

# A Denotational View of Replicated Data Types

Fabio Gadducci<sup>1</sup>, Hernán Melgratti<sup>2,3</sup>, and Christian Roldán<sup>2</sup>(✉)

<sup>1</sup> Dipartimento di Informatica, Università di Pisa, Pisa, Italy

<sup>2</sup> Departamento de Computación, FCEyN,  
Universidad de Buenos Aires, Buenos Aires, Argentina  
croidan@dc.uba.ar

<sup>3</sup> CONICET, Buenos Aires, Argentina


**Abstract.** “Weak consistency” refers to a family of properties concerning the state of a distributed system. One of the key issues in their description is the way in which systems are specified. In this regard, a major advance is represented by the introduction of Replicated Data Types (RDTs), in which the meaning of operators is given in terms of two relations, namely, visibility and arbitration. Concretely, a data type operation is defined as a function that maps visibility and arbitration into a return value. In this paper we recast such standard approaches into a denotational framework in which a data type is seen as a function that maps visibility into admissible arbitrations. This characterisation provides a more abstract view of RDTs that (i) highlights some of the implicit assumptions shared in operational approaches to specification; (ii) accommodates underspecification and refinement; (iii) enables a categorical presentation of RDT and the development of composition operators for specifications.

## 1 Introduction

Distributed systems replicate their state over different nodes in order to satisfy several non-functional requirements, such as performance, availability, and reliability. It then becomes crucial to keep a consistent view of the replicated data. However, this is a challenging task because consistency is in conflict with two common requirements of distributed applications: *availability* (every request is eventually executed) and tolerance to network *partitions* (the system operates even in the presence of failures that prevent communication among components). In fact, it is impossible for a system to simultaneously achieve strong Consistency, Availability and Partition tolerance [6]. Since many domains cannot renounce to availability and network partitions, developers need to cope with weaker notions

---

The first author has been partially supported by CONICET International Cooperation Grant 995/15. Research partially supported by UBACyT project 2014–2017 20020130200092BA and CONICET project PIP 11220130100148CO.



**Fig. 1.** A scenario for the replicated data type **Register**

of consistency by allowing, e.g., replicas to (temporarily) exhibit some discrepancies, as long as they eventually converge to the same state.

This setting challenges the way in which data are specified: states, state transitions and return values should account for the different views that a data item may simultaneously have. Consider a data type **Register** corresponding to a memory cell that is read and updated by using, respectively, operations  $rd$  and  $wr$ . In a replicated scenario, the value obtained when reading a register after two concurrent updates  $wr(0)$  and  $wr(1)$  (i.e., updates taking place over different replicas) is affected by the way in which updates propagate among the different replicas: it is perfectly possible that the result of the read is (i) undefined (when the read is performed over a third replica that has not received any of the updates), (ii) 0 or (iii) 1. Basically, the return value depends on the updates that are seen by that read operation. Choosing the return value is straightforward when a read sees just one update. This is less so if a read is performed over a replica that knows both updates, for allowing all replicas to (consistently) pick one of the available values. A common strategy for registers is that the *last-write wins*, i.e., the last update should be chosen when several concurrent updates are observed. This strategy implicitly assumes that all events in a system can be arranged in a total order. Several recent approaches focus on the operational specification of replicated data types [2–5, 7, 8, 12, 14]. Usually, the specification describes the meaning of an operation in terms of two different relations among events: *visibility*, which explains the causes for each result, and *arbitration*, which totally orders events. Consider the visibility relation  $V$  in Fig. 1a and the arbitrations  $A_1$  and  $A_2$  in Fig. 1b and c, respectively. The meaning of  $rd$  is defined such that  $rd(V, A_1) = 1$  and  $rd(V, A_2) = 0$ . We remark that operational approaches require specifications to be *functional*, i.e., for every operation, visibility and arbitration relation, there exists exactly one return value. In this way operational specifications commit to concrete policies for resolving conflicts.

This work aims at putting on firm grounds the operational approaches for RDTs by giving them a purely functional description and, eventually, a categorical one. In our view, RDTs are functions that map visibility graphs (i.e., configurations) into sets of admissible arbitrations, i.e., all executions that generate a particular configuration. In this setting, a configuration mapped to an empty set of admissible arbitrations stands for an unreachable configuration. We rely on such an abstract view of RDTs to highlight some of the implicit

assumptions shared by most of the operational approaches. In particular, we characterise operational approaches, such as [4, 12], as those specifications that satisfy three properties: besides the evident requirement of being *functional* (i.e., deterministic and total), they must be *coherent* (i.e., larger states are explained as the composition of smaller ones), and *saturated* (e.g., an unobserved operation can be arbitrated in any position, even before the events that it sees). We show this inclusion to be strict and discuss some interesting cases that do not fall in this class. Moreover, we show that functional characterisation elegantly accounts for underspecification and refinement, which are standard notions in data type specification.

Then, we develop a categorical presentation for specifications. We focus on coherent specifications and show that there is a one-to-one correspondence between coherent specifications and a particular class of functors from the category  $I(\mathcal{L})$  of labelled directed acyclic graphs and injective *past-reflecting morphism* (which are the dual notion of tp-morphisms [9]) to the category  $\mathcal{P}(\mathcal{L})$  of sets of paths and path-set morphisms preserving the initial object. As it is standard from classical results on algebraic specification theory, pullbacks and (a weak form of) pushouts in  $I(\mathcal{L})$  provide basic operators for composing specifications, and thus our functorial presentation is the first step towards a denotational semantics of RDTs (see e.g. [1] and the references therein).

The paper has the following structure. Section 2 introduces the basic definitions concerning labelled directed acyclic graphs. Section 3 discusses our functional mechanism for the presentation of Replicated Data Types. Section 4 compares our proposal with the classical operational one [2]. Section 5 illustrates a categorical characterisation for our proposal. Finally, in the closing section we draw some conclusions and highlight further developments.

## 2 Labelled Directed Acyclic Graphs

In this section we recall the basics of labelled directed acyclic graphs, which are used for our description of replicated data types. We rely on countable sets  $\mathcal{E}$  of events  $e, e', \dots, e_1, \dots$  and  $\mathcal{L}$  of labels  $\ell, \ell', \dots, \ell_1, \dots$


**Definition 1 (Labelled Directed Acyclic Graph).** *A Labelled Directed Acyclic Graph (LDAG) over a set of labels  $\mathcal{L}$  is a triple  $\mathbf{G} = \langle \mathcal{E}_{\mathbf{G}}, \prec_{\mathbf{G}}, \lambda_{\mathbf{G}} \rangle$  such that  $\mathcal{E}_{\mathbf{G}}$  is a set of events,  $\prec_{\mathbf{G}} \subseteq \mathcal{E}_{\mathbf{G}} \times \mathcal{E}_{\mathbf{G}}$  is a binary relation whose transitive closure is a strict partial order, and  $\lambda_{\mathbf{G}} : \mathcal{E}_{\mathbf{G}} \rightarrow \mathcal{L}$  is a labeling function. An LDAG  $\mathbf{G}$  is a path if  $\prec_{\mathbf{G}}$  is a strict total order.*

We write  $\mathbb{G}(\mathcal{L})$  and  $\mathbb{P}(\mathcal{L})$  to respectively denote the sets of all LDAGs and paths over  $\mathcal{L}$ . We use  $\mathbf{G}$  to range over  $\mathbb{G}(\mathcal{L})$  and  $\mathbf{P}$  to range over  $\mathbb{P}(\mathcal{L})$ . Moreover, we write  $\prec_{\mathbf{P}}$  instead of  $\prec_{\mathbf{P}}$  to make evident that paths are total orders. We say that  $\mathbf{P}$  is a path over  $\mathcal{E}$  if  $\mathcal{E}_{\mathbf{P}} = \mathcal{E}$  and write  $\mathbb{P}(\mathcal{E}, \lambda)$  for  $\{\mathbf{P} \mid \mathbf{P} \text{ is a path over } \mathcal{E} \text{ and } \lambda_{\mathbf{P}} = \lambda\}$ . We usually omit the subscript  $\mathbf{G}$  (or  $\mathbf{P}$ ) when referring to the elements of  $\mathbb{G}$  (of  $\mathbb{P}$ , respectively) when no confusion arises. We write  $\epsilon$  for the empty LDAG, i.e., such that  $\mathcal{E}_{\epsilon} = \emptyset$ .

**Definition 2 (Morphism).** An LDAG morphism  $\mathbf{f}$  from  $\mathbf{G}$  to  $\mathbf{G}'$ , written  $\mathbf{f} : \mathbf{G} \rightarrow \mathbf{G}'$ , is a mapping  $\mathbf{f} : \mathcal{E}_{\mathbf{G}} \rightarrow \mathcal{E}_{\mathbf{G}'}$  such that  $\lambda_{\mathbf{G}} = \mathbf{f}; \lambda_{\mathbf{G}'}$  and  $\mathbf{e} \prec_{\mathbf{G}} \mathbf{e}'$  implies  $\mathbf{f}(\mathbf{e}) \prec_{\mathbf{G}'} \mathbf{f}(\mathbf{e}')$ .

Hereafter we implicitly consider LDAGs up-to isomorphism, i.e., related by a bijective function that preserves and reflects the underlying relation.

*Example 1.* Consider the set  $\mathcal{L} = \{\langle \text{rd}, 0 \rangle, \langle \text{rd}, 1 \rangle, \langle \text{wr}(0), \text{ok} \rangle, \langle \text{wr}(1), \text{ok} \rangle\}$  of labels describing the operations of a 1-bit register. Each label is a tuple  $\langle \text{op}, \text{rv} \rangle$  where  $\text{op}$  denotes an operation and  $\text{rv}$  its return value. For homogeneity, we associate the return value  $\text{ok}$  to every write operation. Now, take the LDAG over  $\mathcal{L}$  defined as  $\mathbf{G}_1 = \langle \{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}, \prec, \lambda \rangle$  where  $\prec = \{(\mathbf{e}_1, \mathbf{e}_3), (\mathbf{e}_2, \mathbf{e}_3)\}$  and  $\lambda$  is such that  $\lambda(\mathbf{e}_1) = \langle \text{wr}(0), \text{ok} \rangle$ ,  $\lambda(\mathbf{e}_2) = \langle \text{wr}(1), \text{ok} \rangle$ ,  $\lambda(\mathbf{e}_3) = \langle \text{rd}, 0 \rangle$ . A graphical representation of  $\mathbf{G}_1$  is in Fig. 2a. Since we consider LDAGs up-to isomorphism, we do not depict events and write instead the corresponding labels when no confusion arises.  $\mathbf{G}_2$  is an LDAG where  $\prec_{\mathbf{G}_2}$  is empty. Neither  $\mathbf{G}_1$  nor  $\mathbf{G}_2$  is a path, because they are not total orders.  $\mathbf{P}_1$  in Fig. 2c is an LDAG that is also a path. Hereinafter we use undirected arrows when depicting paths and avoid drawing transitions that are obtained by transitivity, as shown in Fig. 2d. All LDAGs in Fig. 2 belong to  $\mathbb{G}(\mathcal{L})$ , but only  $\mathbf{P}_1$  is in  $\mathbb{P}(\mathcal{L})$ .



**Fig. 2.** Two simple LDAGs and two paths.

## 2.1 LDAG Operations

We now present a few operations on LDAGs, which will be used in the following sections. We start by introducing some notation for binary relations. We write  $\text{Id}$  for the identity relation over events and  $\preceq$  for  $\prec \cup \text{Id}$ . We write  $- \prec \mathbf{e}$  (and similarly  $- \preceq \mathbf{e}$ ) for the preimage of  $\mathbf{e}$ , i.e.,  $- \prec \mathbf{e} = \{\mathbf{e}' \mid \mathbf{e}' \prec \mathbf{e}\}$ . We use  $\prec|_{\mathcal{E}}$  for the restriction of  $\prec$  to elements in  $\mathcal{E}$ , i.e.  $\prec|_{\mathcal{E}} = \prec \cap (\mathcal{E} \times \mathcal{E})$ . Analogously,  $\lambda|_{\mathcal{E}}$  is the domain restriction of  $\lambda$  to the elements in  $\mathcal{E}$ . We write  $\mathcal{E}_{\top}$  for the extension of the set  $\mathcal{E}$  with a fresh element, i.e.,  $\mathcal{E}_{\top} = \mathcal{E} \cup \{\top\}$  such that  $\top \notin \mathcal{E}$ .

**Definition 3 (Restriction and Extension).** Let  $\mathbf{G} = \langle \mathcal{E}, \prec, \lambda \rangle$  and  $\mathcal{E}' \subseteq \mathcal{E}$ . We define

- $\mathbf{G}|_{\mathcal{E}'} = \langle \mathcal{E}', \prec|_{\mathcal{E}'}, \lambda|_{\mathcal{E}'} \rangle$  as the restriction of  $\mathbf{G}$  to  $\mathcal{E}'$ ;
- $\mathbf{G}_{\mathcal{E}'}^{\ell} = \langle \mathcal{E}_{\top}, \prec \cup (\mathcal{E}' \times \{\top\}), \lambda[\top \mapsto \ell] \rangle$  as the extension of  $\mathbf{G}$  over  $\mathcal{E}'$  with  $\ell$ .

Restriction obviously lifts to sets  $\mathcal{X}$  of LDAGs, i.e.,  $\mathcal{X}|_{\mathcal{E}} = \{G|_{\mathcal{E}} \mid G \in \mathcal{X}\}$ . We omit the subscript  $\mathcal{E}'$  in  $G_{\mathcal{E}'}$  when  $\mathcal{E}' = \mathcal{E}$ .

*Example 2.* Consider the LDAGs  $G_1$  and  $G_2$  depicted in Fig. 2a and b, respectively. Then,  $G_2 = G_1|_{-\prec_{e_3}}$  and  $G_1 = G_2^{\langle rd, 0 \rangle}$ .

The following operator allows for the combination of several paths and plays a central role in our characterisation of replicated data types.

**Definition 4 (Product).** Let  $\mathcal{X} = \{\langle \mathcal{E}_i, \prec_i, \lambda_i \rangle\}_i$  be a set of paths. The product of  $\mathcal{X}$  is

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

Intuitively, the product of paths is analogous to the synchronous product of transition systems, in which common elements are identified and the remaining ones can be freely interleaved, as long as the original orders are respected.

*Example 3.* Consider the paths  $P_1$  and  $P_2$  in Fig. 3, and assume that they share the event labelled  $\langle wr(2), ok \rangle$ . Their product has two paths  $P_3$  and  $P_4$ , each of them contains the elements of  $P_1$  and  $P_2$  and preserves the relative order of the elements in the original paths. We remark that the product is empty when the paths have incompatible orders. For instance,  $P_3 \otimes P_4 = \emptyset$ .




Fig. 3. Product between two paths.

It is straightforward to show that  $\otimes$  is associative and commutative. Hence, we freely use  $\otimes$  over sets of sets of paths.

### 3 Specifications

We introduce our notion of specification and applies it to some well-known data types.

**Definition 5 (Specification).** A specification  $S$  is a function  $S : \mathbb{G}(\mathcal{L}) \rightarrow 2^{\mathbb{P}(\mathcal{L})}$  such that  $S(\epsilon) = \{\epsilon\}$  and  $\forall G. S(G) \subseteq 2^{\mathbb{P}(\mathcal{E}_G, \lambda_G)}$ .

A specification  $\mathcal{S}$  maps an LDAG (i.e., a visibility relation) to a set of paths (i.e., its admissible arbitrations). Note that  $P \in \mathcal{S}(\mathbf{G})$  is a path over  $\mathcal{E}_{\mathbf{G}}$ , and hence a total order of the events in  $\mathbf{G}$ . However, we do not require  $P$  to be a topological ordering of  $\mathbf{G}$ , i.e.,  $\prec_{\mathbf{G}} \subseteq \prec_P$  may not hold. Although some specification approaches consider only arbitrations that include visibility [5, 7], our definition accommodates also presentations, such as [2, 4], in which arbitrations may not preserve visibility. We focus later on in a few subclasses, such as coherent specifications, in order to establish a precise correspondence with replicated data types. We also remark that it could be the case that  $\mathcal{S}(\mathbf{G}) = \emptyset$ , which means that  $\mathcal{S}$  forbids the configuration  $\mathbf{G}$  (more details in Example 4 below). For technical convenience, we impose  $\mathcal{S}(\epsilon) = \{\epsilon\}$  and disallow  $\mathcal{S}(\epsilon) = \emptyset$ :  $\mathcal{S}$  cannot forbid the empty configuration, which denotes the initial state of a data type.

We now illustrate the specification of some well-known replicated data types.

*Example 4 (Counter).* The data type **Counter** provides operations for incrementing and reading an integer register with initial value 0. A read operation returns the number of increments seen by that read. An increment is always successful and returns the value *ok*. Formally, we consider the set of labels  $\mathcal{L} = \{\langle \mathbf{inc}, ok \rangle\} \cup (\{\mathbf{rd}\} \times \mathbb{N})$ . Then, a **Counter** is specified by  $\mathcal{S}_{Ctr}$  defined such that

$$P \in \mathcal{S}_{Ctr}(\mathbf{G}) \text{ if } \forall e \in \mathcal{E}_{\mathbf{G}}. \lambda(e) = \langle \mathbf{rd}, k \rangle \text{ implies } k = \#\{e' \mid e' \prec_{\mathbf{G}} e \text{ and } \lambda(e') = \langle \mathbf{inc}, ok \rangle\}$$

A visibility graph  $\mathbf{G}$  has admissible arbitrations (i.e.,  $\mathcal{S}_{Ctr}(\mathbf{G}) \neq \emptyset$ ) only when each event  $e$  in  $\mathbf{G}$  labelled by **rd** has a return value  $k$  that matches the number of increments antecedent  $e$  in  $\mathbf{G}$ . We illustrate two cases for the definition of  $\mathcal{S}_{Ctr}$  in Fig. 4. While the configuration in Fig. 4a has admissible arbitrations, the one in Fig. 4b has not, because the unique event labelled by **rd** returns 0 when it is actually preceded by an observed increment. In other words, an execution is not allowed to generate such a visibility graph. We remark that  $\mathcal{S}_{Ctr}$  does not impose any constraint on the ordering  $\prec_P$ .

In fact, a path  $P \in \mathcal{S}_{Ctr}(\mathbf{G})$  does not need to be a topological ordering of  $\mathbf{G}$  as, for instance, the rightmost path in the set of Fig. 4a.

$$\mathcal{S}_{Ctr} \left( \begin{array}{c} \langle \mathbf{inc}, ok \rangle \\ \downarrow \\ \langle \mathbf{rd}, 1 \rangle \end{array} \right) = \left\{ \begin{array}{cc} \langle \mathbf{inc}, ok \rangle & \langle \mathbf{rd}, 1 \rangle \\ | & | \\ \langle \mathbf{rd}, 1 \rangle & \langle \mathbf{inc}, ok \rangle \end{array} \right\} \quad \mathcal{S}_{Ctr} \left( \begin{array}{c} \langle \mathbf{inc}, ok \rangle \\ \downarrow \\ \langle \mathbf{rd}, 0 \rangle \end{array} \right) = \emptyset$$

(a)
(b)

**Fig. 4.** Counter specification.

*Example 5 (Last-write-wins Register).* A **Register** stores a value that can be read and updated. We assume that the initial value of a register is undefined. We take  $\mathcal{L} = \{\langle \mathbf{wr}(k), ok \rangle \mid k \in \mathbb{N}\} \cup (\{\mathbf{rd}\} \times \mathbb{N} \cup \{\perp\})$  as the set of labels. The specification  $S_{lwwR}$  gives the semantics of a register that adopts the last-write-wins strategy.

$$P \in S_{lwwR}(\mathbf{G}) \text{ if } \forall e \in \mathcal{E}_{\mathbf{G}}. \begin{cases} \lambda(e) = \langle \mathbf{rd}, \perp \rangle \text{ implies } \forall e' \prec_{\mathbf{G}} e. \forall k. \lambda(e') \neq \langle \mathbf{wr}(k), ok \rangle \\ \lambda(e) = \langle \mathbf{rd}, k \rangle \text{ implies } \exists e' \prec_{\mathbf{G}} e. \lambda(e') = \langle \mathbf{wr}(k), ok \rangle \text{ and} \\ \forall e'' \prec_{\mathbf{G}} e. e' \prec_P e'' \text{ implies } \forall k'. \lambda(e'') \neq \langle \mathbf{wr}(k'), ok \rangle \end{cases}$$

An LDAG  $\mathbf{G}$  has admissible arbitrations only when each event associated with a read operation returns a previously written value. As per the first condition above, a read operation returns the undefined value  $\perp$  when it does not see any write. By the second condition, a read  $e$  returns a natural number  $k$  when it sees an operation  $e'$  that writes that value  $k$ . In such case, any admissible arbitration  $P$  must order  $e'$  as the greatest (accordingly to  $\prec_P$ ) of all write operations seen by  $e$ .

*Example 6 (Generic Register).* We now define a **Generic Register** that does no commit to a particular strategy for resolving conflicts. We specify this type as follows

$$P \in S_{gR}(\mathbf{G}) \text{ if } \forall e \in \mathcal{E}_{\mathbf{G}}. \begin{cases} \lambda(e) = \langle \mathbf{rd}, \perp \rangle \text{ implies } \forall e' \prec_{\mathbf{G}} e. \forall k. \lambda(e') \neq \langle \mathbf{wr}(k), ok \rangle \\ \lambda(e) = \langle \mathbf{rd}, k \rangle \text{ implies } \exists e' \prec_{\mathbf{G}} e. \lambda(e') = \langle \mathbf{wr}(k), ok \rangle \text{ and} \\ \forall e''. \lambda(e'') = \langle \mathbf{rd}, k'' \rangle \text{ and } - \prec_{\mathbf{G}} e = - \prec_{\mathbf{G}} e'' \text{ implies } k = k'' \end{cases}$$

As in Example 5, the return value of a read corresponds to a written value seen by that read, but the specification does not determine which value should be chosen. We require instead that all read operations with the same causes (i.e.,  $- \prec_{\mathbf{G}} e = - \prec_{\mathbf{G}} e'$ ) have the same result. Since this condition is satisfied by any admissible configuration  $\mathbf{G}$ , it ensures convergence. The fact that convergence is explicitly required contrasts with approaches like [2, 4], where on the contrary convergence is ensured automatically by considering only deterministic specifications. We remark that for the deterministic cases, e.g., Examples 4 and 5, we do not need to explicitly require convergence.

### 3.1 Refinement

Refinement is a standard approach in data type specification, which allows for a hierarchical organisation that goes from abstract descriptions to concrete implementations. The main benefit of refinement relies on the fact that applications can be developed and reasoned about in terms of abstract types, which hide implementation details and leave some freedom for the implementation. Consider the specification  $S_{gR}$  of the **Generic Register** introduced in Example 6,

which only requires a policy for conflict resolution that ensures convergence. On the contrary, the specification  $S_{lwwR}$  in Example 5 explicitly states that concurrent updates must be resolved by adopting the last-write-wins policy. Since the latter policy ensures convergence, we would like to think about  $S_{lwwR}$  as a refinement of  $S_{gR}$ . We characterise refinement in our setting as follows.

**Definition 6 (Refinement).** *Let  $S_1, S_2$  be specifications. We say that  $S_1$  refines  $S_2$  and we write  $S_1 \sqsubseteq S_2$  if  $\forall G. S_1(G) \subseteq S_2(G)$ .*

*Example 7.* It can be easily checked that  $P \in S_{lwwR}(G)$  implies  $P \in S_{gR}(G)$  for any  $G$ . Consequently,  $S_{lwwR}$  is a refinement of  $S_{gR}$ .

*Example 8.* Consider the data type **Set**, which provides (among others) the operations **add**, **rem** and **lookup** for respectively 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 of elements (see [13] for a detailed discussion). We illustrate two possible alternatives by considering the execution scenario depicted in Fig. 5. A reasonable semantics for **lookup** over  $G$  and  $P$  would fix the result  $V$  as either  $\emptyset$  or  $\{1\}$ . In fact, under the last-write-wins policy, the specification prescribes that **lookup** returns  $\{1\}$  in this scenario. Differently, the strategy of 2P-Sets<sup>1</sup> establishes that the result is  $\emptyset$ .

The following definition provides a specification for an abstract data type **Set** that allows (among others) any of the above policies.

$$P \in S_{Set}(G) \text{ if } \forall e \in \mathcal{E}_G. \lambda(e) = \langle \text{lookup}, V \rangle \text{ implies } B_e \subseteq V \subseteq A_e \text{ and } \text{Conv}_{e,V}$$

where

$$\begin{aligned} A_e &= \{k \mid e' \prec_G e \text{ and } \lambda(e') = \langle \text{add}(k), ok \rangle\} \\ B_e &= A_e \setminus \{k \mid e' \prec_G e \text{ and } \lambda(e') = \langle \text{rem}(k), ok \rangle\} \\ \text{Conv}_{e,V} &= \forall e'. \lambda(e') = \langle \text{lookup}, V' \rangle \text{ and } - \prec_G e = - \prec_G e' \text{ implies } V = V' \end{aligned}$$

The set  $A_e$  contains the elements added to (and possibly removed from) the set seen by  $e$  while  $B_e$  contains those elements for which  $e$  sees no removal. Thus, the condition  $B_e \subseteq V \subseteq A_e$  states that **lookup** returns a set that contains at least all the elements added but not removed (i.e., in  $B_e$ ). However, the return value  $V$  may contain elements that have been added and removed (the choice is left unspecified). Condition **Conv** ensures convergence, similarly to the specification of  $S_{gR}$  in Example 6.

Then, a concrete resolution policy such as 2P-Sets can be specified as follows

$$P \in S_{2P-Sets}(G) \text{ if } \forall e \in \mathcal{E}_G. \lambda(e) = \langle \text{lookup}, V \rangle \text{ implies } V = B_e$$

Clearly,  $S_{2P-Sets}$  is a refinement of  $S_{Sets}$ . Other policies can be specified analogously.

<sup>1</sup> In 2P-Sets, additions of elements that have been previously removed have no effect.






Fig. 5. A scenario for the replicated data type **Set**

### 3.2 Classes of Specifications

We now discuss two properties of specifications. Firstly, we look at specifications for which the behaviour of larger computations matches that of their shorter prefixes.

**Definition 7 (Past-Coherent Specification).** *Let  $S$  be a specification. We say that  $S$  is past-coherent (briefly, coherent) if*

$$\forall G. S(G) = \bigotimes_{e \in \mathcal{L}_G} S(G|_{\rightarrow_e})$$

Note that coherence implies that  $S(G)|_{\rightarrow_e} \subseteq S(G|_{\rightarrow_e})$ . Intuitively, sub-paths are obtained from the interleaving of the paths belonging to the associated sub-specifications.

*Example 9.* The specifications in Examples 4, 5 and 6 are all coherent, because their definitions are in terms of restrictions of the LDAGs. Now consider the specification  $S$  defined such that the equalities in Fig. 6 hold.  $S$  is not coherent

$$\begin{array}{ll}
 S(\langle o_1, v_1 \rangle) = \{ \langle o_1, v_1 \rangle \} & S(\langle o_1, v_1 \rangle \langle o_2, v_2 \rangle) = \left\{ \begin{array}{c} \langle o_1, v_1 \rangle \\ | \\ \langle o_2, v_2 \rangle \end{array} \right\} \\
 S(\langle o_2, v_2 \rangle) = \{ \langle o_2, v_2 \rangle \} & \text{(b)} \\
 \text{(a)} & 
 \end{array}$$
  

$$S\left( \begin{array}{cc} \langle o_1, v_1 \rangle & \langle o_2, v_2 \rangle \\ \searrow & \swarrow \\ & \langle o_3, v_3 \rangle \end{array} \right) = \left\{ \begin{array}{cc} \langle o_1, v_1 \rangle & \langle o_2, v_2 \rangle \\ | & | \\ \langle o_2, v_2 \rangle & \langle o_1, v_1 \rangle \\ | & | \\ \langle o_3, v_3 \rangle & \langle o_3, v_3 \rangle \end{array} \right\}$$

(c)

Fig. 6. A non-coherent specification.

because the arbitrations for the LDAG in Fig. 6b should contain all the interleavings for the paths associated with its sub-configurations, as depicted in Fig. 6a. Instead, note that the arbitration of  $\langle o_2, v_2 \rangle$  before  $\langle o_1, v_1 \rangle$  in the leftmost path on Fig. 6c would not hinder coherence by itself, even if it is not allowed by the sub-configuration in Fig. 6b.

A second class of specifications is concerned with saturation. Intuitively, a saturated specification allows every top element on the visibility to be arbitrated in any position. We first introduce the notion of saturation for a path.

**Definition 8 (Path Saturation).** *Let  $P$  be a path and  $\ell$  a label. We write  $\text{sat}(P, \ell)$  for the set of paths obtained by saturating  $P$  with respect to  $\ell$ , defined as follows*

$$\text{sat}(P, \ell) = \{Q \mid Q \in \mathbb{P}(\mathcal{E}_{P\ell}, \lambda_{P\ell}) \text{ and } Q|_{\mathcal{E}_P} = P\}$$

A path  $P$  saturated with a label  $\ell$  generates the set of all paths obtained by placing a new event labelled by  $\ell$  in any position within  $P$ . A saturated specification thus extends a computation by adding a new operation that can be arbitrated in any position.

**Definition 9 (Saturated Specification).** *Let  $S$  be a specification. We say that  $S$  is saturated if*

$$\forall \langle G, P \rangle, \ell. P \in S(G^\ell)|_{\mathcal{E}_G} \text{ implies } \text{sat}(P, \ell) \subseteq S(G^\ell)$$

*Example 10.* The specifications in Examples 4, 5 and 6 are all saturated because a new event  $e$  can be arbitrated in any position. In fact, the specifications in Examples 4 and 6 do not use any information about arbitration, while the specification in Example 5 constrains arbitrations only for events that are not maximal. Figure 7 shows a specification that is not saturated because it does not allow to arbitrate the top event (the one labelled  $\langle \text{rd}, 1 \rangle$ ) as the first operation in the path. In a saturated specification, the equality in Fig. 4a should hold. We remark that the specification is coherent although it is not saturated.

$$\begin{array}{l} S(\langle \text{inc}, ok \rangle) = \{ \langle \text{inc}, ok \rangle \} \\ S(\langle \text{rd}, 1 \rangle) = \{ \langle \text{rd}, 1 \rangle \} \end{array} \qquad S\left(\begin{array}{c} \langle \text{inc}, ok \rangle \\ \downarrow \\ \langle \text{rd}, 1 \rangle \end{array}\right) = \left\{ \begin{array}{c} \langle \text{inc}, ok \rangle \\ | \\ \langle \text{rd}, 1 \rangle \end{array} \right\}$$

**Fig. 7.** A non-saturated specification

## 4 Replicated Data Type

In this section we show that our proposal can be considered as (and it is actually more general than) a model for the operational description of RDTs as given in [2,4]. We start by recasting the original definition of RDT (as given in [2, Definition 4.5]) in terms of LDAGs. As hinted in the introduction, the meaning of each operation of an RDT is specified in terms of a context, written  $\mathbb{C}$ , which is a pair  $\langle \mathbf{G}, \mathbf{P} \rangle$  such that  $\mathbf{P} \in \mathbb{P}(\mathcal{E}_{\mathbf{G}}, \lambda_{\mathbf{G}})$ . We write  $\mathbb{C}(\mathcal{L})$  for the set of contexts over  $\mathcal{L}$ , and fix a set  $\mathcal{O}$  of operations and a set  $\mathcal{V}$  of values. Then, the operational description of RDTs in [2,4] can be formulated as follows.

**Definition 10 (Replicated Data Type).** *A Replicated Data Type (RDT) is a function  $\mathcal{F} : \mathcal{O} \times \mathbb{C}(\mathcal{O}) \rightarrow \mathcal{V}$ .*

In words, for any visibility graph  $\mathbf{G}$  and arbitration  $\mathbf{P}$ , the specification  $\mathcal{F}$  indicates the result of executing the operation  $\text{op}$  over  $\mathbf{G}$  and  $\mathbf{P}$ , which is  $\mathcal{F}(\text{op}, \langle \mathbf{G}, \mathbf{P} \rangle)$ .

*Example 11.* The data type **Counter** introduced in Example 4 is formally specified in [2,4] as follows

$$\begin{aligned} \mathcal{F}_{ctr}(\text{inc}, \langle \mathbf{G}, \mathbf{P} \rangle) &= ok \\ \mathcal{F}_{ctr}(\text{rd}, \langle \mathbf{G}, \mathbf{P} \rangle) &= \#\{e \mid e \in \mathbf{G} \text{ and } \lambda(e) = \text{inc}\} \end{aligned}$$

Given a context  $\langle \mathbf{G}, \mathbf{P} \rangle$  in  $\mathbb{C}(\mathcal{O} \times \mathcal{V})$ , we may check whether the value associated with each operation matches the definition of a particular RDT. This notion is known as *return value consistency* [2, Definition 4.8]. In order to relate contexts with and without return values, we use the following notation: given  $\mathbf{G} \in \mathbb{G}(\mathcal{O} \times \mathcal{V})$ , by  $\bar{\mathbf{G}} \in \mathbb{G}(\mathcal{O})$  we denote the LDAG obtained by projecting the labels of  $\mathbf{G}$  in the obvious way.

**Definition 11 (Return Value Consistent).** *Let  $\mathcal{F}$  be an RDT and  $\langle \mathbf{G}, \mathbf{P} \rangle \in \mathbb{C}(\mathcal{O} \times \mathcal{V})$  a context. We say that  $\mathcal{F}$  is Return Value Consistent (RVAL) over  $\mathbf{G}$  and  $\mathbf{P}$  and we write  $\text{RVAL}(\mathcal{F}, \mathbf{G}, \mathbf{P})$  if  $\forall e \in \mathcal{E}_{\mathbf{G}}. \lambda(e) = \langle \text{op}, v \rangle$  implies  $\mathcal{F}(\text{op}, \bar{\mathbf{G}}|_{-\prec_e}, \bar{\mathbf{P}}|_{-\prec_e}) = v$ . Moreover, we define*

$$\text{PRVAL}(\mathcal{F}, \mathbf{G}) = \{\mathbf{P} \mid \text{RVAL}(\mathcal{F}, \mathbf{G}, \mathbf{P})\}$$


*Example 12.* Consider the RDT  $\mathcal{F}_{ctr}$  introduced in Example 11. The context in Fig. 8a is RVAL consistent while the one in Fig. 8b is not because  $\mathcal{F}_{ctr}$  requires **rd** to return the number of **inc** operations seen by that read, which in this case should be 2.

The following result states that return value consistent paths are all coherent, in the sense that they match the behaviour allowed for any shorter configuration.

**Lemma 1.** *Let  $\mathcal{F}$  be an RDT and  $\mathbf{G}$  an LDAG. Then*

$$\text{PRVAL}(\mathcal{F}, \mathbf{G}) = \bigotimes_{e \in \mathcal{E}_{\mathbf{G}}} \text{PRVAL}(\mathcal{F}, \mathbf{G}|_{-\prec_e}).$$

As for coherent specifications, the property  $\text{PRVAL}(\mathcal{F}, \mathbf{G})|_{-\prec_e} \subseteq \text{PRVAL}(\mathcal{F}, \mathbf{G}|_{-\prec_e})$  also holds for return value consistent paths.



**Fig. 8.** RVAL consistency for  $\mathcal{F}_{ctr}$ .

#### 4.1 Deterministic Specifications

We now focus on the relation between our notion of specification, as introduced in Definition 5, and the operational description of RDTs, as introduced in [2,4] and formalised in Definition 10 in terms of LDAGs. Specifically, we characterise a proper subclass of specifications that precisely correspond to RDTs.

For this section we restrict our attention to specifications over the set of labels  $O \times \mathcal{V}$ , i.e.,  $S : \mathbb{G}(O \times \mathcal{V}) \rightarrow 2^{\mathbb{P}(O \times \mathcal{V})}$ .

**Definition 12 (Total Specification).** *Let  $S$  be a specification. We say that  $S$  is total if*

$$\forall \langle \bar{G}, \bar{P} \rangle, \text{op}. \exists G_1, v. \bar{G} = \bar{G}_1 \wedge \bar{P} \in \overline{S(G_1^{\langle \text{op}, v \rangle})} \Big|_{\mathcal{E}_{G_1}}$$

Intuitively, a specification is total when every projection over  $O$  of a context in  $\mathbb{C}(O \times \mathcal{V})$ , as represented by  $\langle \bar{G}, \bar{P} \rangle \in \mathbb{C}(O)$ , can be extended with the execution of any operation of the data type. This is formalised by stating that for any operation  $\text{op}$  and any admissible arbitration (sequence of operations)  $\bar{P}$  of a configuration  $\bar{G}$  (once more, labelled only with operations), then  $\bar{P}$  can be extended into an admissible arbitration of the configuration  $\overline{G_1^{\langle \text{op}, v \rangle}}$ , where  $G_1$  is just one of the possible configurations (the one labelled with the correct return values) whose projection corresponds to  $\bar{G}$ .

We remark that a total specification does not prevent the definition of an operation that admits more than one return value in certain configurations, i.e.,  $v$  in Definition 12 does not need to be unique. For instance, consider the **Generic Register** in Example 6, in which operation  $\text{rd}$  may return any of the causally-independent, previously written values. Albeit being total, the specification for  $\text{rd}$  is not deterministic. On the contrary, a specification is deterministic if an operation executed over a configuration admits at most one return value, as formally stated below.

**Definition 13 (Deterministic Specification).** *Let  $S$  be a specification. We say that  $S$  is deterministic if*

$$\forall G, \text{op}, v, v'. v \neq v' \text{ implies } \overline{S(G^{\langle \text{op}, v \rangle})} \Big|_{\mathcal{E}_G} \cap \overline{S(G^{\langle \text{op}, v' \rangle})} \Big|_{\mathcal{E}_G} = \emptyset$$

A weaker notion for determinism could allow the result for an added operation to depend also on the given admissible path. We say that a specification  $S$  is *value-deterministic* if

$$\forall \mathbf{G}, \text{op}, v, v'. v \neq v' \wedge \mathbf{G} \neq \epsilon \text{ implies } S(\mathbf{G}^{(\text{op}, v)}) \Big|_{\mathcal{E}_{\mathbf{G}}} \cap S(\mathbf{G}^{(\text{op}, v')}) \Big|_{\mathcal{E}_{\mathbf{G}}} = \emptyset$$

Finally, we say that a specification is *functional* if it is both deterministic and total.

*Example 13.* Figure 9 shows a value-deterministic specification. Although a read operation that follows an increment may return two different values, such difference is explained by the previous computation: in one case the increment succeeds while in the other fails. The specification is however not deterministic because it admits a sequence of operations to be decorated with different return values.

$$\begin{array}{cc} S \left( \begin{array}{c} \langle \text{inc}, \text{ok} \rangle \\ \downarrow \\ \langle \text{rd}, 1 \rangle \end{array} \right) = \left\{ \begin{array}{c} \langle \text{inc}, \text{ok} \rangle \\ | \\ \langle \text{rd}, 1 \rangle \end{array} \right\} & S \left( \begin{array}{c} \langle \text{inc}, \text{fail} \rangle \\ \downarrow \\ \langle \text{rd}, \perp \rangle \end{array} \right) = \left\{ \begin{array}{c} \langle \text{inc}, \text{fail} \rangle \\ | \\ \langle \text{rd}, \perp \rangle \end{array} \right\} \\ \text{(a)} & \text{(b)} \end{array}$$

**Fig. 9.** A value-deterministic and coherent specification.

*Example 14.* It is straightforward to check that the specifications in Examples 4 and 5 are deterministic. On the contrary, the specification of the **Generic Register** in Example 6 is not even value-deterministic. It suffices to consider a configuration in which a read operation sees two different written values. Similarly, **Set** in Example 8 is not deterministic.

The lemma below states a simple criterion for determinism.

**Lemma 2.** *Let  $S$  be a coherent and deterministic specification. Then*

$$\forall \mathbf{G}_1, \mathbf{G}_2. \overline{\mathbf{G}_1} = \overline{\mathbf{G}_2} \text{ implies } \mathbf{G}_1 = \mathbf{G}_2 \vee \overline{S(\mathbf{G}_1)} \cap \overline{S(\mathbf{G}_2)} = \emptyset$$

So, if two configurations are annotated with the same operations yet with different values, then their admissible paths are already all different if we disregard return values.

## 4.2 Correspondence Between RDTs and Specifications

This section establishes the connection between RDTs and specifications. We first introduce a mapping from RDTs to specifications.

**Definition 14.** Let  $\mathcal{F}$  be an RDT. We write  $\mathbb{S}(\mathcal{F})$  for the specification associated with  $\mathcal{F}$ , defined as follows

$$\mathbb{S}(\mathcal{F})(\mathbf{G}) = \text{PRVAL}(\mathcal{F}, \mathbf{G})$$

Next result shows that RDTs correspond to specifications that are coherent, functional and saturated.

**Lemma 3.** For every RDT  $\mathcal{F}$ ,  $\mathbb{S}(\mathcal{F})$  is coherent, functional, and saturated.

The inverse mapping from specifications to RDTs is defined below.

**Definition 15.** Let  $S$  be a specification. We write  $\mathbb{F}(S)$  for the RDT associated with  $S$ , defined as follows

$$\mathbb{F}(S)(\text{op}, \bar{\mathbf{G}}, \bar{\mathbf{P}}) = v \text{ if } \exists \mathbf{G}_1. \bar{\mathbf{G}} = \bar{\mathbf{G}}_1 \wedge \bar{\mathbf{P}} \in S(\overline{\mathbf{G}_1^{\langle \text{op}, v \rangle}}) \Big|_{\mathcal{E}_{\mathbf{G}_1}}$$

Note that  $\mathbb{F}(S)$  may not be well-defined for some  $S$ , e.g. when  $S$  is not deterministic. The following lemma states the conditions under which  $\mathbb{F}(S)$  is well-defined.

**Lemma 4.** For every coherent and functional specification  $S$ ,  $\mathbb{F}(S)$  is well-defined.

The following two results show that RDTs are a particular class of specifications, and hence, provide a fully abstract characterisation of operational RDTs.

**Theorem 1.** For every coherent, functional, and saturated specification  $S$ ,  $S = \mathbb{S}(\mathbb{F}(S))$ .

**Theorem 2.** For every RDT  $\mathcal{F}$ ,  $\mathcal{F} = \mathbb{F}(\mathbb{S}(\mathcal{F}))$ .

The above characterisation implies that there are data types that cannot be specified as operational RDTs. Consider e.g. **Generic Register** and **Set**, as introduced respectively in Examples 6 and 8. As noted in Example 14, they are not deterministic. Hence, they cannot be translated as RDTs. We remark that a non-deterministic specification does not imply a non-deterministic conflict resolution, but it allows for underspecification.

## 5 A Categorical Account of Specifications

In the previous sections we provided a functional characterisation of RDTs. We now proceed on to a denotational account of our formalism by providing a categorical foundation which is amenable to the building of a family of operators on specifications.

## 5.1 Composing LDAGs

We start by considering a sub-class of morphisms between LDAGs, which account for the evolution of visibility relation by reflecting the information about observed events.

**Definition 16 (Past-Reflecting Morphism).** *Let  $G_1$  and  $G_2$  be LDAGs and  $f : G_1 \rightarrow G_2$  an LDAG morphism. We say that  $f$  is past-reflecting if*

$$\forall e \in \mathcal{E}_{G_1}. f(- \prec e) = \bigcup_{e' \in - \prec f(e)} e'.$$

We can concisely write  $f(- \prec e) = - \prec f(e)$  and spell out the definition as

$$\forall e \in \mathcal{E}_{G_1}. \forall e_2 \in G_2|_{- \prec f(e)}. \exists e_1 \in G_1|_{- \prec e}. f(e_1) = e_2$$

It is noteworthy that this requirement boils down to (the dual of) what are called *tp-morphisms* in the literature on algebraic specification theory, which are an instance of *open maps* [9]. As we will see, this property is going to be fundamental in obtaining a categorical characterisation of coherent specifications.

Now, let  $\mathcal{G}(\mathcal{L})$  be the category whose objects are LDAGs and arrows are past-reflecting morphisms, and  $I(\mathcal{L})$  the sub-category whose arrows are injective morphisms.

**Proposition 1 (LDAG Pullbacks/Pushouts).** *The category  $\mathcal{G}(\mathcal{L})$  of LDAGs and past-reflecting morphisms has (strict) initial object, pullbacks and pushouts along monos.*

Note that pushout squares along monos are also pullback ones. As often the case, the property concerning pushouts does not hold in  $I(\mathcal{L})$ , even if a weak form does, since monos are stable under pushouts in  $\mathcal{G}(\mathcal{L})$ . For the time being, we just remark that these properties guarantee a degree of modularity for our formalism.

We need a last definition before giving a categorical presentation.

**Definition 17 (Downward closure).** *Let  $G = \langle \mathcal{E}, \prec, \lambda \rangle$  be an LDAG and  $\mathcal{E}' \subseteq \mathcal{E}$ . We say that  $\mathcal{E}'$  is downward closed if*

$$\forall e \in \mathcal{E}'. - \prec e \subseteq \mathcal{E}'.$$

It is easy to show that for any past-reflecting morphism  $f : G_1 \rightarrow G_2$  the image of  $\mathcal{E}_{G_1}$  along  $f$  is downward closed. Should  $f$  be injective, we strengthen the relationship.

**Lemma 5.** *An injective morphism  $f : G_1 \rightarrow G_2$  is past-reflecting if and only if*

1.  $f(e_1) \prec_{G_2} f(e_2)$  implies  $e_1 \prec_{G_1} e_2$ ;
2.  $\bigcup_{e \in \mathcal{E}_{G_1}} f(e)$  is downward closed.

This result tells us that past-reflecting injective morphisms  $\mathbf{f} : \mathbf{G}_1 \rightarrow \mathbf{G}_2$  are uniquely characterised as such by the properties of the image of  $\mathcal{E}_1$  with respect to  $\mathbf{G}_2$ .

Now, while the initial object of both  $\mathcal{G}(\mathcal{L})$  and  $I(\mathcal{L})$  is the empty graph  $\epsilon$ , the pullback in the latter has an easy characterisation, thanks to the previous lemma. Indeed, let  $\mathbf{f}_i : \mathbf{G}_i \rightarrow \mathbf{G}$  be past-preserving injective morphisms, assuming the functions on elements to be identities for the sake of simplicity, and let  $\mathcal{E} = \mathcal{E}_{\mathbf{G}_1} \cap \mathcal{E}_{\mathbf{G}_2}$ . Then,  $\mathbf{G}_1|_{\mathcal{E}} = \mathbf{G}_2|_{\mathcal{E}}$  and they correspond (with the obvious morphisms) to the pullback of  $\mathbf{f}_1$  and  $\mathbf{f}_2$ .

## 5.2 The Model Category

We now move to define the model category.

**Definition 18 (Morphism Saturation).** *Let  $\mathbb{P}(\mathcal{E}_1, \lambda_1)$  and  $\mathbb{P}(\mathcal{E}_2, \lambda_2)$  be sets of paths and  $\mathbf{f} : \mathcal{E}_1 \rightarrow \mathcal{E}_2$  an injective function such that  $\lambda_1 = \mathbf{f}; \lambda_2$ . The saturation function  $\mathbf{sat}(-, \mathbf{f})$  is defined as follows*

$$\mathbf{sat}(\mathbf{P}, \mathbf{f}) = \{\mathbf{Q} \mid \mathbf{Q} \in \mathbb{P}(\mathcal{E}_2, \lambda_2) \text{ and } \mathbf{P} = \mathbf{Q}|_{\mathbf{f}(\mathcal{E}_1)}\}$$

That is, each  $\mathbf{Q}$  is the image of  $\mathbf{P}$  via a morphism with underlying function  $\mathbf{f}$ . We can exploit saturation in order to get a simple definition of our model category.

**Definition 19 (Path-Set 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  $f : \mathcal{X}_1 \rightarrow \mathcal{X}_2$  is an injective function  $f : \mathcal{E}_1 \rightarrow \mathcal{E}_2$  such that  $\lambda_1 = f; \lambda_2$  and*

$$\mathcal{X}_2 \subseteq \bigcup_{\mathbf{P} \in \mathcal{X}_1} \mathbf{sat}(\mathbf{P}, \mathbf{f})$$

The property can be stated as

$$\forall \mathbf{P}_2 \in \mathcal{X}_2. \exists \mathbf{P}_1 \in \mathcal{X}_1. \mathbf{P}_2 \in \mathbf{sat}(\mathbf{P}_1, \mathbf{f})$$

thus each path in  $\mathbf{P}_2$  is related to a (unique) path in  $\mathbf{P}_1$  via a morphism induced by  $\mathbf{f}$ . Let  $\mathcal{P}(\mathcal{L})$  be the category whose objects are sets of paths over the same elements and labelling (i.e., subsets of  $\mathbb{P}(\mathcal{E}, \lambda)$  for some  $\mathcal{E}$  and  $\lambda$ ), and arrows are path-set morphisms.

**Proposition 2 (Path Pullbacks/Pushouts).** *The category  $\mathcal{P}(\mathcal{L})$  of sets of paths and path-set morphisms has (strict) initial object and pullbacks.*

As for  $I(\mathcal{L})$ , also  $\mathcal{P}(\mathcal{L})$  admits a weak form of pushouts along monos.

*Remark 1.* The initial object is the set in  $2^{\mathbb{P}(\emptyset, \lambda_\emptyset)}$  including only the empty path  $\epsilon$ . As for pullbacks, let  $\mathbf{f}_i : \mathcal{X}_i \rightarrow \mathcal{X}$  be path-set morphisms, assuming the functions on elements to be identities for the sake of simplicity, and let  $\mathcal{E} = \mathcal{E}_1 \cap \mathcal{E}_2$ . Then, the pullback is the set  $\mathcal{X}_1|_{\mathcal{E}} \cup \mathcal{X}_2|_{\mathcal{E}}$  in  $2^{\mathbb{P}(\mathcal{E}, \lambda)}$  with  $\lambda = \lambda_1|_{\mathcal{E}} = \lambda_2|_{\mathcal{E}}$ . As for pushouts, let  $\mathbf{f} : \mathcal{X} \rightarrow \mathcal{X}_i$  be injective path-set morphisms, assuming the functions on elements to be identities for the sake of simplicity, and  $\mathcal{E} = \mathcal{E}_1 \cup \mathcal{E}_2$ . Then, the “weak” pushout is the set  $\mathcal{X}_1 \otimes \mathcal{X}_2$  in  $2^{\mathbb{P}(\mathcal{E}, \lambda)}$  with  $\lambda$  the extension of  $\lambda_1$  and  $\lambda_2$ .



### 5.3 A Categorical Correspondence

It is now time to move towards our categorical characterisation of specifications.

First, let us restrict our attention to functors  $F : I(\mathcal{L}) \rightarrow \mathcal{P}(\mathcal{L})$  that preserve the underlying set of objects, i.e., such that the underlying function on objects  $Ob_F$  maps an LDAG  $\mathbf{G}$  into a subset of  $2^{\mathbb{P}(\mathcal{E}_{\mathbf{G}}, \lambda_{\mathbf{G}})}$  (and preserves the underlying function on path-set morphisms). We also say that  $F$  is coherent if  $F(\mathbf{G}) = \bigotimes_{e \in \mathcal{E}_{\mathbf{G}}} F(\mathbf{G}|_{-\preceq e})$  for all LDAGS  $\mathbf{G}$ . Thus, any such functor  $F$  that preserves the initial object (i.e.,  $F(\epsilon) = \{\epsilon\}$ ) gives raise to a specification: it just suffices to consider the object function  $Ob_F : Ob_{I(\mathcal{L})} \rightarrow Ob_{\mathcal{P}(\mathcal{L})}$ .

**Proposition 3.** *Let  $F : I(\mathcal{L}) \rightarrow \mathcal{P}(\mathcal{L})$  be a (coherent) functor preserving the initial object. Then  $Ob_F$  is a (coherent) specification.*

For the inverse we need an additional lemma.

**Lemma 6.** *Let  $\mathcal{S}$  be a coherent specification and  $\mathcal{E} \subseteq \mathcal{E}_{\mathbf{G}}$  downward closed. Then  $\mathcal{S}(\mathbf{G})|_{\mathcal{E}} \subseteq \mathcal{S}(\mathbf{G}|_{\mathcal{E}})$ .*

The lemma above immediately implies the following result.

**Proposition 4.** *A coherent specification  $\mathcal{S}$  induces a coherent functor  $\mathbb{M}(\mathcal{S}) : I(\mathcal{L}) \rightarrow \mathcal{P}(\mathcal{L})$  preserving the initial object such that  $Ob_{\mathbb{M}(\mathcal{S})} = \mathcal{S}$ .*

By using Propositions 3 and 4 we can state the main result of this section.

**Theorem 3.** *There is a bijection between coherent specifications and coherent functors  $I(\mathcal{L}) \rightarrow \mathcal{P}(\mathcal{L})$  preserving the initial object.*

## 6 Conclusions and Future Works

Our contribution proposes a denotational view of replicated data types. While most of the traditional approaches are operational in flavour [4, 7, 8], we strived for a formalism for specifications which could exploit the classical tools of algebraic specification theory. More precisely, we associate to each configuration (i.e., visibility) a set of admissible arbitrations. Differently from those previous approaches, our presentation naturally accommodates non-deterministic specifications and enables abstract definitions allowing for different strategies in conflict resolution. Our formulation brings into light some properties held by mainstream specification formalisms: beside the obvious property of functionality, they also satisfy coherence and saturation. A coherent specification can neither prescribe an arbitration order between events that are unrelated by visibility nor allow for additional arbitrations over past events when a configuration is extended (i.e., a new top element is added to visibility). Instead, a saturated specification cannot impose any constraint to the arbitration of top elements. Note that saturation does not hold when requiring that admissible arbitrations should be also topological orderings of visibility. Hence, the approaches in [2, 4] generate

specifications that are not saturated. We remark that this relation between visibility and arbitration translates in a quite different property in our setting, and this suggests that consistency models defined as relations between visibility and arbitration (e.g., *monotonic* and *causal* consistency) could have alternative characterisations. We plan to explore these connections in future works.

Another question concerns coherence, which prevents a specification from choosing an arbitration order on events that are unrelated by visibility and forbids, e.g., the definition of strategies that arbitrate first the events coming from a particular replica. Consequently, it becomes natural to look for those RDTs and consistency models that are the counterpart of non-coherent specifications, still preserving some suitable notion of causality between events. We do believe that the weaker property  $\mathcal{S}(\mathbf{G})|_{\rightarrow e} \subseteq \mathcal{S}(\mathbf{G}|_{\rightarrow e})$  (that is, no additional arbitration over past events when a configuration is extended) is a worthwhile alternative, accommodating for many examples that impose less restrictions on the set of admissible paths (hence, that may allow more freedom to the arbitration).

These issues might be further clarified by our categorical presentation. Our proposal is inspired by current work on the semantics of nominal calculi [11], and it shares similarities with [10], since our category  $\mathcal{G}$  is the sub-category of their  $\mathbf{FinSet}^{\rightarrow}$  with past-reflecting morphisms. The results on Sect. 5 focus on a functorial characterisation of specifications. We chose an easy way out for establishing the bijection between functors and specifications by restricting the possible object functions and by defining coherence “on the nose”, (i.e., by considering functors  $F$  such that  $F(\mathbf{G}) \subseteq 2^{\mathbb{P}(\mathcal{E}_{\mathbf{G}}, \lambda_{\mathbf{G}})}$  and  $F(\mathbf{G}) = \bigotimes_{e \in \mathcal{E}_{\mathbf{G}}} F(\mathbf{G}|_{\rightarrow e})$ ), since requiring the specification to be coherent is needed in order to obtain the functor in Proposition 4. A proper characterisation should depend on the properties of  $F$  over the arrows of  $\mathcal{G}$  (such as pullback/pushout preservation), instead of the properties of the objects in its image on  $\mathcal{P}$ .

The same categorical presentation may shed light on suitable operators on specifications. Indeed, this is the usual situation when providing a functorial semantics for a language (see e.g. [1], and the references therein, among many others), and intuitively we have already a freshness operator  $F^\ell(\mathbf{G}) = F(\mathbf{G}^\ell)$ , along the lines of edge allocation in [10]. We plan to extend these remarks into a full-fledged algebra for specifications.

## References

1. Bonchi, F., Buscemi, M.G., Ciancia, V., Gadducci, F.: A presheaf environment for the explicit fusion calculus. *J. Autom. Reasoning* **49**(2), 161–183 (2012)
2. Burckhardt, S.: Principles of eventual consistency. *Found. Trends Program. Lang.* **1**(1–2), 1–150 (2014)
3. Burckhardt, S., Gotsman, A., Yang, H.: Understanding eventual consistency. Technical Report MSR-TR-2013-39, Microsoft Research (2013)
4. 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)

5. 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. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2015)
6. Gilbert, S., Lynch, N.: Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services. SIGACT News **33**(2), 51–59 (2002)
7. Gotsman, A., Yang, H.: Composite replicated data types. In: Vitek, J. (ed.) ESOP 2015. LNCS, vol. 9032, pp. 585–609. Springer, Heidelberg (2015). doi:[10.1007/978-3-662-46669-8\\_24](https://doi.org/10.1007/978-3-662-46669-8_24)
8. Gotsman, A., Yang, H., Ferreira, C., Najafzadeh, M., Shapiro, M.: 'Cause i'm strong enough: reasoning about consistency choices in distributed systems. In: Bodík, R., Majumdar, R. (eds.) POPL 2016, pp. 371–384. ACM (2016)
9. Joyal, A., Nielson, M., Winskel, G.: Bisimulation and open maps. In: LICS 1993, pp. 418–427. IEEE (1993)
10. Montanari, U., Sammartino, M.: A network-conscious  $\pi$ -calculus and its coalgebraic semantics. Theor. Comput. Sci. **546**, 188–224 (2014)
11. Pitts, A.M.: Nominal Sets: Names and Symmetry in Computer Science. Cambridge University Press, Cambridge (2013)
12. 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, Heidelberg (2011). doi:[10.1007/978-3-642-24550-3\\_29](https://doi.org/10.1007/978-3-642-24550-3_29)
13. Shapiro, M., Preguiça, N., Baquero, C., Zawirski, M.: A comprehensive study of convergent and commutative replicated data types. Technical Report RR-7506, Inria-Centre Paris-Rocquencourt (2011)
14. 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)