langle 6, 5 \rangle$ .

**Example 4.12** (*Implementation of the PN-Counter*). The RLTS for the implementation of the replica  $\mathbf{r}$  of the RDT PNCounter specified in Example 2.3 can be defined as  $\mathcal{Q}^r = \langle \Sigma, \oplus, 1, \rightsquigarrow^r \rangle$  where

- $\Sigma = \mathbb{N}^{\mathcal{R}} \times \mathbb{N}^{\mathcal{R}}$ , i.e., states are pairs  $(p, n)$ , where the mappings  $p$  and  $n$  respectively record the number of increments and decrements performed over each replica.
- The combinator  $\oplus$  is defined as  $(p, n) \oplus (p', n') = (\max\{p, p'\}, \max\{n, n'\})$ .
- The identity  $1$  is the pair of constant mappings to  $0$ .
- The set of mutators is  $\mathcal{M} = \{\langle \text{inc}, \text{ok} \rangle, \langle \text{dec}, \text{ok} \rangle\}$ . Then,  $\rightsquigarrow^r$  is defined as the least relation satisfying rule  $(\text{rcv})$  and the following ones

$$\frac{}{(p, n) \rightsquigarrow^{\langle \text{inc}, \text{ok} \rangle} (p[r \mapsto p(r) + 1], n)} \text{(INC)}$$

$$\frac{}{(p, n) \rightsquigarrow^{\langle \text{dec}, \text{ok} \rangle} (p, n[r \mapsto p(r) + 1])} \text{(DEC)}$$

As for the GCounter, we can conclude that  $\langle \Sigma, \oplus, 1 \rangle$  is a monoidal RLTS. Also, the set of barbs is defined as  $\mathcal{B} = \{\langle \text{rd}, k \rangle \mid k \in \mathbb{N}\}$ , and  $\downarrow$  is given by

$$(p, n) \downarrow \langle \text{rd}, k \rangle \iff k = \sum_{s \in \mathcal{R}} p(s) - n(s).$$

**Example 4.13** (Implementation of the Multi-value Register). We present an implementation for the Multi-value Register specified in Example 2.4 based on the optimised implementation proposed in [24]. The implementation associates a single scalar (logical) clock and a vector clock to each written value to determine if two writes are concurrent or causally ordered. A scalar clock is a natural number and a version vector is just a mapping from replicas to natural numbers. The state of a replica consists of a pair  $(S, v)$  where  $S \in 2^{\mathcal{R} \times \mathbb{N} \times \mathcal{V}}$  is a set of tagged values and  $v \in \mathbb{N}^{\mathcal{R}}$  a version vector. An element  $(r, c, k) \in S$  indicates that the value  $k$  has been written over  $r$  at the logical time  $c$ . It is assumed that  $(\mathcal{V}, \leq)$  is a partial order. Then, the order among tagged elements is given by

$$(r, c, k) \leq (r', c', k') \iff (r = r' \wedge (c, k) \leq_{\text{lex}} (c', k')) \vee (r \neq r' \wedge k \leq k')$$

where  $\leq_{\text{lex}}$  stands for lexicographical order. Consequently, the maximal concurrent written values in a set  $S \in 2^{\mathcal{R} \times \mathbb{N} \times \mathcal{V}}$  are the maximal elements  $\|S\|$  defined as  $\{u \in S \mid \nexists v \in S, u \leq v\}$ .

Then, for definition of  $\mathcal{Q}^r = \langle \Sigma, \oplus, 1, \rightsquigarrow^r \rangle$  we take

- $\Sigma = \{(\|S\|, v) \mid S \in 2^{\mathcal{R} \times \mathbb{N} \times \mathcal{V}}, v \in \mathbb{N}^{\mathcal{R}}, \forall (r, c, k) \in S. 1 \leq c \leq v(r)\}$ , i.e., each state consists of the set of maximal concurrent written values tagged with a logical clock consistent with the version vector  $v$ .
- The state combinator  $\oplus$  is defined as follows

$$(S, v) \oplus (S', v') = (\|(S \cap S') \cup S_{|v'} \cup S'_{|v}\|, \max\{v, v'\})$$

where  $S_{|v} = \{(r, c, k) \mid (r, c, k) \in S \wedge c > v(r)\}$

- the identity  $1$  is the pair  $(\emptyset, \perp)$  where  $\perp$  is the constant mapping to  $0$ .
- The set of mutators is  $\mathcal{M} = \{\langle \mathbf{wr}(k), ok \rangle \mid k \in \mathcal{V}\}$ . Then,  $\rightsquigarrow^r$  is defined as the least relation satisfying rule  $(\text{RCV})$  and the following one

$$\frac{}{(S, v) \rightsquigarrow^{\langle \mathbf{wr}(k), ok \rangle} \{ \{ (r, v(r) + 1, k) \}, v[r \mapsto v(r) + 1] \}} \quad (\text{WR})$$

It is immediate from its definition that  $\oplus$  is commutative. We show that  $1 = (\emptyset, \perp)$  is the identity as follows.

$$\begin{aligned} (S, v) \oplus (\emptyset, \perp) &= (\|(S \cap \emptyset) \cup S_{|\perp} \cup \emptyset_{|v}\|, \max\{v, \perp\}) && \text{by the def. of } \oplus \\ &= (\|S_{|\perp} \cup \emptyset_{|v}\|, \max\{v, \perp\}) && \text{by the def. of } \cup \text{ and } \cap \\ &= (\|S_{|\perp}\|, \max\{v, \perp\}) && \text{because } \emptyset_{|v} = \emptyset \\ &= (\|S_{|\perp}\|, v) && \text{because } \perp \text{ is the identity} \end{aligned}$$

It remains to show that  $\|S_{|\perp}\| = S$ . Since  $(S, v) \in \Sigma$ , for all  $\forall(r, c, k) \in S.1 \leq c \leq v(r)$ . Therefore,  $\forall(r, c, k) \in S, c > 0 = \perp(r)$  holds. Hence,  $S_{|\perp} = S$ . Moreover,  $(S, v) \in \Sigma$  implies that there exists  $S'$  such that  $S = \|S'\|$ . Therefore,  $\|S\| = S$  because  $\|-\|$  is idempotent. Hence, we can conclude that  $\mathcal{Q}^r = \langle \Sigma, \oplus, 1, \rightsquigarrow^r \rangle$  is a commutative monoid. We now inspect the conditions in Definition 4.3

- 1 if  $(S_i, v_i) \rightsquigarrow^{\langle \mathbf{wr}(k), ok \rangle} (S'_i, v'_i)$  then we have  $S'_i = \{(k, r, v(r) + 1)\}$  and  $v'_i = v_i[r \mapsto v(r) + 1]$  for  $i = 1, 2$ . Then

$$(S'_1, v'_1) \oplus (S'_2, v'_2) = (\|(S'_1 \cap S'_2) \cup S'_{1|v'_2} \cup S'_{2|v'_1}\|, \max\{v'_1, v'_2\})$$

Firstly, we note that  $\max\{v'_1, v'_2\} = \max\{v_1, v_2\}[r \mapsto \max\{v_1, v_2\}(r) + 1]$ . For  $\|(S'_1 \cap S'_2) \cup S'_{1|v'_2} \cup S'_{2|v'_1}\|$ , we have three cases

- $v_1(r) = v_2(r)$ . Then  $S'_1 = S'_2 = \{(r, v_1(r) + 1, k)\} = S'_1 \cap S'_2$ . Moreover  $S'_{1|v'_2} = S'_{2|v'_1} = \emptyset$  because  $S'_i = (r, v_1(r) + 1, k)$  and clearly  $v_1(r) + 1 \neq v'_i(r) = v_i(r) + 1$ . Hence

$$(S_1, v'_1) \oplus (S_2, v'_2) = (\{(r, v_1(r) + 1, k)\}, \max\{v'_1, v'_2\})$$

Therefore  $(S_1, v_1) \oplus (S_2, v_2) \rightsquigarrow^{\langle \mathbf{wr}(k), ok \rangle} (S'_1, v'_1) \oplus (S'_2, v'_2)$ .

- $v_1(r) > v_2(r)$ . Then  $S'_i = \{(r, v_i(r) + 1, k)\}$  and  $S'_1 \cap S'_2 = \emptyset$ . Moreover  $S'_{1|v'_2} = \{(r, v_i(r) + 1, k)\}$  and  $S'_{2|v'_1} = \emptyset$ . Hence

$$(S_1, v'_1) \oplus (S_2, v'_2) = (\{(r, v_1(r) + 1, k)\}, \max\{v'_1, v'_2\})$$

Therefore  $(S_1, v_1) \oplus (S_2, v_2) \rightsquigarrow^{\langle \mathbf{wr}(k), ok \rangle} (S'_1, v'_1) \oplus (S'_2, v'_2)$ .

- $v_2(r) > v_1(r)$ . Follows as in the previous case.

2 Follows along the lines of the previous case.

The set of barbs is defined as  $\mathcal{B} = \{\langle \text{rd}, S \rangle_r \mid S \in 2^{\mathcal{V}}\}$ , and  $\downarrow$  is given by

$$(S, v) \downarrow \langle \text{rd}, S' \rangle_r \iff S' = \{k \mid (s, c, k) \in S\}$$

**Example 4.14** (Implementation of the Last-write wins Register). We now describe the implementation based on timestamps proposed in [21] for the Last-write wins Register introduced in Example 2.5. Let  $(\mathbb{T}, <)$  be the totally-ordered set of timestamps with minimum  $t_0$ . Then, the corresponding RLTS for a replica  $r$  is defined as  $\mathcal{Q}^r = \langle \Sigma, \oplus, 1, \rightsquigarrow^r \rangle$  where

- $\Sigma = \mathbb{T} \times (\mathcal{V} \cup \{\perp\})$ , i.e., the state  $(t, k)$  of a replica contains the current value  $k$  of the register and its associated timestamp. We assume  $(\mathcal{V}, \leq)$  to be a total order and establish that  $\perp \leq k$  for all  $k \in \mathcal{V}$ ; moreover, we consider  $\mathbb{T} \times (\mathcal{V} \cup \{\perp\})$  lexicographically ordered, i.e.,  $(t, k) \leq (t', k')$  when either  $t < t'$  or  $t = t'$  and  $k \leq k'$ .
- The state combinator  $\oplus$  is defined as follows

$$(t, k) \oplus (t', k') = \max\{(t, k), (t', k')\}$$

- the identity  $1$  is the pair  $(t_0, \perp)$ .
- The set of mutators is  $\mathcal{M} = \{\langle \text{wr}(k), \text{ok} \rangle \mid k \in \mathbb{N}\}$ . Then,  $\rightsquigarrow^r$  is defined as the least relation satisfying rule  $(\text{rcv})$  and the following one

$$\frac{t' \in \mathbb{T}}{(t, k) \rightsquigarrow^r \langle \text{wr}(k'), \text{ok} \rangle \max\{(t, k), (t', k')\}} \quad (\text{wr})$$

Note that the transition system is non-deterministic, because the target of the arrow depends on the chosen  $t'$ .

It immediately follows from the definitions that  $\mathcal{Q}^r = \langle \Sigma, \oplus, 1, \rightsquigarrow^r \rangle$  is a commutative monoid. We now inspect the conditions in Definition 4.3

- 1 if  $(t_i, k_i) \rightsquigarrow^r \langle \text{wr}(k), \text{ok} \rangle (t'_i, k'_i)$  then either  $(t'_i, k'_i) = (t_i, k_i)$  or  $(t_i, k_i) < (t'_i, k'_i)$  and  $k'_i = k$ . Then, define  $(t', k') = (t'_1, k'_1) \oplus (t'_2, k'_2) = \max\{(t'_1, k'_1) \oplus (t'_2, k'_2)\}$ . Since  $(t_i, k_i) \leq (t', k')$  for  $i = 1, 2$ , we conclude that  $(t_1, k_1) \oplus (t_1, k_1) \leq (t', k')$  and therefore  $(t_1, k_1) \oplus (t_1, k_1) \rightsquigarrow^r \langle \text{wr}(k), \text{ok} \rangle (t', k')$ .

2 Follows along the lines of the previous case.

The set of barbs is defined as  $\mathcal{B} = \{\langle \text{rd}, k \rangle \mid k \in \mathbb{N}\}$ , and  $\downarrow$  is given by

$$(t, k) \downarrow \langle \text{rd}, k' \rangle \iff k' = k$$



**Example 4.15** (Implementation of the *OR-Set*). The behaviour of a single replica according to the implementation proposed in [2] is given by the RLTS  $\mathcal{Q}^r = \langle \Sigma, \oplus, 1, \rightsquigarrow^r \rangle$  defined as follows

- $\Sigma = \{(v, w) \mid v \in \mathbb{N}^{\mathcal{R}}, w \in \mathbb{N}^{\mathcal{V} \times \mathcal{R}}, w(k, r) \leq v(r)\}$ , i.e., states are pairs  $(v, w)$ , where  $v$  is a version vector, and  $w$  maps each pair  $(k, r)$  to the most recent version of  $r$  in which  $k$  has been added to  $r$ . If  $w(k, r) = 0$  then  $k$  does not belong to the set according to  $r$ .
- The combination of  $(v, w)$  and  $(v', w')$  handles conflicting information in the mappings  $w$  and  $w'$  for the same element. There is a conflict between  $w$  and  $w'$  for  $(k, s)$  when one mapping indicates that  $k$  is present in the replica  $s$  and the other does not, i.e.,  $w(k, s) > 0$  and  $w'(k, s) = 0$  or vice versa. Such conflicts are resolved by using the version vectors  $v$  and  $v'$ , as formalised by the operation  $+$ :  $(\mathbb{N}^{\mathcal{R}} \times \mathbb{N}^{\mathcal{V} \times \mathcal{R}}) \times (\mathbb{N}^{\mathcal{R}} \times \mathbb{N}^{\mathcal{V} \times \mathcal{R}}) \rightarrow \mathbb{N}^{\mathcal{V} \times \mathcal{R}}$ , defined such that

$$(v, w) + (v', w')(k, s) = \begin{cases} w(k, s) & v(s) > v'(s) \\ w'(k, s) & v(s) < v'(s) \\ \min\{w(k, s), w'(k, s)\} & \text{otherwise} \end{cases}$$

Then, the state combinator  $\oplus$  is defined as

$$(v, w) \oplus (v', w') = (\max\{v, v'\}, (v, w) + (v', w'))$$

- The identity  $1$  is the pair  $(\perp, \perp)$  where  $\perp$  are suitable constant maps to  $0$ .
- The set of mutators is

$$\mathcal{M} = \{\langle \mathbf{add}(k), ok \rangle \mid k \in \mathcal{V}\} \cup \{\langle \mathbf{rem}(k), ok \rangle \mid k \in \mathcal{V}\}$$

Then,  $\rightsquigarrow^r$  is defined as the least relation satisfying rule  $(\text{RCV})$  and the following ones

$$\frac{}{(v, w) \rightsquigarrow^r (v[\mathbf{r} \mapsto v(\mathbf{r}) + 1], w[(k, \mathbf{r}) \mapsto v(\mathbf{r}) + 1])} \text{(ADD)}$$

$$\frac{}{(v, w) \rightsquigarrow^r (v, w[\forall s \in \mathcal{R}. (k, s) \mapsto 0])} \text{(REM)}$$

By rule (ADD), if  $k$  is added to  $r$ , then  $r$  changes its local state by (i) creating a new version, i.e., the entry  $v(r)$  is updated to  $v(r) + 1$ , and (ii) recording that  $k$  has been added in the newest version of  $r$ , i.e.  $w(k, r)$  is updated to  $v(r) + 1$ . By rule (REM), if  $k$  is removed from the set, then all entries in the second mapping associated with  $k$  are set to  $0$  to record the elimination.

First, note that  $+$  is a commutative operator. Hence,  $\mathcal{Q}^r = \langle \Sigma, \oplus, 1, \rightsquigarrow^r \rangle$  is a commutative monoid. We now inspect the conditions in Definition 4.3

1 if  $(v_i, w_i) \xrightarrow{\langle \text{add}(k), ok \rangle}^r (v'_i, w'_i)$  then we have  $v'_i = v_i[r \mapsto v_i(r) + 1]$  and  $w'_i = w_i[(k, r) \mapsto v_i(r) + 1]$  for  $i = 1, 2$ .

Then, define  $(v', w') = (v'_1, w'_1) \oplus (v'_2, w'_2) = (\max\{v'_1, v'_2\}, (v'_1, w'_1) + (v'_2, w'_2))$ . Note that  $\max\{v'_1, v'_2\} = \max\{v_1, v_2\}[r \mapsto \max\{v_1, v_2\}(r) + 1]$ . Also, for all  $s \in \mathcal{R}$  and  $k' \in \mathbb{N}$ , if  $s \neq r$  and  $k' \neq k$  we have

$$((v'_1, w'_1) + (v'_2, w'_2))(k', s) = ((v_1, w_1) + (v_2, w_2))(k', s)$$

Additionally

$$((v'_1, w'_1) + (v'_2, w'_2))(k, r) = \max\{v_1(r), v_2(r)\} + 1$$

Hence

$$(v'_1, w'_1) + (v'_2, w'_2) = ((v_1, w_1) + (v_2, w_2))[(k, r) \mapsto \max\{v_1(r), v_2(r)\} + 1]$$

Consequently,  $(v_1, w_1) \oplus (v_2, w_2) \xrightarrow{\langle \text{add}(k), ok \rangle}^r (v', w')$ .

The case for the mutator  $\langle \text{rem}(k), ok \rangle$  follows analogously.

2 Follows along the lines of the previous case.

The set of barbs is  $\mathcal{B} = \{\langle \text{lookup}, S \rangle \mid S \in 2^{\mathcal{V}}\}$ , and  $\downarrow$  is given by

$$(v, w) \downarrow \langle \text{lookup}, S \rangle \iff S = \{k \mid \exists s \in \mathcal{R}. w(k, s) > 0\}$$

When performing a lookup, the replica  $r$  returns the set  $S$  of all  $k$  that, according to  $w$ , have been added to some replica  $s$  (i.e.,  $w(k, s) > 0$  holds)

**Example 4.16** (Implementation of the 2P-Set). The behaviour of a single replica according to the implementation proposed in [21] is given by the RLTS  $\mathcal{Q}^r = \langle \Sigma, \oplus, 1, \rightsquigarrow^r \rangle$  defined as follows

- $\Sigma = 2^{\mathcal{V}} \times 2^{\mathcal{V}}$ , i.e., states are pairs  $(A, R)$ , where  $A$  and  $R$  are respectively the sets of added and removed elements.
- The combination of  $(A, R)$  and  $(A', R')$  simply consists in the point-wise union of sets, i.e.,  $(A, R) \oplus (A', R') = (A \cup A', R \cup R')$ .
- The identity 1 is the pair  $(\emptyset, \emptyset)$ .
- The set of mutators is

$$\mathcal{M} = \{\langle \text{add}(k), ok \rangle \mid k \in \mathcal{V}\} \cup \{\langle \text{rem}(k), ok \rangle \mid k \in \mathcal{V}\}$$

Then,  $\rightsquigarrow^r$  is defined as the least relation satisfying rule (RCV) and the following ones

$$\frac{}{(A, R) \rightsquigarrow^r (A \cup \{k\}, E)} \text{(ADD)}$$

$$\frac{}{(A, R) \rightsquigarrow^r (A, E \cup \{k\})} \text{(REM)}$$

According to (ADD), an addition simply extends the set of added elements with  $k$ ; analogously, a removal extends the set of removed events (ADD).

The implementation is analogous to the PNCOUNTER; consequently, we can show that  $\mathcal{Q}$  is a monoidal RLTS.

The set of barbs is defined as  $\mathcal{B} = \{\langle \text{lookup}, S \rangle \mid S \in 2^{\mathcal{V}}\}$ , and  $\downarrow$  is given by

$$(A, R) \downarrow \langle \text{lookup}, S \rangle \iff S = A \setminus R$$

When performing a lookup, the replica  $r$  returns the elements that have been added but not removed.

## 5. From specifications to LTS

We can exploit the structure of coherent functors to recover an operational interpretation of specifications. In the following, we consider a coherent functor  $\mathbb{F} : \mathbf{PIDag}(\mathcal{L}) \rightarrow \mathbf{SPath}(\mathcal{L})$ . We will recover the corresponding LTS via a standard categorical construction, namely its *category of elements*  $\mathbf{E}(\mathbb{F})$  (see e.g. [18, Pag. 41]) as follows. Given a pr-morphism  $\mathbf{f} : \mathbf{G} \rightarrow \mathbf{G}_1$  and  $\mathbf{P} \in \mathbb{F}(\mathbf{G})$ , we denote by  $\mathbb{F}(\mathbf{f})(\mathbf{P})$  the set of paths in  $\mathbb{F}(\mathbf{G}_1)$  that are in the “image” of  $\mathbf{P}$  via  $\mathbb{F}(\mathbf{f})$ , formally specified as  $\mathbb{F}(\mathbf{G}_1) \cap \text{sat}(\mathbf{P}, \mathbf{f})$ .

**Definition 5.1** (Category of elements). *The category of elements  $\mathbf{E}(\mathbb{F})$  of  $\mathbb{F}$  is obtained as follows*

- objects are pairs  $\langle \mathbf{G}, \mathbf{P} \rangle$ , such that  $\mathbf{G} \in \mathbf{PIDag}(\mathcal{L})$  and  $\mathbf{P} \in \mathbb{F}(\mathbf{G})$ ;
- arrows  $\mathbf{f} : \langle \mathbf{G}, \mathbf{P} \rangle \rightarrow \langle \mathbf{G}_1, \mathbf{P}_1 \rangle$  are pr-morphisms  $\mathbf{f} : \mathbf{G} \rightarrow \mathbf{G}_1$  such that  $\mathbf{P}_1 \in \mathbb{F}(\mathbf{f})(\mathbf{P})$ .

Intuitively, arrows in  $\mathbf{E}(\mathbb{F})$  stand for the possible ways a path in  $\mathbb{F}(\mathbf{G})$  can evolve according to  $\mathbb{F}(\mathbf{f})$ . The category  $\mathbf{E}(\mathbb{F})$  is clearly an LTS, since each category is so. We note that our way of distilling an LTS is similar to how one obtains an LTS from a relation presheaf [23, Definition 4.1].

**Example 5.2.** *Consider the functor  $\mathbb{M}(\mathcal{S}_{mvR})$  induced by the coherent specification  $\mathcal{S}_{mvR}$  in Example 2.4. An object  $\langle \mathbf{G}, \mathbf{P} \rangle$  of the category of elements  $\mathbf{E}(\mathbb{M}(\mathcal{S}_{mvR}))$  represents a state of the RDT where the events in the visibility graph  $\mathbf{G}$  are arbitrated according to  $\mathbf{P} \in \mathcal{S}_{mvR}(\mathbf{G})$ . An arrow  $\mathbf{f} : \langle \mathbf{G}, \mathbf{P} \rangle \rightarrow \langle \mathbf{G}_1, \mathbf{P}_1 \rangle$  in  $\mathbf{E}(\mathbb{M}(\mathcal{S}_{mvR}))$  describes a computation where the visibility  $\mathbf{G}$  is extended to  $\mathbf{G}_1$  and the arbitration  $\mathbf{P}$  to  $\mathbf{P}_1$ . For instance, take the graphs  $\mathbf{G}_2$  and  $\mathbf{G}_3$  in Fig. 4c*

and Fig. 4d, and the unique pr-morphism  $\mathbf{f}: \mathbf{G}_2 \rightarrow \mathbf{G}_3$ . If  $\mathbf{P}_2$  is a total order of the events in  $\mathbf{G}_2$ , then there is a ps-morphism  $\mathbf{f}: \{\mathbf{P}_2\} \rightarrow \mathbf{sat}(\mathbf{P}_2, \mathbf{f})$ . Moreover,  $\mathcal{S}_{mvR}(\mathbf{G}_3) \cap \mathbf{sat}(\mathbf{P}_2, \mathbf{f}) = \mathbf{sat}(\mathbf{P}_2, \mathbf{f})$  because  $\mathcal{S}_{mvR}$  imposes no constraint on the admissible arbitrations of a consistent visibility. Therefore, there is a morphism  $\mathbf{f}: \langle \mathbf{G}_2, \mathbf{P}_2 \rangle \rightarrow \langle \mathbf{G}_3, \mathbf{P}_3 \rangle$  for any  $\mathbf{P}_3 \in \mathbf{sat}(\mathbf{f}, \mathbf{P}_2)$  in  $\mathbf{E}(\mathbb{M}(\mathcal{S}_{mvR}))$ .

A pr-morphism may not induce an arrow in the category of elements, as  $\mathbf{f}: \mathbf{G}_2 \rightarrow \mathbf{G}_4$  with  $\mathbf{G}_2$  from Fig. 4c and  $\mathbf{G}_4$  its root extension along  $\langle \mathbf{rd}, \{1\} \rangle$ . Indeed,  $\mathcal{S}_{mvR}(\mathbf{G}_4) = \emptyset$ , and hence  $\mathbb{M}(\mathcal{S}_{mvR})(\mathbf{f})(\mathbf{P}) = \emptyset$  for any  $\mathbf{P} \in \mathcal{S}_{mvR}(\mathbf{G}_2)$ .

An analogous situation occurs when the specification restricts the allowed arbitrations, as  $\mathcal{S}_{lwwR}$  in Example 2.5. Consider the root extension  $\mathbf{f}: \mathbf{G} \rightarrow \mathbf{G}_1$  along  $\langle \mathbf{rd}, \{2\} \rangle$ , with  $\mathbf{G}_1$  as in Fig. 4a and  $\mathbf{G}$  is  $\mathbf{G}_1$  without the event  $\langle \mathbf{rd}, \{2\} \rangle$ . Then we have  $\mathcal{S}_{lwwR}(\mathbf{G}_1) = \{\mathbf{P}_1\}$ , and  $\mathcal{S}_{lwwR}(\mathbf{G})$  contains two paths:  $\mathbf{P}$ , which keeps the order of writes as in  $\mathbf{P}_1$ , and  $\mathbf{P}'$ , which inverts it. The latter cannot be extended to any path in  $\mathcal{S}_{lwwR}(\mathbf{G}_1)$ , as writes are in the wrong order. In fact, there is no  $\mathbf{f}: \langle \mathbf{G}, \mathbf{P}' \rangle \rightarrow \langle \mathbf{G}_1, \mathbf{P}_1 \rangle$  in  $\mathbf{E}(\mathbb{M}(\mathcal{S}_{lwwR}))$ . Contrastingly, we have that  $\mathbf{f}: \langle \mathbf{G}, \mathbf{P}' \rangle \rightarrow \langle \mathbf{G}_1, \mathbf{P}'' \rangle$  is an arrow of  $\mathbf{E}(\mathbb{M}(\mathcal{S}_{mvR}))$  for any  $\mathbf{P}'' \in \mathbf{sat}(\mathbf{f}, \mathbf{P}')$ , as the order of writes is irrelevant for  $\mathcal{S}_{mvR}$ .

### 5.1. One- and Multi-replica LTSs

We now recover one-replica LTSs in our setting (see Definition 2.10), by observing that the way visibility is augmented along a transition can be formalised as a root  $\ell$ -extension. In fact, the one-replica LTS precisely corresponds to a sub-category of  $\mathbf{E}(\mathbb{F})$  consisting only of such extensions.

**Proposition 5.3.** *Let  $\mathbf{E}_o(\mathbb{F})$  be  $\mathbf{E}(\mathbb{F})$  restricted to root  $\ell$ -extensions, for all  $\ell \in \mathcal{L}$ . Then the one-replica LTS coincides with the LTS for  $\mathbf{E}_o(\mathbb{F})$ .*

This is easily seen: Each root  $\ell$ -extension corresponds uniquely to a  $\ell$ -labelled one-replica transition, and between any two graphs there is at most one root extension, so that also the label is implicitly recovered.

We now go a step further, and recover the multi-replica LTS of Definition 2.12 by exploiting the structure of coherent functors. Recall that, given replica states  $\langle \mathbf{G}_i, \mathbf{P}_i \rangle$ ,  $i = 1, 2$ , we assume that  $\mathbf{G}_1$  and  $\mathbf{G}_2$  are compatible when forming the composite state. In our categorical setting, compatibility is formalised as the existence of a span  $\mathbf{f}_i: \mathbf{G} \rightarrow \mathbf{G}_i$  of mono pr-morphisms such that  $\mathcal{E}_{\mathbf{G}} = \mathcal{E}_{\mathbf{G}_1} \cap \mathcal{E}_{\mathbf{G}_2}$  (thus, shared nodes have the same labels and predecessors). Notably, the union  $\mathbf{G}_1 \sqcup \mathbf{G}_2$  and the obvious morphisms are the pushout for this span.

Before our main characterisation result, we need two technical lemmas. The first says that certain pushouts in  $\mathbf{SPath}(\mathcal{L})$  can be decomposed as pushouts over singleton path sets. The second says that every extension is determined by a root extension along the same label.

**Lemma 5.4 (Decomposition).** *Consider the following diagrams in  $\mathbf{SPath}(\mathcal{L})$*

$$\begin{array}{ccc}
 \mathcal{X} & \xleftarrow{\mathbf{f}_2} & \mathcal{X}_2 \\
 \mathbf{f}_1 \downarrow & & \downarrow \mathbf{f}_3 \\
 \mathcal{X}_1 & \xleftarrow{\mathbf{f}_4} & \mathcal{X}_1 \otimes \mathcal{X}_2
 \end{array}
 \qquad
 \begin{array}{ccc}
 \{\mathbf{P}\} & \xleftarrow{\bar{\mathbf{f}}_2} & \{\mathbf{P}_2\} \\
 \bar{\mathbf{f}}_1 \downarrow & & \downarrow \bar{\mathbf{f}}_3 \\
 \{\mathbf{P}_1\} & \xleftarrow{\bar{\mathbf{f}}_4} & \mathbf{P}_1 \otimes \mathbf{P}_2
 \end{array}$$

If the diagram on the left is a pushout, then for all  $P_1 \in \mathcal{X}_1$  and  $P_2 \in \mathcal{X}_2$  there are pushouts as shown on the right such that  $\mathbf{f}_i$  and  $\bar{\mathbf{f}}_i$  have the same underlying function on events.

**Lemma 5.5.** *Let  $\mathbf{f} : \mathbf{G} \rightarrow \mathbf{G}_1$  be a pr-morphism in  $\mathbf{PIDag}(\mathcal{L})$ . Then it is an  $\ell$ -extension if and only if there exists a pushout in  $\mathbf{PDag}(\mathcal{L})$*

$$\begin{array}{ccc} \bar{\mathbf{G}} & \xleftarrow{\bar{\mathbf{f}}} & \bar{\mathbf{G}}_1 \\ \downarrow & & \downarrow \\ \mathbf{G} & \xleftarrow{\mathbf{f}} & \mathbf{G}_1 \end{array}$$

such that  $\bar{\mathbf{f}}$  is a root  $\ell$ -extension.

We now show that transitions of the multi-replica LTS are those corresponding to  $\ell$ -extensions. Intuitively, an  $\ell$ -extension describes a “local” augmentation of a graph, corresponding to a step of computation of a single replica.

**Proposition 5.6.** *Let  $\mathbf{E}_m(\mathbb{F})$  be  $\mathbf{E}(\mathbb{F})$  restricted to  $\ell$ -extensions, for all  $\ell \in \mathcal{L}$ . Then the multi-replica LTS coincides with the LTS for  $\mathbf{E}_m(\mathbb{F})$ .*

*Proof.* We first show that states in  $\mathbf{E}(\mathbb{F})$  – and thus in  $\mathbf{E}_m(\mathbb{F})$ , as they have the same objects – are of the required form. If  $\mathbf{G}_1$  and  $\mathbf{G}_2$  are compatible, by Lemma 3.31 we obtain  $\mathbb{F}(\mathbf{G}_1 \sqcup \mathbf{G}_2) \cong \mathbb{F}(\mathbf{G}_1) \otimes \mathbb{F}(\mathbf{G}_2)$ , from which we have that for any  $\langle \mathbf{G}_1 \sqcup \mathbf{G}_2, \mathbf{P} \rangle$ , as in Definition 5.1, it holds  $\mathbf{P} \in \mathbb{F}(\mathbf{G}_1) \otimes \mathbb{F}(\mathbf{G}_2)$ , as required.

In order to show that multi-replica transitions correspond to morphisms of  $\mathbf{E}_m(\mathbb{F})$ , we will show that these morphisms are generated according to the rule (COMP), i.e., every  $\ell$ -extension in  $\mathbf{E}_m(\mathbb{F})$  is generated from a suitable rooted  $\ell$ -extension, and vice versa. We show one direction, the other one is analogous. Let  $\bar{\mathbf{f}} : \langle \mathbf{G}_1, \mathbf{P}_1 \rangle \rightarrow \langle \mathbf{G}'_1, \mathbf{P}'_1 \rangle$  be a rooted  $\ell$ -extension. For any graph  $\mathbf{G}_2$  compatible with  $\mathbf{G}_1$  we have the following pushout in  $\mathbf{PDag}(\mathcal{L})$

$$\begin{array}{ccc} \mathbf{G}_1 & \xleftarrow{\bar{\mathbf{f}}} & \mathbf{G}'_1 \\ \downarrow & & \downarrow \\ \mathbf{G}_1 \sqcup \mathbf{G}_2 & \xleftarrow{\mathbf{f}} & \mathbf{G}'_1 \sqcup \mathbf{G}_2 \end{array}$$

(note that  $\mathbf{G}'_1 \sqcup \mathbf{G}_1 \sqcup \mathbf{G}_2 \cong \mathbf{G}'_1 \sqcup \mathbf{G}_2$ , as  $\mathbf{G}_1$  is a sub-graph of  $\mathbf{G}'_1$ ), which implies by Lemma 5.5 that  $\mathbf{f}$  is a  $\ell$ -extension. We now have to show that for all  $\mathbf{P}' \in \mathbf{P} \otimes \mathbf{P}'_1$  there exist morphisms  $\mathbf{f} : \langle \mathbf{G}_1 \sqcup \mathbf{G}_2, \mathbf{P} \rangle \rightarrow \langle \mathbf{G}'_1 \sqcup \mathbf{G}_2, \mathbf{P}' \rangle$  in  $\mathbf{E}_m(\mathbb{F})$ . By coherence of  $\mathbb{F}$  and Lemma 3.31, we have that the pushout above is mapped by  $\mathbb{F}$  to a pushout in  $\mathbf{SPath}(\mathcal{L})$  of the form

$$\begin{array}{ccc} \mathbb{F}(\mathbf{G}_1) & \xleftarrow{\mathbb{F}(\bar{\mathbf{f}})} & \mathbb{F}(\mathbf{G}'_1) \\ \downarrow & & \downarrow \\ \mathbb{F}(\mathbf{G}_1 \sqcup \mathbf{G}_2) & \xleftarrow{\mathbb{F}(\mathbf{f})} & \mathbb{F}(\mathbf{G}_1 \sqcup \mathbf{G}_2) \otimes \mathbb{F}(\mathbf{G}'_1) \end{array}$$

By Lemma 5.4, this implies the existence of a pushout for the span  $\{P\} \leftarrow \{P_1\} \hookrightarrow \{P'_1\}$ , where the bottom leg is a ps-morphism  $\{P\} \hookrightarrow P \otimes P'_1$  with the same underlying function on events as  $\mathbb{F}(\mathbf{f})$ . Therefore, by definition of ps-morphism and  $P \in \mathbb{F}(\mathbf{G}_1 \sqcup \mathbf{G}_2)$ , and by  $P'_1 \in \mathbb{F}(\mathbf{G}'_1)$ , we obtain

$$P \otimes P'_1 \subseteq \text{sat}(P, \mathbb{F}(\mathbf{f})) \cap \mathbb{F}(\mathbf{G}_1 \sqcup \mathbf{G}_2) \otimes \mathbb{F}(\mathbf{G}'_1)$$

Noting that the formula on the right coincides with  $\mathbb{F}(\mathbf{f})(P)$ , we obtain the existence of the desired morphisms in  $\mathbf{E}_m(\mathbb{F})$ , by Definition 5.1.  $\square$

**Example 5.7.** *Consider once more the category  $\mathbf{E}(\mathbb{M}(\mathcal{S}_{mvR}))$  discussed in Example 5.2. By Proposition 5.3, the behaviour of a single replica is characterised by morphisms associated with root extensions. The morphism  $\mathbf{f}: \langle \mathbf{G}_2, P_2 \rangle \rightarrow \langle \mathbf{G}_3, P_3 \rangle$  described in Example *refex:catelem* corresponds to a one-replica transition with label  $\langle \mathbf{rd}, \{2, 3\} \rangle$ . In fact, its underlying pr-morphism  $\mathbf{f}: \mathbf{G}_2 \rightarrow \mathbf{G}_3$  is a root extension, accounting for the occurrence of the event  $\langle \mathbf{rd}, \{2, 3\} \rangle$  that sees any other event in the configuration; this may happen locally if all events in other replicas have been already propagated. Contrastingly, the extension  $\mathbf{f}': P_1 \rightarrow \mathbf{G}_2$  accounts for a new event  $\langle \mathbf{wr}(3), \text{ok} \rangle$  that is unaware of any other event, and hence, executed in a completely different replica. This corresponds to multi-replica transitions  $\mathbf{f}': \langle P_1, P \rangle \rightarrow \langle \mathbf{G}_2, P' \rangle$ , where  $P'$  arbitrates the additional event  $\langle \mathbf{wr}(3), \text{ok} \rangle$  anywhere in  $P$ . The local one-replica execution originating this transition is obtained via Lemma 5.5: It is  $\mathbf{f}: \langle \emptyset, \emptyset \rangle \rightarrow \langle \langle \mathbf{wr}(3), \text{ok} \rangle, \langle \mathbf{wr}(3), \text{ok} \rangle \rangle$  (here we write  $\langle \mathbf{wr}(3), \text{ok} \rangle$  for the one-node graph/path), with the underlying pr-morphism  $\emptyset \rightarrow \langle \mathbf{wr}(3), \text{ok} \rangle$  a root morphism.*

## 5.2. Context LTS

So far, the category of elements allowed to recast the set-theoretical presentation of one- and multi-replica LTSs. However, its strength is in allowing to obtain a new LTS that is reminiscent of the category of contexts à la Leifer-Milner [17], where arrows represent contexts enabling a transition from the source to the target of the arrow. Here, observations are pairs of an event plus an embedding that records how the resulting local visibility embeds into the global one. These additional observations will be needed for defining a correct notion of simulation.

**Definition 5.8.** *The context LTS is obtained by taking elements  $\langle \mathbf{G}, P \rangle$  of  $\mathbf{E}(\mathbb{F})$  as states, and by labelled transitions triples*

$$\langle \mathbf{G}, P \rangle \xrightarrow{\langle \mathbf{f}, \bar{\mathbf{f}} \rangle} \langle \mathbf{G}_1, P_1 \rangle$$

*such that  $\mathbf{f}: \langle \mathbf{G}, P \rangle \rightarrow \langle \mathbf{G}_1, P_1 \rangle$  and  $\bar{\mathbf{f}}: \langle \bar{\mathbf{G}}_1, \bar{P}_1 \rangle \rightarrow \langle \mathbf{G}_1, P_1 \rangle$  are arrows of  $\mathbf{E}(\mathbb{F})$  and there exists a pushout in  $\mathbf{PDag}(\mathcal{L})$*

$$\begin{array}{ccc} \bar{\mathbf{G}} & \hookrightarrow & \bar{\mathbf{G}}_1 \\ \downarrow & & \downarrow \bar{\mathbf{f}} \\ \mathbf{G} & \xrightarrow{\mathbf{f}} & \mathbf{G}_1 \end{array}$$

Note that each arrow  $\mathbf{f} : \langle \mathbf{G}, \mathbf{P} \rangle \rightarrow \langle \mathbf{G}_1, \mathbf{P}_1 \rangle$  of  $\mathbf{E}(\mathbb{F})$  induces at least one labelled transition (it suffices to consider for  $\bar{\mathbf{f}}$  the identity of  $\mathbf{G}_1$ ), but they could actually be more. In fact, all labels can be obtained constructively, since in  $\mathbf{PDag}(\mathcal{L})$  pushouts along monos are also pullbacks, and the arrow  $[\mathbf{f}, \bar{\mathbf{f}}] : \mathbf{G} + \bar{\mathbf{G}}_1 \rightarrow \mathbf{G}_1$  uniquely induced by the coproduct must be epi.

Also note that if we restrict to consider only injections, the pair  $\langle \mathbf{f}, \bar{\mathbf{f}} \rangle$  is uniquely characterized by  $\langle \bar{\mathbf{G}}_1, \bar{\mathbf{P}}_1 \rangle$ . In order to simplify some definitions, in the following we abuse notation and denote as  $\langle \bar{\mathbf{G}}_1, \bar{\mathbf{P}}_1 \rangle$  such a label  $\langle \mathbf{f}, \bar{\mathbf{f}} \rangle$ .

Finally, it is noteworthy that the context LTS includes also the one- and multi-replica, as stated by the result below.

**Lemma 5.9.** *The LTS for  $\mathbf{E}_o(\mathbb{F})$  ( $\mathbf{E}_m(\mathbb{F})$ ) coincides with the restriction of the context LTS to transitions whose labels are pairs  $\langle \mathbf{f}, \text{id} \rangle$ , where  $\mathbf{f}$  is a root extension (an extension, respectively).*

For the former, note that if  $\mathbf{f}$  is a root extension, then a mono pr-morphism forming a pushout square has to be an isomorphism. Instead, should  $\mathbf{f}$  be an extension, a few alternatives for the second component of the label are available, such as taking for  $\bar{\mathbf{G}}_1$  the smallest graph such that  $\mathbf{G} \sqcup \bar{\mathbf{G}}_1 = \mathbf{G}_1$ . However, the choice is immaterial for our later results on simulation, and further abusing notation we simply denote as  $\ell$  a label  $\langle \mathbf{f}, \text{id} \rangle$  such that  $\mathbf{f}$  is an  $\ell$ -extension.

**Example 5.10.** *Consider the multi-replica transition arrow  $\mathbf{f}' : \langle \mathbf{P}_1, \mathbf{P} \rangle \rightarrow \langle \mathbf{G}_2, \mathbf{P}' \rangle$  of Example 5.7. Since we have the arrow  $\mathbf{f} : \langle \emptyset, \emptyset \rangle \rightarrow \langle \langle \mathbf{wr}(3), \text{ok} \rangle, \langle \mathbf{wr}(3), \text{ok} \rangle \rangle$  in the category of elements,  $\mathbf{f}'$  yields the following context LTS transition*

$$\langle \mathbf{P}_1, \mathbf{P} \rangle \xrightarrow{\langle \mathbf{f}', \bar{\mathbf{f}}' \rangle} \langle \mathbf{G}_2, \mathbf{P}' \rangle$$

where  $\bar{\mathbf{f}}'$  is the embedding of the one-node graph  $\langle \mathbf{wr}(3), \text{ok} \rangle$  into  $\mathbf{G}_2$ . As mentioned, we can just use  $\langle \langle \mathbf{wr}(3), \text{ok} \rangle, \langle \mathbf{wr}(3), \text{ok} \rangle \rangle$  as label, since this is uniquely determined. This label conveys the information about the resulting visibility and arbitration pair for the acting replica.

### 5.3. Implementation correctness via simulation

We are now ready to characterise implementation correctness as a simulation relation between the context LTS and a RLTS implementation.

**Definition 5.11** (Implementation correctness). *Let  $\mathcal{S}$  be a specification,  $\mathcal{C}_{\mathcal{S}}$  the context LTS, and  $\mathcal{I}_{\mathcal{S}} = \{\mathcal{I}_r\}_{r \in \mathcal{R}}$  the RLTS implementation. An implementation relation  $\mathcal{R}_{\mathcal{S}} = \{\mathcal{R}_r\}_{r \in \mathcal{R}}$  is a  $\mathcal{R}$ -family of relations such that  $\mathcal{R}_r$  relates states in  $\mathcal{I}_r$  and  $\mathcal{C}_{\mathcal{S}}$  and if  $(\sigma, \langle \mathbf{G}, \mathbf{P} \rangle) \in \mathcal{R}_r$  then*

1. if  $\sigma \downarrow b$  then  $\exists \mathbf{G}', \mathbf{P}'$  such that  $\langle \mathbf{G}, \mathbf{P} \rangle \xrightarrow{b} \langle \mathbf{G}', \mathbf{P}' \rangle$  and  $(\sigma, \langle \mathbf{G}', \mathbf{P}' \rangle) \in \mathcal{R}_r$ ;
2. if  $\sigma \xrightarrow{m} \sigma'$  then  $\exists \mathbf{G}', \mathbf{P}'$  such that  $\langle \mathbf{G}, \mathbf{P} \rangle \xrightarrow{m} \langle \mathbf{G}', \mathbf{P}' \rangle$  and  $(\sigma', \langle \mathbf{G}', \mathbf{P}' \rangle) \in \mathcal{R}_r$ ;
3. if  $\sigma \xrightarrow{\sigma'} \sigma''$  then  $\exists \mathbf{G}', \mathbf{G}'', \mathbf{P}', \mathbf{P}''$  such that  $\langle \mathbf{G}, \mathbf{P} \rangle \xrightarrow{\langle \mathbf{G}', \mathbf{P}' \rangle} \langle \mathbf{G}'', \mathbf{P}'' \rangle$ ,  $(\sigma', \langle \mathbf{G}', \mathbf{P}' \rangle) \in \mathcal{R}_{\mathcal{S}}$ , and  $(\sigma'', \langle \mathbf{G}'', \mathbf{P}'' \rangle) \in \mathcal{R}_r$ .

We write  $\sim_S$  for the largest implementation relation.

The notion of implementation correctness coincides with the one given in [9, Definition 21]. We may introduce one-replica implementation relations, and the associated notion of correctness, just by suitably restricting the context LTS, that is, by requiring  $b$  and  $m$  in items 1 and 2 above to arise from a root extension instead of an extension.

**Example 5.12.** *We now look at the GCounter RDT and we exhibit a one-replica implementation relation for it.*

*Restricting the context LTS for  $\mathcal{S}_{GC}$  to its one-replica labels, we find the following three types of transitions*

$$T1: \langle \mathbf{G}, \mathbf{P} \rangle \xrightarrow{\langle \text{rd}, k \rangle} \langle \mathbf{G}^{\langle \text{rd}, k \rangle}, \mathbf{P}' \rangle, \text{ arising from a rooted } \langle \text{rd}, k \rangle\text{-extension } \mathbf{G} \rightarrow \mathbf{G}^{\langle \text{rd}, k \rangle}.$$

$$T2: \langle \mathbf{G}, \mathbf{P} \rangle \xrightarrow{\langle \text{inc}, ok \rangle} \langle \mathbf{G}^{\langle \text{inc}, ok \rangle}, \mathbf{P}' \rangle, \text{ arising from a rooted } \langle \text{inc}, ok \rangle\text{-extension } \mathbf{G} \rightarrow \mathbf{G}^{\langle \text{inc}, ok \rangle}.$$

$$T3: \langle \mathbf{G}, \mathbf{P} \rangle \xrightarrow{\langle \mathbf{G}', \mathbf{P}' \rangle} \langle \mathbf{G}'', \mathbf{P}'' \rangle, \text{ arising from a suitable pushout in } \mathbf{PDag}(\mathcal{L}).$$

*Given a RTLS implementation  $\mathcal{I}_{\mathcal{S}_{GC}}$ , defined according to Example 4.10, let  $v_r$  denote the state  $v$  in  $\mathcal{I}_r$ . We have that the following family of relations*

$$\mathcal{R}_r = \left\{ (v_r, \langle \mathbf{G}, \mathbf{P} \rangle) \left| \begin{array}{l} \exists f: \mathcal{E}_{\mathbf{G}} \rightarrow \mathcal{R} \text{ s.t.} \\ \forall \mathbf{s} \in \mathcal{R}. v_r(\mathbf{s}) = \#\{\mathbf{e} \mid f(\mathbf{e}) = \mathbf{s} \wedge \lambda_{\mathbf{G}} = \langle \text{inc}, ok \rangle\} \end{array} \right. \right\}$$

*is a one-replica implementation relation. The proof is a straightforward adaptation of [9, Example 16] and is carried out by case analysis, noting that the transition types T1-T3 above correspond to the tree cases of Definition 5.11.*

## 6. Implementation model

In this section we present our model for implementations. Similarly to what we did for specifications, the aim is to characterise RLTSs by means of suitable functors. Our models are based on a power-domain construction, modelling non-determinism.

### 6.1. Implementations as functors over power-monoids

Our model for implementations is inspired by [23], where LTSs are given as functors from the free monoid over labels, represented as a one-object category, to a category of non-deterministic computations. This allows modelling sequences of transitions as compositions of computations. In our setting, we use a *one-replica* category representing the free monoid of actions of a replica.

**Definition 6.1** (One-replica category). *The one-replica category  $\mathbf{IR}$  is the category with one object and whose arrows are words over  $\mathcal{M}$ .*



**Example 6.2.** *The one-replica category  $\mathbf{IR}$  for the Multi-value Register in Ex. 4.14, has a unique object  $r$  and the arrows  $w : r \rightarrow r$  with  $w \in \{\langle \mathbf{wr}(k), ok \rangle \mid k \in \mathcal{V}\}^*$  with  $\iota_r = \epsilon : r \rightarrow r$ .*

In order to capture the behaviour of a set of replicas  $\mathcal{R}$ , we need to account for the fact that single replicas must show the “same” behaviour. For instance, in the *Multi-value Register* implementation LTS (see Ex. 4.13), the (WRITE) operation on two replicas  $r$  and  $s$  have to return exactly the same set of version vectors  $\mathbb{N}^{\mathcal{R}}$ , up to a swapping of  $r$  and  $s$  in their domain.

We formalise this constraint by introducing the category  $\mathbf{IR}(\mathcal{R})$ , containing  $\#\mathcal{R}$  isomorphic copies of  $\mathbf{IR}$ . That is, that category comes equipped with isomorphisms  $\iota_{r,s} : r \rightarrow s$  for each  $r, s \in \mathcal{R}$  such that  $\iota_{r,r} = id_r = \iota_{r,s}; \iota_{s,r}$ . Furthermore, we require a *naturality* condition, namely, for all words  $w$  over  $\mathcal{M}$  we have  $\iota_{r,s}; w = w; \iota_{r,s}$ . This constraint precisely enforces the requirement on the behaviour of the single replicas, which are now naturally isomorphic. For the sake of clarity, we usually suffix arrows associated to elements of  $\mathcal{M}$  with the replica they belong to, e.g.  $m : r \rightarrow r$  is denoted as  $m_r$ .

**Example 6.3.** *Consider again the Multi-value Register and two replicas, i.e.,  $\mathcal{R} = \{r, s\}$ . Then,  $\mathbf{IR}(\mathcal{R})$  has two objects  $r$  and  $s$  and the arrows  $w : r_1 \rightarrow r_2$  with  $r_1, r_2 \in \mathcal{R}$  and  $w \in \{\langle \mathbf{wr}(k), ok \rangle \mid k \in \mathcal{V}\}^*$ . We identify  $\iota_{r_1, r_2}$  with  $\epsilon : r_1 \rightarrow r_2$  for  $r_1, r_2 \in \mathcal{R}$ . We write  $\langle \mathbf{wr}(k), ok \rangle_r$  for the arrow  $\langle \mathbf{wr}(k), ok \rangle : r \rightarrow r$ , i.e., the execution of the mutator  $\langle \mathbf{wr}(k), ok \rangle$  over the replica  $r$ . Hence,  $\langle \mathbf{wr}(k_1), ok \rangle_r; \langle \mathbf{wr}(k_2), ok \rangle_s$  stands for a computation that consists of two writes: Firstly,  $k_1$  is written over  $r$ , then,  $k_2$  is written over  $s$ .*

We now move to define a category where the arrows of  $\mathbf{IR}(\mathcal{R})$  are interpreted as non-deterministic computations. The key idea is to generalise the development of Section 4.1 and represent the state-space as a monoid, and transitions as monoid endomorphisms. The first step is to consider the category  $\mathbf{Mon}$  of monoids and monoid homomorphisms. This is not enough, since here the arrows of  $\mathbf{IR}(\mathcal{R})$  would be interpreted as functions mapping each state to a single state, i.e., as deterministic computations. To represent nondeterministic computations, we consider the *powerset functor*  $\mathbf{P} : \mathbf{Set} \rightarrow \mathbf{Set}$ , mapping a set to its subsets. It is well-known that this functor forms a monad (here we omit unit and multiplication), and that it lifts to monoids, i.e., to a functor  $\mathbf{P} : \mathbf{Mon} \rightarrow \mathbf{Mon}$ . We can then consider the associated *Kleisli* category  $\mathbf{P}(\mathbf{Mon})$ , where objects are those of  $\mathbf{Mon}$  and arrows  $\mathcal{M}_1 \rightarrow \mathcal{M}_2$  are monoid homomorphisms  $\mathcal{M}_1 \rightarrow \mathbf{P}(\mathcal{M}_2)$ , representing relations that respect the monoid structure. In our setting, these will be used to represent transition relations of monoidal RLTSs.

**Definition 6.4** (Functorial implementation). *A functorial implementation of  $\mathcal{R}$  is a functor  $\mathbb{I} : \mathbf{IR}(\mathcal{R}) \rightarrow \mathbf{P}(\mathbf{Mon})$ .*

An implementation functor thus maps each replica into isomorphic monoids of possible states, and the arrows of a replica are mapped to relations over such sets of states, with the additional property of being compositional with

respect to state composition. More precisely,  $\mathbb{I}(m_r)(\sigma)$  is the set of states of the replica  $r$  that are reachable from  $\sigma$  after observing  $m_r$ . The naturality of the isomorphisms  $\iota_{r,s}$  ensures that the replicas exhibit the “same” behaviour, yet up to isomorphism, which takes care of the possible permutations among replicas.

**Example 6.5.** *An implementation functor for the Multi-value Register over a set of replicas  $\mathcal{R}$  can be defined as follows*

- We assume that all replicas are implemented in the same way (i.e., they only differ in their identifier). For this reason, their states exhibit the same structure, i.e., the implementation monoid  $\langle \Sigma, \oplus, 1 \rangle$  defined in Example 4.13. Hence, the implementation functor maps each replica is mapped to same implementation monoid, i.e.,  $\forall r. \mathbb{I}(r) = \langle \Sigma, \oplus, 1 \rangle$ .
- Each arrow  $\langle \mathbf{wr}(k), ok \rangle : r \rightarrow r$  corresponding to the execution of the mutator  $\langle \mathbf{wr}(k), ok \rangle$  over the replica  $r$  is mapped to a transition relation corresponding to that replica, i.e.,

$$\mathbb{I}(\langle \mathbf{wr}(k), ok \rangle_r)(S, v) = \{(S', v') \mid (S, v) \xrightarrow{\langle \mathbf{wr}(k), ok \rangle}^r (S', v'), \\ S' = (r, v(r) + 1, k), v' = v[r \mapsto v(r) + 1]\}$$

- $\mathbb{I}(\iota_{r,s}) = \varphi_{r,s}$  are RLTS isomorphisms.

We are now ready to give the main correspondence theorem of this section.

**Theorem 6.6.** *There is a one-to-one correspondence between monoidal, state-deterministic RLTS implementations and functorial implementations.*

*Proof.* We first show that a monoidal, state-deterministic RLTS implementation  $\mathcal{I}$  over  $\mathcal{R}$  induces a functor  $\mathbb{I}(\mathcal{I}): \mathbf{IR}(\mathcal{R}) \rightarrow \mathbf{P}(\mathbf{Mon})$ . Suppose  $\mathcal{I} = \{\langle \Sigma^r, \oplus^r, 1^r, \rightsquigarrow^r \rangle\}_{r \in \mathcal{R}}$ , and let  $\varphi_{r,s}$  be the isomorphism between  $\mathcal{I}_r$  and  $\mathcal{I}_s$ . Then we define

$$\begin{aligned} \mathbb{I}(\mathcal{I})(r) &= \langle \Sigma^r, \oplus^r, 1^r \rangle \\ \mathbb{I}(\mathcal{I})(m_r)(\sigma) &= \{\sigma' \mid \sigma \xrightarrow{m}^r_{\mathcal{M}} \sigma'\} \\ \mathbb{I}(\mathcal{I})(\iota_{r,s}) &= \varphi_{r,s} \end{aligned}$$

with the obvious identity mappings. This functor is well-defined, in fact:  $\mathbb{I}(\mathcal{I})(m_r)$  is a monoid homomorphism by monoidality of  $\mathcal{I}$ ;  $\mathbb{I}(\mathcal{I})(m_r; m'_r) = \mathbb{I}(\mathcal{I})(m_r); \mathbb{I}(\mathcal{I})(m'_r)$ , as both sides amount to making a  $m$  transition followed by a  $m'$  transition in  $\mathcal{I}_r$ ; and finally, the equations  $\iota_{r,s}; w = w; \iota_{r,s}$  in  $\mathbf{IR}(\mathcal{R})$  are preserved, because they amount to  $\mathbb{I}(\mathcal{I})(\iota_{r,s}) = \varphi_{r,s}$  being a RLTS homomorphism, hence commuting with the transition structures of  $\mathcal{I}_r$  and  $\mathcal{I}_s$ .

Although  $\mathbb{I}(\mathcal{I})$  does not explicitly represent transitions labelled by states (these do not occur as arrows of  $\mathbf{IR}(\mathcal{R})$ ), they are implicitly represented in the target category  $\mathbf{P}(\mathbf{Mon})$  as the monoidal structure. In fact, since the  $\mathcal{I}_r$  are

all state-deterministic,  $\rightsquigarrow$  is the closure of  $\rightsquigarrow_{\mathcal{M}}$  with respect to the monoidal operation.

We conclude by observing that the correspondence given above is actually bijective: two distinct RLTSs will yield two distinct functors, and every functor  $\mathbf{IR}(\mathcal{R}) \rightarrow \mathbf{P}(\mathbf{Mon})$  can be turned into a RLTS using the equations above.  $\square$

## 6.2. From functorial implementations to LTSs

We can distill an LTS from a functorial implementation similarly to how we obtained context LTSs from functorial specifications. Given a functorial implementation  $\mathbb{I}$ , we can define its category of elements  $\mathbf{E}(\mathbb{I})$  as in Definition 5.1. It is rewritten here for the sake of clarity.

**Definition 6.7** (Category of elements,  $\mathbb{I}$ ). *The category of elements  $\mathbf{E}(\mathbb{I})$  of  $\mathbb{I}$  is obtained as*

- states are pairs  $\langle r, \sigma \rangle$ , such that  $r \in \mathcal{R}$  and  $\sigma \in \mathbb{I}(r)$ ;
- arrows  $g : \langle r, \sigma \rangle \rightarrow \langle s, \sigma' \rangle$  are arrows  $g : r \rightarrow s$  such that  $\sigma' \in \mathbb{I}(g)(\sigma)$ .

It suffices to apply a simpler machinery than the one used for context LTSs.

**Definition 6.8** (Implementation LTS). *The implementation LTS is obtained by taking elements  $\langle r, \sigma \rangle$  of  $\mathbf{E}(\mathbb{I})$  as states, and by labelled transitions triples*

$$\langle r, \sigma \rangle \xrightarrow{m_r} \langle r, \sigma' \rangle$$

such that  $m_r : \langle r, \sigma \rangle \rightarrow \langle r, \sigma' \rangle$  is an arrow of  $\mathbf{E}(\mathbb{I})$ .

We restricted to transitions over the same replica, but of course this is just for convenience, since all our examples fit into this pattern.

We note that an implementation LTS is in general smaller than the corresponding RLTS, as it lacks explicit state-labelled transitions; these instead are captured by the monoidal structure of the state-space.

## 7. Conclusions and further works

In our paper we considered RDTs, and we laid out the basis for an algebraic characterisation of their operational semantics as well as of their implementation correctness in terms of (higher-order) simulation. The core of our contribution lies precisely in the formalism behind such characterisations. Our proposal builds on [12, 10] and improves [11, 9] and similar set-theoretical characterisations, which are now made precise and recast into standard notions from the literature on algebraic specifications, thus allowing for the use of a large body of methods and techniques in the analysis of RDTs. We offered a few examples for showing the adequateness of our proposal, even if its strength need to be further checked by a larger number of case studies.

In order to stress the methodological points, we adopted some simplifications. The most notable is the removal of the `snd` label from our transition systems.

Indeed, in our examples, and, in fact, in most case studies we are aware of, a replica always spawns a full copy of itself, thus from the point of view of simulation it is irrelevant, and it would be in any case captured by the identity arrow on the category of replicas. The modelling of replica communication [9], where the action of sending plays a larger role, is the subject of ongoing work.

Our construction of transition systems out of a category of elements follows an already established pattern for presheaves and simulation, most notably in [23]. The distilling of labels is clearly reminiscent of the *contexts as labels* paradigm advanced by Leifer and Milner [17], and it would fit in its less constrained version proposed in [4]. Since this was not the main methodological issue of the paper, we adopted a presentation requiring some ingenuity.

Our set-theoretical definition of the operational semantics of an implementation focuses on the algebraic structure of replica states. In this respect, it shares aims with the approach presented in [1] for implementation of *Commutative Replicated Data Types* (CRDTs). There, the chosen structure for states are join-semilattices, instead of monoids. As a consequence, mutators in [1] are required to be inflations on the applied states while mutators in our setting must preserve the monoidal structure of the states. As far as we know, the approach presented in this paper is the first one that exploits a monoidal structure for LTSs to obtain a functorial characterisation of implementations.

## References

- [1] Baquero, C., Almeida, P.S., Cunha, A., Ferreira, C.: Composition in state-based replicated data types. *Bullettion of the EATCS* **123** (2017)
- [2] Bieniussa, A., Zawirski, M., Preguiça, N.M., Shapiro, M., Baquero, C., Balesgas, V., Duarte, S.: An optimized conflict-free replicated set. *CoRR* **abs/1210.3368** (2012)
- [3] Blackburn, P., de Rijke, M., Venema, Y.: *Modal Logic*, Cambridge Tracts in Theoretical Computer Science, vol. 53. Cambridge University Press (2001)
- [4] Bonchi, F., Gadducci, F., Monreale, G.V.: A general theory of barbs, contexts, and labels. *ACM Transactions on Computational Logic* **15**(4), 35:1–35:27 (2014)
- [5] Bouajjani, A., Enea, C., Hamza, J.: Verifying eventual consistency of optimistic replication systems. In: Jagannathan, S., Sewell, P. (eds.) *POPL 2014*. pp. 285–296. ACM (2014)
- [6] Burckhardt, S., Gotsman, A., Yang, H.: Understanding eventual consistency. *Tech. Rep. MSR-TR-2013-39*, Microsoft Research (2013)
- [7] Burckhardt, S., Gotsman, A., Yang, H., Zawirski, M.: Replicated data types: specification, verification, optimality. In: Jagannathan, S., Sewell, P. (eds.) *POPL 2014*. pp. 271–284. ACM (2014)

- [8] Cerone, A., Bernardi, G., Gotsman, A.: A framework for transactional consistency models with atomic visibility. In: Aceto, L., de Frutos-Escrig, D. (eds.) CONCUR 2015. LIPIcs, vol. 42, pp. 58–71. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2015)
- [9] Gadducci, F., Melgratti, H., Roldán, C.: On the semantics and implementation of replicated data types. *Science of Computer Programming* **167**, 91–113 (2018)
- [10] Gadducci, F., Melgratti, H., Roldán, C., Sammartino, M.: Implementation correctness for replicated data types, categorically. In: Pun, V.K.I., Stolz, V., Simão, A. (eds.) ICTAC 2020. pp. 283–303. Springer (2020)
- [11] Gadducci, F., Melgratti, H.C., Roldán, C.: A denotational view of replicated data types. In: Jacquet, J., Massink, M. (eds.) COORDINATION 2017. LNCS, vol. 10319, pp. 138–156. Springer (2017)
- [12] Gadducci, F., Melgratti, H.C., Roldán, C., Sammartino, M.: A categorical account of replicated data types. In: Chattopadhyay, A., Gastin, P. (eds.) FSTTCS 2019. LIPIcs, vol. 150, pp. 42:1–42:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2019)
- [13] Gotsman, A., Burckhardt, S.: Consistency models with global operation sequencing and their composition. In: Richa, A.W. (ed.) DISC 2017. LIPIcs, vol. 91, pp. 23:1–23:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2017)
- [14] Gotsman, A., Yang, H.: Composite replicated data types. In: Vitek, J. (ed.) ESOP 2015. LNCS, vol. 9032, pp. 585–609. Springer (2015)
- [15] Kaki, G., Earanky, K., Sivaramakrishnan, K.C., Jagannathan, S.: Safe replication through bounded concurrency verification. In: OOPSLA 2018. PACMPL, vol. 2, pp. 164:1–164:27. ACM (2018)
- [16] Lanese, I., Pérez, J.A., Sangiorgi, D., Schmitt, A.: On the expressiveness and decidability of higher-order process calculi. *Information and Computation* **209**(2), 198–226 (2011)
- [17] Leifer, J.J., Milner, R.: Deriving bisimulation congruences for reactive systems. In: Palamidessi, C. (ed.) CONCUR 2000. LNCS, vol. 1877, pp. 243–258. Springer (2000)
- [18] MacLane, S., Moerdijk, I.: *Sheaves in geometry and logic: A first introduction to topos theory*. Springer (1992)
- [19] Milner, R., Sangiorgi, D.: Barbed bisimulation. In: Kuich, W. (ed.) ICALP 1992. LNCS, vol. 623, pp. 685–695. Springer (1992)
- [20] Shapiro, M., Preguiça, N., Baquero, C., Zawirski, M.: Conflict-free replicated data types. In: Défago, X., Petit, F., Villain, V. (eds.) SSS 2011. LNCS, vol. 6976, pp. 386–400. Springer (2011)

- [21] Shapiro, M., Preguiça, N., Baquero, C., Zawirski, M.: A comprehensive study of convergent and commutative replicated data types. Tech. Rep. RR-7506, Inria-Centre Paris-Rocquencourt (2011)
- [22] Sivaramakrishnan, K.C., Kaki, G., Jagannathan, S.: Declarative programming over eventually consistent data stores. In: Grove, D., Blackburn, S. (eds.) PLDI 2015. pp. 413–424. ACM (2015)
- [23] Sobociński, P.: Relational presheaves, change of base and weak simulation. *Computer and System Sciences* **81**(5), 901–910 (2015)
- [24] Zawirski, M., Baquero, C., Bieniusa, A., Preguiça, N.M., Shapiro, M.: Eventually consistent register revisited. In: Alvaro, P., Bessani, A. (eds.) PaPoC@EuroSys 2016. pp. 9:1–9:3. ACM (2016)