

A Categorical Account of Replicated Data Types

Fabio Gadducci 

Dipartimento di Informatica, Università di Pisa, Italia
<https://www.di.unipi.it/~gadducci>
fabio.gadducci@unipi.it

Hernán Melgratti 

Departamento de Computación, Universidad de Buenos Aires, Argentina
ICC-CONICET-UBA, Buenos Aires, Argentina
<https://lafhis.dc.uba.ar/~melgratti>
hmelgra@dc.uba.ar

Christian Roldán 

Departamento de Computación, Universidad de Buenos Aires, Argentina
<https://lafhis.dc.uba.ar/~croidan>
croidan@dc.uba.ar

Matteo Sammartino 

Department of Computer Science, University College London, UK
<https://matteosammartino.com>
m.sammartino@ucl.ac.uk

Abstract

Replicated Data Types (RDTs) have been introduced as a suitable abstraction for dealing with weakly consistent data stores, which may (temporarily) expose multiple, inconsistent views of their state. In the literature, RDTs are commonly specified in terms of two relations: visibility, which accounts for the different views that a store may have, and arbitration, which states the logical order imposed on the operations executed over the store. Different flavours, e.g., operational, axiomatic and functional, have recently been proposed for the specification of RDTs. In this work, we propose an algebraic characterisation of RDT specifications. We define categories of visibility relations and arbitrations, show the existence of relevant limits and colimits, and characterize RDT specifications as functors between such categories that preserve these additional structures.

2012 ACM Subject Classification Theory of computation → Program semantics; Software and its engineering → General programming languages

Keywords and phrases Replicated data type, Specification, Functorial characterisation

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2019.42

Funding Research partially supported by the MIUR PRIN 2017FTXR7S “IT-MaTTeR’S”, by the EU H2020 RISE programme under the Marie Skłodowska-Curie grant agreement 778233, by the UBACyT projects 20020170100544BA and 20020170100086BA, by the PIP project 11220130100148CO, by the Leverhulme Prize PLP-2016-129, and by the EPSRC grant EP/S028641/1.

Acknowledgements We thank the reviewers for their careful reading and insightful comments.

1 Introduction

The CAP theorem establishes that a distributed data store can simultaneously provide two of the following three properties: consistency, availability, and tolerance to network partitions [8]. A weakly consistent data store prioritises availability and partition tolerance over consistency. As a consequence, a weakly consistent data store may (temporarily) expose multiple, inconsistent views of its state; hence, the behaviour of operations may depend on the particular view over which they are executed. Replicated data types (RDTs) have been proposed as suitable data type abstractions for weakly consistent data stores. The specification of such data types usually takes into account the particular views over which



© Fabio Gadducci, Hernán Melgratti, Christian Roldán, and Matteo Sammartino;
licensed under Creative Commons License CC-BY

39th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2019).

Editors: Arkadev Chattopadhyay and Paul Gastin; Article No. 42; pp. 42:1–42:15



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

order in \mathcal{X}_1 . In this work we show that a large class of specifications, dubbed *coherent*, can be characterised functorially. Roughly, a coherent specification accounts for those RDTs such that the arbitrations associated with a visibility relation can be obtained by extending arbitrations associated with “smaller” visibilities: as illustrated in [6], they correspond to what are called *return value consistent* RDTs in [4]. We establish a bijection between functors and specifications, showing that a coherent specification induces a functor from $\mathbf{PIDag}(\mathcal{L})$ into $\mathbf{SPath}(\mathcal{L})$ that preserves colimits and binary pullbacks and vice versa.

The paper has the following structure. Section 2 offers some preliminaries on categories of relations, which are used for proposing some basic results on categories of graphs and paths in Section 3. Section 4 recalls the set-theoretical presentation of RDTs introduced in [6]. Section 5 introduces our semantical model, the category of set of paths, describing some of its basic properties with respect to limits and colimits. In Section 6 we present some categorical operators for RDTs, which are used in Section 7 to present our main characterisation results. The paper is closed with some final remarks, a comparison of the proposed constructions with those presented in [7], and some hints towards future work.

2 Preliminaries on Relations

Relations. Given a finite set E , a (binary) *relation* ρ over E is a subset $\rho \subseteq E \times E$ of the cartesian product of E with itself. We use the pair $\langle E, \rho \rangle$ to denote a relation ρ over E , in order to always have the set of events explicit, and simply \emptyset to denote the empty relation.

A subset $E' \subseteq E$ is *downward closed* with respect to ρ if $\forall e \in E, e' \in E'. e \rho e'$ implies $e \in E'$ and, when ρ is clear, we write $\lfloor e \rfloor$ for the smallest downward closed set including $e \in E$.

► **Definition 1** ((Binary Relation) Morphisms). *A (binary relation) morphism $f : \langle E, \rho \rangle \rightarrow \langle T, \gamma \rangle$ is a function $f : E \rightarrow T$ such that*

$$\forall e, e' \in E. e \rho e' \text{ implies } f(e) \gamma f(e')$$

A morphism $f : \langle E, \rho \rangle \rightarrow \langle T, \gamma \rangle$ is past-reflecting (shortly, pr-morphism) if

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

Note that both classes of morphisms are closed under composition: we denote as \mathbf{Bin} the category of relations and their morphisms and \mathbf{PBin} the sub-category of pr-morphisms.

► **Lemma 2** (Characterising pr-morphisms). *Let $f : \langle E, \rho \rangle \rightarrow \langle T, \gamma \rangle$ be a morphism. If*

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

then it is a pr-morphism. If f is injective, then the converse holds.

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

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

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

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

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

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

42:4 A Categorical Account of Replicated Data Types

Clearly, **Bin** has both finite limits and finite colimits, which are computed point-wise as in **Set**. The structure is largely lifted to **PBin**.

► **Proposition 3** (Properties of **PBin**). *The inclusion functor $\mathbf{PBin} \rightarrow \mathbf{Bin}$ reflects finite colimits and binary pullbacks.*

In other words, since **Bin** has finite limits and finite colimits, finite colimits and binary pullbacks in **PBin** always exist and are computed as in **Bin**. There is e.g. no terminal object, since morphisms in **Bin** into the singleton are clearly not past-reflecting.

Monos in **Bin** are just morphisms whose underlying function is injective, and similarly in **PBin**, so that the inclusion functor preserves (and reflects) them.

► **Lemma 4** (Monos under pushouts). *Pushouts in **Bin** (and thus in **PBin**) preserve monos.*

We now introduce labelled relations. Consider the forgetful functors $U_r : \mathbf{Bin} \rightarrow \mathbf{Set}$ and $U_p : \mathbf{PBin} \rightarrow \mathbf{Set}$, the latter factoring through the inclusion functor $\mathbf{PBin} \rightarrow \mathbf{Bin}$. Given a set \mathcal{L} of labels, we consider the comma categories $\mathbf{Bin}(\mathcal{L}) = U_r \downarrow \mathcal{L}$ and $\mathbf{PBin}(\mathcal{L}) = U_p \downarrow \mathcal{L}$: finite colimits and binary pullbacks always exist and are essentially computed as in **Bin**.

Explicitly, an object in $U_r \downarrow \mathcal{L}$ is a triple (E, ρ, λ) for a labeling function $\lambda : E \rightarrow \mathcal{L}$. A label-preserving morphism $(E, \rho, \lambda) \rightarrow (E', \rho', \lambda')$ is a morphism $f : (E, \rho) \rightarrow (E', \rho')$ such that $\forall s \in E. \lambda(s) = \lambda'(f(s))$. Moreover, finite colimits and binary pullbacks exist and are computed as in **Bin**. Similar properties hold for the objects and the morphisms of $U_p \downarrow \mathcal{L}$.

3 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 5** (**PDag**). ***PDag** is the full sub-category of **PBin** whose objects are directed acyclic graphs.*

In other terms, objects are relations whose transitive closures are *strict* partial orders.

► **Remark 6**. The full sub-category of **Bin** whose objects are directed acyclic graphs is not suited for our purposes, since e.g. it does not admit pushouts, not even along monos. The one with pr-morphisms is much more so, still remaining computationally simple.

► **Proposition 7** (Properties of **PDag**). *The inclusion functor $\mathbf{PDag} \rightarrow \mathbf{PBin}$ reflects finite colimits and binary pullbacks.*

We now move to consider *paths*, i.e., relations that are total orders.

► **Definition 8** (**Path**). ***Path** is the full sub-category of **Bin** whose objects are paths.*

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.

► **Proposition 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 $U_{rp} : \mathbf{Path} \rightarrow \mathbf{Set}$ and $U_{pd} : \mathbf{PDag} \rightarrow \mathbf{Set}$: for a set of labels \mathcal{L} we denote $\mathbf{PDag}(\mathcal{L}) = U_{rp} \downarrow \mathcal{L}$ and $\mathbf{Path}(\mathcal{L}) = U_{pd} \downarrow \mathcal{L}$. Once more, finite colimits and binary pullbacks always exist and are essentially computed as in **Bin**.

4 Replicated Data Type Specification

We briefly recall the set-theoretical model of replicated data types (RDT) introduced in [6]. Our main result is its categorical characterisation, which is given in the following sections.

First, some notation. We denote a graph as the triple $\langle \mathcal{E}, \prec, \lambda \rangle$ and a path as the triple $\langle \mathcal{E}, \leq, \lambda \rangle$, in order to distinguish them. Moreover, given a graph $\mathbf{G} = \langle \mathcal{E}, \prec, \lambda \rangle$ and a subset $\mathcal{E}' \subseteq \mathcal{E}$, we denote by $\mathbf{G}|_{\mathcal{E}'}$ the obvious restriction (and the same for a path \mathbf{P}).

We now define a product operation on a set of paths $\mathcal{X} = \{\langle \mathcal{E}_i, \leq_i, \lambda_i \rangle\}_i$. First, we say that the paths of a set \mathcal{X} are *compatible* if $\forall e, i, j. e \in \mathcal{E}_i \cap \mathcal{E}_j$ implies $\lambda_i(e) = \lambda_j(e)$.

► **Definition 10** (Product). *Let \mathcal{X} be a set of compatible paths. The product of \mathcal{X} is*

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

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. A set of sets of paths $\mathcal{X}_1, \mathcal{X}_2, \dots$ is compatible if $\bigcup_i \mathcal{X}_i$ is so. In such case we can define the product $\bigotimes_i \mathcal{X}_i$ as $\bigotimes \bigcup_i \mathcal{X}_i$.

Now, let us further denote with $\mathbb{G}(\mathcal{L})$ and $\mathbb{P}(\mathcal{L})$ the sets of (finite) graphs and (finite) paths, respectively, labelled over \mathcal{L} and with ϵ the empty graph. Also, when the set of labels \mathcal{L} is chosen, we let $\mathbb{G}(\mathcal{E}, \lambda)$ and $\mathbb{P}(\mathcal{E}, \lambda)$ the sets of graphs and paths, respectively, whose elements are those in \mathcal{E} and are labelled by $\lambda : \mathcal{E} \rightarrow \mathcal{L}$.

► **Definition 11** (Specifications). *A specification \mathcal{S} is a function $\mathcal{S} : \mathbb{G}(\mathcal{L}) \rightarrow 2^{\mathbb{P}(\mathcal{L})}$ such that $\mathcal{S}(\epsilon) = \{\epsilon\}$ and $\forall \mathbf{G}. \mathcal{S}(\mathbf{G}) \in 2^{\mathbb{P}(\mathcal{E}_{\mathbf{G}}, \lambda_{\mathbf{G}})}$.*

In other words, a specification \mathcal{S} maps a graph (interpreted in terms of the visibility relation of a RDT) to a set of paths (that is, the admissible arbitrations of the RDT). Indeed, note that $\mathbf{P} \in \mathcal{S}(\mathbf{G})$ is a path over $\mathcal{E}_{\mathbf{G}}$, hence a total order of the events in \mathbf{G} .

As shown in [6], Definition 11 offers an alternative characterisation of RDTs [4] for a suitable choice of the set of labels. In particular, an RDT boils down to a specification labelled over pairs $\langle \text{operation}, \text{value} \rangle$ that is *saturated* and *past-coherent*. The former property is a technical one: roughly, if \mathbf{G}' is an extension of \mathbf{G} with a fresh event e , then the admissible arbitrations that a saturated specification \mathcal{S} assigns to \mathbf{G}' (i.e., the set of paths $\mathcal{S}(\mathbf{G}')$) are included in the admissible arbitrations of \mathbf{G} saturated with respect to e , i.e., all the paths that extends a path in $\mathcal{S}(\mathbf{G})$ with e inserted at an arbitrary position. Coherence instead is fundamental and expresses that admissible arbitrations of a visibility graph can be obtained by composing the admissible arbitrations of smaller visibilities.

► **Definition 12** ((Past-)Coherent Specification). *Let \mathcal{S} be a specification. We say that \mathcal{S} is past-coherent (briefly, coherent) if*

$$\forall \mathbf{G} \neq \epsilon. \mathcal{S}(\mathbf{G}) = \bigotimes_{e \in \mathcal{E}_{\mathbf{G}}} \mathcal{S}(\mathbf{G}|_{\{e\}}).$$

Explicitly, 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 associated with its sub-graphs $\mathbf{G}|_{\{e\}}$.

Next example illustrates a coherent specification for the `Register` RDT.

► **Example 13** (Register). Fix the set of labels $\mathcal{L} = \{\langle \text{wr}(k), \text{ok} \rangle, \langle \text{rd}, k \rangle \mid k \in \mathbb{N}\} \cup \{\langle \text{rd}, \perp \rangle\}$. Then, the specification of the RDT **Register** is given by the function \mathcal{S}_{lwwR} defined as

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

Intuitively, a visibility graph \mathbf{G} is mapped to a non-empty set of arbitrations (i.e., $\mathcal{S}_{lwwR}(\mathbf{G}) \neq \emptyset$) only when each event e in \mathbf{G} associated with a read operation has a return value k that matches the value written by the greatest event e' (according to \prec_P). The result of a read is undefined (i.e., \perp) when it does not see any write (first condition).

5 The model category

In order to provide a categorical characterisation of coherent specifications, we must first define precisely the model category. So far, we know that its objects have to be sets of compatible paths. We fix a set of labels \mathcal{L} , and we first look at a free construction for paths, and then we turn our attention to morphisms.

5.1 Saturation

► **Definition 14** (Path saturation). Let P be a path and $f : (\mathcal{E}_P, \lambda_P) \rightarrow (\mathcal{E}, \lambda)$ a function preserving labels. The saturation of P along f is defined as

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

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

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

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

► **Definition 16** (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

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

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 λ is fully characterised as the restriction of λ_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 17.** *Let $\mathcal{X}_1 \subseteq \mathbb{P}(\mathcal{E}_1, \lambda_1)$ be a set of paths and $\mathbf{f} : (\mathcal{E}_1, \lambda_1) \rightarrow (\mathcal{E}_2, \lambda_2)$ a function preserving labels. Then $\mathcal{X}_1 \subseteq \text{ret}(\text{sat}(\mathcal{X}_1, \mathbf{f}), \mathbf{f})$. If \mathbf{f} is injective, then the equality holds.*

► **Lemma 18.** *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 $\mathcal{X}_2 \subseteq \text{sat}(\text{ret}(\mathcal{X}_2, \mathbf{f}), \mathbf{f})$.*

We say that an injective function \mathbf{f} is *saturated* with respect to \mathcal{X}_2 if the equality holds.

► **Example 19.** Consider the set of paths \mathcal{X}_1 and \mathcal{X}_2 and the pr-morphism \mathbf{f} below

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

the underlying function \mathbf{f} (defined in Example 15) is *not* saturated with respect to \mathcal{X}_2 because

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

In fact, the ps-morphism $\mathbf{f} : \mathcal{X}_1 \rightarrow \mathcal{X}_2$ only adds the new event $\langle \text{rd}, 2 \rangle$ on top of the path in \mathcal{X}_1 , thus making it a *topological* ps-morphism (see Section 7.3 later on).

5.2 From saturation to categories

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

► **Definition 20** (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 function $\mathbf{f} : (\mathcal{E}_1, \lambda_1) \rightarrow (\mathcal{E}_2, \lambda_2)$ preserving labels such that $\mathcal{X}_2 \subseteq \text{sat}(\mathcal{X}_1, \mathbf{f})$.*

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.

► **Example 21.** Consider the following three sets and the function \mathbf{f} from Example 15

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

Now, \mathbf{f} induces a ps-morphism $\mathbf{f} : \mathcal{X}_1 \rightarrow \mathcal{X}_2$ because $\mathcal{X}_2 \subseteq \text{sat}(\mathcal{X}_1, \mathbf{f})$ (the latter is shown in Example 15). 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 $\langle \text{rd}, 2 \rangle$.

► **Definition 22** (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 23** (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. We have a 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: 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 $\mathbf{f}'_i : \mathcal{E}_1 \times_{\mathcal{E}} \mathcal{E}_2 \rightarrow \mathcal{E}_i$ 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 $\mathbf{SPath}(\mathcal{L})$. ◀

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.

► **Lemma 24.** *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. ◀

► **Example 25.** Consider the following ps-morphisms

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

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

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

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.

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

6 Structure and Operators for Visibility

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 \mathbf{G} is *rooted* if there exists a (necessarily unique) event $e \in \mathcal{E}_{\mathbf{G}}$ such that $\mathbf{G} = \mathbf{G}|_{\{e\}}$.

► **Definition 26** (Extension). *Let $\mathbf{G} = \langle \mathcal{E}, \prec, \lambda \rangle$ and $\mathcal{E}' \subseteq \mathcal{E}$. We define the extension of \mathbf{G} over \mathcal{E}' with ℓ as the graph $\mathbf{G}_{\mathcal{E}'}^{\ell} = \langle \mathcal{E}_{\top}, \prec \cup (\mathcal{E}' \times \{\top\}), \lambda[\top \mapsto \ell] \rangle$.*

Here, \mathcal{E}_{\top} denotes the extension of the set \mathcal{E} with a new event \top , labelled into ℓ . Intuitively, $\mathbf{G}_{\mathcal{E}'}^{\ell}$ is obtained by adding to the visibility relation \mathbf{G} a new event “seeing” some events in \mathcal{E}' . We call the inclusion $\mathbf{G} \rightarrow \mathbf{G}_{\mathcal{E}'}^{\ell}$ an *extension morphism*. Should $\mathbf{G}_{\mathcal{E}'}^{\ell}$ be rooted, we call it a *root extension* of \mathbf{G} , and the associated inclusion a *root extension morphism*.

► **Proposition 27.** *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 $\mathbf{f} : \mathbf{G}_1|_{[e]} \rightarrow \mathbf{G}_1$. ◀

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 direct acyclic graphs and monic pr-morphisms. Note that the chosen morphism \mathbf{f} in the proof of Proposition 27 is mono, since morphisms in $\mathbf{PDag}(\mathcal{L})$ are 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 $\mathbf{PIDag}(\mathcal{L})$.

We first need a technical lemma.

► **Lemma 28** (Monos under pushouts, 2). *Pushouts in $\mathbf{PDag}(\mathcal{L})$ preserve monos.*

We can then state an important characterisation of $\mathbf{PIDag}(\mathcal{L})$.

► **Proposition 29.** *$\mathbf{PIDag}(\mathcal{L})$ 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 also 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 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 $\mathbf{G}_2|_{[e_1]} \cap \mathbf{G}_2|_{[e_2]} \rightarrow \mathbf{G}_2|_{[e_1]}$. ◀

7 A categorical correspondence

It is now the time for moving towards our categorical characterisation of specifications. In this section we will show that coherent specifications induce functors preserving the relevant categorical structure (soundness) and, conversely, that a certain class of functors (basically, those preserving finite colimits and binary pullbacks) induce coherent specifications (completeness). Finally, we will prove that these functions between functors and specifications are mutually inverse, establishing a one-to-one correspondence (up-to isomorphism).

We first provide a simple technical result for coherent specifications.

► **Lemma 30.** *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 $\mathcal{S}(\mathbf{G})|_{\mathcal{E}} = (\bigotimes_{e \in \mathcal{E}_{\mathbf{G}}} \mathcal{S}(\mathbf{G}|_{[e]}))|_{\mathcal{E}}$ and $\mathcal{S}(\mathbf{G}|_{\mathcal{E}}) = \bigotimes_{e \in \mathcal{E}} \mathcal{S}(\mathbf{G}|_{[e]})$. Note that $(\bigotimes_{e \in \mathcal{E}_{\mathbf{G}}} \mathcal{S}(\mathbf{G}|_{[e]}))|_{\mathcal{E}} \subseteq \bigotimes_{e \in \mathcal{E}} \mathcal{S}(\mathbf{G}|_{[e]})$ 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). ◀

7.1 Soundness

The notion of specification introduced in Definition 11 is oblivious to the existence of morphisms between graphs. In the following we impose a minimal consistency requirement, i.e., that a specification maps isomorphic graphs to isomorphic sets of paths, along the same isomorphism on events. That is, if there exists an isomorphism in **PDag** from \mathbf{G}_1 to \mathbf{G}_2 with underlying bijection $\mathbf{f} : \mathcal{E}_{\mathbf{G}_1} \rightarrow \mathcal{E}_{\mathbf{G}_2}$, then for all specifications \mathcal{S} there is an isomorphism in **SPath**(\mathcal{L}) from $\mathcal{S}(\mathbf{G}_1)$ to $\mathcal{S}(\mathbf{G}_2)$ with the same underlying function.

► **Proposition 31** (functors induced by specifications). *A coherent specification \mathcal{S} induces a 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 $\mathbf{f} : (\mathcal{E}_{\mathbf{G}}, \lambda_{\mathbf{G}}) \hookrightarrow (\mathcal{E}_{\mathbf{G}'}, \lambda_{\mathbf{G}'})$. The proof boils down to showing that \mathbf{f} really is a 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 30 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 $P \in \mathcal{S}(\mathbf{G}')$. Since $P|_{\mathcal{E}_{\mathbf{G}}} \in \mathcal{S}(\mathbf{G})$, we have $P \in \mathbf{sat}(\mathcal{S}(\mathbf{G}), \mathbf{f})$, because saturation adds missing events – namely those in $\mathcal{E}_{\mathbf{G}'} \setminus \mathcal{E}_{\mathbf{G}}$ – to $P|_{\mathcal{E}_{\mathbf{G}}}$ in all possible ways. Therefore we can conclude $\mathcal{S}(\mathbf{G}') \subseteq \mathbf{sat}(\mathcal{S}(\mathbf{G}), \mathbf{f})$. ◀

It is a well-known fact that the category of sets and injective functions lacks pushouts. The same also holds for **PIDag**(\mathcal{L}). However, recall now that pushouts in **PDag**(\mathcal{L}) preserve monos (Lemma 28). Thus in the following we say that a functor $\mathbb{F} : \mathbf{PIDag}(\mathcal{L}) \rightarrow \mathbf{SPath}(\mathcal{L})$ weakly preserves finite pushouts (and in fact, finite colimits) if any commuting square in **PIDag**(\mathcal{L}) that is a pushout (via the inclusion functor) in **PDag**(\mathcal{L}) is mapped by \mathbb{F} to a pushout in **SPath**(\mathcal{L}).

► **Theorem 32.** *Let \mathcal{S} be a coherent specification. The induced functor $\mathbb{M}(\mathcal{S}) : \mathbf{PIDag}(\mathcal{L}) \rightarrow \mathbf{SPath}(\mathcal{L})$ weakly preserves finite colimits and preserves binary pullbacks.*

Proof. The initial object is easy, since it holds by construction. As for pushouts and pullbacks: since \mathcal{S} is coherent, it boils down to Lemma 24. ◀

7.2 Completeness

It is now time for moving to the completeness results of our work, showing (a few alternatives on) how to obtain a specification from a functor.

► **Theorem 33.** *Let $\mathbb{F} : \mathbf{PIDag}(\mathcal{L}) \rightarrow \mathbf{SPath}(\mathcal{L})$ be a functor such that $\mathbb{F}(\mathbf{G}) \subseteq \mathbb{P}(\mathcal{E}_{\mathbf{G}}, \lambda_{\mathbf{G}})$. If \mathbb{F} weakly preserves finite colimits and preserves binary pullbacks, it induces a coherent specification $\mathcal{S}(\mathbb{F})$.*

Proof. Let $\mathbb{S}(\mathbb{F})(\mathbb{G}) = \mathbb{F}(\mathbb{G})$. We shall show that $\mathbb{F}(\mathbb{G})$ is coherent. Consider the following pushout in $\mathbf{PDag}(\mathcal{L})$

$$\begin{array}{ccc} \mathbb{G}|_{[e_1] \cap [e_2]} & \hookrightarrow & \mathbb{G}|_{[e_2]} \\ \downarrow & & \downarrow \\ \mathbb{G}|_{[e_1]} & \hookrightarrow & \mathbb{G}|_{[e_1] \cup [e_2]} \end{array} \quad (7.1)$$

Since \mathbb{F} preserves pullbacks, thus monos, and weakly preserves pushouts, this diagram is mapped by \mathbb{F} to the following pushout in $\mathbf{SPath}(\mathcal{L})$

$$\begin{array}{ccc} \mathbb{F}(\mathbb{G}|_{[e_1] \cap [e_2]}) & \hookrightarrow & \mathbb{F}(\mathbb{G}|_{[e_2]}) \\ \downarrow & & \downarrow \\ \mathbb{F}(\mathbb{G}|_{[e_1]}) & \hookrightarrow & \mathbb{F}(\mathbb{G}|_{[e_1] \cup [e_2]}) \end{array} \quad (7.2)$$

where all underlying functions between events are inclusions. By Lemma 24 we have that

$$\mathbb{F}(\mathbb{G}|_{[e_1] \cup [e_2]}) \simeq \mathbb{F}(\mathbb{G}|_{[e_1]}) \otimes \mathbb{F}(\mathbb{G}|_{[e_2]})$$

Since clearly $\mathbb{G} = \mathbb{G}|_{\bigcup_{e \in \mathcal{E}_{\mathbb{G}}} [e]}$, by associativity of pushouts we obtain coherence

$$\mathbb{F}(\mathbb{G}) \simeq \bigotimes_{e \in \mathcal{E}_{\mathbb{G}}} \mathbb{F}(\mathbb{G}|_{[e]})$$

Isomorphism preservation follows from \mathbb{F} being a functor. ◀

Combined with Theorem 32, the result above intuitively tells us that the coherence of a specification roughly corresponds to the weak preservation of colimits. However, the set-theoretical requirement $\mathbb{F}(\mathbb{G}) \subseteq \mathbb{P}(\mathcal{E}_{\mathbb{G}}, \lambda_{\mathbb{G}})$ is still unsatisfactory, yet apparently unavoidable, because a generic \mathbb{F} could associate *any* set of paths to a graph. We can sharpen the result by requiring functors to preserve specific properties for suitable arrows of $\mathbf{PDag}(\mathcal{L})$. The candidates are root extension morphisms, given the properties shown in Section 6. In order to define the functors, we also need to consider a suitable subset of the arrows of $\mathbf{SPath}(\mathcal{L})$.

► **Definition 34** (Saturated specifications). *Let \mathcal{S} be a specification. It is saturated if for all graphs \mathbb{G} and extensions $\mathbb{G}_{\mathcal{E}}^{\ell}$ the inclusion $\mathbf{f} : \mathcal{E}_{\mathbb{G}} \rightarrow \mathcal{E}_{\mathbb{G}_{\mathcal{E}}^{\ell}}$ is saturated with respect to $\mathcal{S}(\mathbb{G}_{\mathcal{E}}^{\ell})$ (see Lemma 18), that is*

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

A *saturation ps-morphism* (along ℓ) is a saturated ps-morphism $\mathbf{f} : \mathcal{X}_1 \rightarrow \mathcal{X}_2$ with underlying function $(\mathcal{E}, \lambda) \rightarrow (\mathcal{E}_{\top}, \lambda[\top \mapsto \ell])$. We can now prove an instance of Theorem 33 concerning saturated specifications.

► **Proposition 35.** *Let $\mathbb{F} : \mathbf{PDag}(\mathcal{L}) \rightarrow \mathbf{SPath}(\mathcal{L})$ be a functor mapping root extension morphisms into saturation ps-morphisms (along the same labels). If \mathbb{F} weakly preserves finite colimits, it induces a saturated, coherent specification $\mathbb{S}(\mathbb{F})$.*

Proof. We first show that \mathbb{F} preserves monos, which renders the assumption of Theorem 33 about preservation of pullbacks redundant. We will essentially follow the proof of Proposition 29. Given $\mathbf{f} : \mathbb{G}_1 \rightarrow \mathbb{G}_2$ in $\mathbf{PDag}(\mathcal{L})$, we proceed by induction on the cardinality of $\mathcal{E}_{\mathbb{G}_2}$. If $\mathcal{E}_{\mathbb{G}_2}$ is \emptyset , i.e., it is the initial object, then \mathbf{f} is the identity on \emptyset , and the claim follows by functors preserving identities. Suppose now that \mathbb{G}_2 is rooted with root e . If $e \in \text{img}(\mathbf{f})$, then \mathbf{f} again is the identity. Otherwise, \mathbf{f} can be decomposed as $\mathbb{G}_1 \rightarrow (\mathbb{G}_2 \setminus e) \rightarrow \mathbb{G}_2$: the

left-most one satisfies the induction hypothesis, and the right-most one is a root extension morphism, which by hypothesis is mapped to a (monic) saturation ps-morphism. Therefore, by functoriality of \mathbb{F} , the claim holds for the composition of these morphisms. If \mathbf{G}_2 is not rooted, then \mathbf{f} can be similarly decomposed as $\mathbf{G}_1 \rightarrow \mathbf{G}_2|_{[e]} \rightarrow \mathbf{G}_2$. By induction the claim holds for the left-most morphism. The right-most one is obtained via a pushout of the form (7.1), which is mapped by \mathbb{F} to a pushout of the form (7.2), because \mathbb{F} (weakly) preserves finite colimits. By induction hypothesis, the span of this pushout consists of monic ps-morphisms, therefore we use Lemma 24 to conclude that the pushout morphisms are monic as well, hence the right-most morphism satisfies our claim. Again, the claim for the whole of \mathbf{f} follows from functoriality of \mathbb{F} . A similar inductive argument can be used to show that $\mathbb{F}(\mathbf{G})$ is a set of paths over $(\mathcal{E}_G, \lambda_G)$ (up to a label-preserving isomorphism of events). Therefore we can now re-use the proof of Theorem 33 and obtain that $\mathbb{S}(\mathbb{F})$ is a coherent specification.

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

7.3 More Completeness

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

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

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

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

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

► **Example 37.** A topological specification \mathcal{S}_{topR} for a **Register** can be defined as \mathcal{S}_{lwwR} in Example 13 with the additional requirement that paths are topological orderings of visibilities

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

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

► **Definition 38** (Topological specifications). *Let \mathcal{S} be a specification. It is topological if*

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

A *topological ps-morphism* (along ℓ) is a ps-morphism $\mathbf{f} : \mathcal{X}_1 \rightarrow \mathcal{X}_2$ with underlying function $(\mathcal{E}, \lambda) \rightarrow (\mathcal{E}_\top, \lambda[\top \mapsto \ell])$ such that $\mathcal{X}_2 = \mathbf{ext}(\mathbf{pre}(\mathcal{S}(\mathcal{X}_2), \mathbf{f}), \mathbf{f})$. The name is directly reminiscent of what are called topological RDTs in [10, 5], and in fact it similarly guarantees that arbitrations preserve the visibility order. We can thus prove another instance of Theorem 33, now concerning topological specifications.

► **Proposition 39.** *Let $\mathbb{F} : \mathbf{PIDag}(\mathcal{L}) \rightarrow \mathbf{SPath}(\mathcal{L})$ be a functor mapping root extension morphisms into topological ps-morphisms (along the same labels). If \mathbb{F} weakly preserves finite colimits, it induces a topological, coherent specification $\mathbb{S}(\mathbb{F})$.*

7.4 Interchangeability of Functors and Coherent Specifications

The connection between the construction of Theorem 32 and Theorem 33 is quite tight, and in fact induces a one-to-one correspondence between functors and coherent specifications.

► **Theorem 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 functor verifying the hypothesis of Theorem 33. 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 $\mathbb{M}'(\mathbb{G}) = \mathbb{S}(\mathbb{F})(\mathbb{G}) = \mathbb{F}(\mathbb{G})$, therefore we can define $\varphi_{\mathbb{G}} = \text{Id}_{\mathbb{F}(\mathbb{G})}$. We need to prove that it is natural, which in this case amounts to showing $\mathbb{M}'(\mathbf{f}) = \mathbb{F}(\mathbf{f})$, for $\mathbf{f} : \mathbb{G} \rightarrow \mathbb{G}'$ 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 $(\mathcal{E}_{\mathbb{G}}, \lambda_{\mathbb{G}}) \rightarrow (\mathcal{E}_{\mathbb{G}'}, \lambda_{\mathbb{G}'})$.

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, $\mathbb{S}(\mathbb{M}(\mathcal{S}))(\mathbb{G}) = \mathbb{M}(\mathcal{S})(\mathbb{G}) = \mathcal{S}(\mathbb{G})$. ◀

The one-to-one correspondence can be lifted to the specific classes of saturated/topological coherent specifications and to the functors of Proposition 35/Proposition 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.

► **Remark 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 $\mathcal{S}_1(\mathbb{G}) \subseteq \mathcal{S}_2(\mathbb{G})$ for all graphs \mathbb{G} . However, this is a very concrete characterisation: it would be more general to check for the existence of a ps-morphism $\mathcal{S}_2(\mathbb{G}_2) \rightarrow \mathcal{S}_1(\mathbb{G}_1)$ whose underlying function $\mathbf{f} : \mathcal{E}_{\mathbb{G}_2} \rightarrow \mathcal{E}_{\mathbb{G}_1}$ 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 $\mathbb{M}(\mathcal{S}_2) \rightarrow \mathbb{M}(\mathcal{S}_1)$.

8 Conclusions and Further Works

In this paper we have provided a functorial characterisation of RDT specifications. Our starting point is the denotational approach proposed in [7, 6], in which RDT specifications are associated with functions mapping visibility graphs into sets of admissible arbitrations that are also saturated and coherent, and where a preliminary functorial correspondence was proposed. In this paper we streamlined and expanded the latter result. We considered the category $\mathbf{PDag}(\mathcal{L})$ of labelled, acyclic graphs and pr-morphisms for representing visibility graphs. We equip $\mathbf{PDag}(\mathcal{L})$ with operators that model the evolution of visibility graphs and we show that the sub-category $\mathbf{PIDag}(\mathcal{L})$ of monic morphisms can be generated by the subset of root extensions via pushouts. For arbitrations, we take $\mathbf{SPath}(\mathcal{L})$, which is the category of sets of labelled, total orders and ps-morphisms. Then, we show that each coherent

specification mapping isomorphic graphs into isomorphic set of paths induces a functor $\mathbb{M}(\mathcal{S}) : \mathbf{PIDag}(\mathcal{L}) \rightarrow \mathbf{SPath}(\mathcal{L})$. Conversely, we prove that a functor $\mathbb{F} : \mathbf{PIDag}(\mathcal{L}) \rightarrow \mathbf{SPath}(\mathcal{L})$ that preserves finite colimits and binary pullbacks induces an coherent specification $\mathbb{S}(\mathbb{F})$. Moreover, $\mathbb{M}(\mathcal{S})$ and $\mathbb{S}(\mathbb{F})$ are shown to be inverses of each other.

With respect to the categorical results expressed in [7], besides the additional characterisation of topological specifications, the key improvement has been the proof that the coherence of specifications has a precise counterpart in terms of the weak preservation of colimits on their functorial presentations, as stated by Theorem 32 and Theorem 33. We thus removed the set-theoretical requirements occurring e.g. in [7, Section 5.3], as witnessed by the definition of *coherent* functor there. We believe that this purely functorial characterisation of RDTs, as further witnessed by Proposition 35 and Proposition 39, provides an ideal setting for the development of techniques for handling RDT composition, as briefly pointed out by the functorial characterisation of refinement between specifications. Our long term goal is to equip RDT specifications with a set of operators that enables us to specify and reason about complex RDTs compositionally, i.e., in terms of their constituent parts. We aim at providing a uniform formal treatment to the compositional approaches proposed in [1, 10, 12].

References

- 1 Carlos Baquero, Paulo Sérgio Almeida, Alcino Cunha, and Carla Ferreira. Composition in state-based replicated data types. *Bulletin of the EATCS*, 123, 2017.
- 2 Ahmed Bouajjani, Constantin Enea, and Jad Hamza. Verifying eventual consistency of optimistic replication systems. In Suresh Jagannathan and Peter Sewell, editors, *POPL 2014*, pages 285–296. ACM, 2014.
- 3 Sebastian Burckhardt, Alexey Gotsman, and Hongseok Yang. Understanding eventual consistency. Technical Report MSR-TR-2013-39, Microsoft Research, 2013.
- 4 Sebastian Burckhardt, Alexey Gotsman, Hongseok Yang, and Marek Zawirski. Replicated data types: specification, verification, optimality. In Suresh Jagannathan and Peter Sewell, editors, *POPL 2014*, pages 271–284. ACM, 2014.
- 5 Andrea Cerone, Giovanni Bernardi, and Alexey Gotsman. A framework for transactional consistency models with atomic visibility. In Luca Aceto and David de Frutos-Escrig, editors, *CONCUR 2015*, volume 42 of *LIPICs*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2015.
- 6 Fabio Gadducci, Hernán Melgratti, and Christian Roldán. On the semantics and implementation of replicated data types. *Science of Computer Programming*, 167:91–113, 2018.
- 7 Fabio Gadducci, Hernán C. Melgratti, and Christian Roldán. A denotational view of replicated data types. In Jean-Marie Jacquet and Mieke Massink, editors, *COORDINATION 2017*, volume 10319 of *LNCS*, pages 138–156. Springer, 2017.
- 8 Seth Gilbert and Nancy Lynch. Brewer’s conjecture and the feasibility of consistent, available, partition-tolerant web services. *SIGACT News*, 33(2):51–59, June 2002.
- 9 Alexey Gotsman and Sebastian Burckhardt. Consistency models with global operation sequencing and their composition. In Andréa W. Richa, editor, *DISC 2017*, volume 91 of *LIPICs*, pages 23:1–23:16. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017.
- 10 Alexey Gotsman and Hongseok Yang. Composite replicated data types. In Jan Vitek, editor, *ESOP 2015*, volume 9032 of *LNCS*, pages 585–609. Springer, 2015.
- 11 Gowtham Kaki, Kapil Earanky, K. C. Sivaramakrishnan, and Suresh Jagannathan. Safe replication through bounded concurrency verification. In *OOPSLA 2018*, volume 2 of *PACMPL*, pages 164:1–164:27. ACM, 2018.
- 12 Adriaan Leijnse, Paulo Sérgio Almeida, and Carlos Baquero. Higher-order patterns in replicated data types. In *PaPoC 2019*. ACM, 2019.

- 13 Marc Shapiro, N. Preguiça, Carlos Baquero, and Marek Zawirski. Conflict-free replicated data types. In Xavier Défago, Franck Petit, and Vincent Villain, editors, *SSS 2011*, volume 6976 of *LNCS*, pages 386–400. Springer, 2011.
- 14 K. C. Sivaramakrishnan, Gowtham Kaki, and Suresh Jagannathan. Declarative programming over eventually consistent data stores. In David Grove and Steve Blackburn, editors, *PLDI 2015*, pages 413–424. ACM, 2015.