

A Categorical Account for the Specification of Replicated Data Type

Fabio Gadducci 

Dipartimento di Informatica, Università di Pisa, Italia

fabio.gadducci@unipi.it

Hernán Melgratti 

Departamento de Computación, Universidad de Buenos Aires, Argentina

ICC-CONICET-UBA, Argentina

hmelgra@dc.uba.ar

Christian Roldán 

Departamento de Computación, Universidad de Buenos Aires, Argentina

croldan@dc.uba.ar

Matteo Sammartino 

Department of Computer Science, University College London, UK

m.sammartino@ucl.ac.uk

Abstract

Replicated Data Types (RDTs) have been introduced as suitable abstractions 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 a categorical 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 General and reference → General literature; General and reference

Keywords and phrases Replicated data type, Specification, Functorial characterisation

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2019.23

1 Introduction

The CAP theorem establishes that a distributed data store can simultaneously provide two of the following three properties: consistency, availability, and tolerate 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 (RDT) 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 operations are executed. A view is usually represented by a *visibility* relation, which is a binary, acyclic relation over executed operations (a.k.a. *events*). The state of a store is described instead as total order over events, called *arbitrations*, which describes the way in which conflicting concurrent operations are resolved. Different specification approaches for RDTs build on the notions of visibility and arbitration [2, 3, 4, 5, 7, 9, 11, 13, 14]. A purely functional approach for the specification of RDTs has been presented in [7], where an RDT is associated with a function that maps visibility into set of arbitrations.



© John Q. Public and Joan R. Public;

licensed under Creative Commons License CC-BY

39th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science .

Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:13

Leibniz International Proceedings in Informatics



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

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

(a) Specification of a Counter

(b) Non-admissible arbitrations

■ Figure 1 Counter specifications

45 For illustration purposes, consider the RDT **Counter**, which maintains an integer value
 46 and provides two operations for incrementing and reading its current value. The specification
 47 of **Counter** states that every increment is always successful, while the expected result for
 48 a read operation is the number of increments seen by that read, regardless of the order
 49 in which such operations are arbitrated. Following the approach in [7], the RDT **Counter**
 50 is specified as a function \mathcal{S}_{Ctr} that maps visibility relations into sets of arbitrations. For
 51 instance, Figure 1a illustrates the case for a visibility relation that involves an increment
 52 event seen by a read event: events are depicted by pairs $\langle \text{operation}, \text{expected_result} \rangle$,
 53 and **inc** stands for increment and **rd** for read. Note that the expected result for the read
 54 event is 1, which coincides with the quantity of events labelled by $\langle \text{inc}, \text{ok} \rangle$ seen by that
 55 read. The function \mathcal{S}_{Ctr} maps that visibility relation into a set containing two arbitrations
 56 of the events, i.e., two total orders, for the events of the visibility relation. We remark that
 57 arbitration does not mean real time ordering, but just a way in which a store can totally
 58 order events, which may not respect the causal order of operations. In fact, the second
 59 arbitration orders the read event before the increment one, despite the first event causally
 60 depends on the second one. Figure 1b shows instead a case in which the specification maps a
 61 visibility relation into an empty set of arbitrations, which means that such visibility relation
 62 is not accepted by the specification. Basically, that visibility is rejected because there is a
 63 read event that sees an increment event but returns 0 instead of the expected result 1.

64 This work develops the approach suggested in [7] for the categorical characterisation
 65 of RDT specifications. We consider the category $\mathbf{PIDag}(\mathcal{L})$ of labelled, directed acyclic
 66 graphs and *pr-morphisms*, i.e., label-preserving morphisms that reflect directed edges, and
 67 the category $\mathbf{SPath}(\mathcal{L})$ of sets of labelled, total orders and *ps-morphisms*, i.e., morphisms
 68 between set of paths. A ps-morphism $f : \mathcal{X}_1 \rightarrow \mathcal{X}_2$ from a set of paths \mathcal{X}_1 to a set of path
 69 \mathcal{X}_2 states that any total order in \mathcal{X}_2 can be obtained by extending some total order in
 70 \mathcal{X}_1 . In this work we show that a large class of specifications, dubbed *iso-coherent*, can be
 71 characterised functorially. Roughly, a coherent specification accounts for those RDTs such that
 72 the arbitrations associated with a visibility relation can be obtained by extending arbitrations
 73 associated with “smaller” visibilities. An iso-coherent specification is a coherent specification
 74 that maps isomorphic graphs into isomorphic sets of paths. We establish a bijection between
 75 functors and specifications, showing that an iso-coherent specification induces a functor from
 76 $\mathbf{PIDag}(\mathcal{L})$ into $\mathbf{SPath}(\mathcal{L})$ that preserves colimits and binary pullbacks and vice versa.

77 The paper has the following structure. Section 2 offers some preliminaries on categories
 78 of relations, which are used for proposing some basic results on categories of graphs and
 79 paths in Section 3. Section 4 recalls the set-theoretical presentation of RDTs introduced in [7].
 80 Section 5 introduces our semantical model, the category of set of paths, describing some of its
 81 basic properties with respect to limits and colimits. On Section 6 we presents some categorical
 82 operators for RDTs, which are used in Section 7 to present our main characterisation results.
 83 The paper is closed with some final remarks and some hints towards future works.

84 2 Preliminaries on Relations

85 **Relations.** Given a set E , a (binary) *relation* over E is a sub-set $\rho \subseteq E \times E$ of the cartesian
86 product of E with itself. We write $\langle E, \rho \rangle$ for a relation over E , and \emptyset to denote the empty
87 relation. The downward closure of $E' \subseteq E$ is a set such that $\forall e \in E, e' \in E'. e \rho e'$ implies $e \in E'$.

88 In addition, we write $\lfloor e \rfloor$ to stand for the downward closure of a single element 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

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

89 Note that both classes of morphisms are closed under composition: we denote as **Rel** the
90 category of relations and their morphisms and **PRel** the sub-category of pr-morphisms.

91 ► **Lemma 2** (Characterising pr-morphisms). *Let $f : \langle E, \rho \rangle \rightarrow \langle T, \gamma \rangle$ be a morphism. If it is
92 order-reflecting and downward closed, that is*

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

95 *then it is a pr-morphism. If f is injective, then the vice-versa holds.*

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

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

100 In other words, since **Rel** has finite limits and finite colimits, finite colimits and binary
101 pullbacks in **PRel** always exist and are computed as in **Rel**. There is no terminal object,
102 since morphisms in **Rel** into the singleton are clearly not past-reflecting.

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

105 ► **Lemma 4** (Monos under pushouts). *Pushouts in **Rel** (and thus in **PRel**) preserve monos.*

106 We now introduce labelled relations. Consider the forgetful functors $U_r : \mathbf{Rel} \rightarrow \mathbf{Set}$ and
107 $U_p : \mathbf{PRel} \rightarrow \mathbf{Set}$, the latter factoring through the inclusion functor $\mathbf{PRel} \rightarrow \mathbf{Rel}$. Chosen a
108 set \mathcal{L} of labels, we consider the comma categories $\mathbf{Rel}(\mathcal{L}) = U_r \downarrow \mathcal{L}$ and $\mathbf{PRel}(\mathcal{L}) = U_p \downarrow \mathcal{L}$:
109 it is well known that all the relevant structure is preserved in such comma categories.

110 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
111 label-preserving morphism $(E, \rho, \lambda) \rightarrow (E', \rho', \lambda')$ is a morphism $f : (E, \rho) \rightarrow (E', \rho')$ such that
112 $\forall s \in E. \lambda(s) = \lambda'(f(s))$. Moreover, finite limits and finite colimits are computed as in **Rel**.
113 The same characterisation also holds for the objects and the morphisms of $U_p \downarrow \mathcal{L}$.

114 3 Categories of Graphs and Paths

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

117 ► **Definition 5** (Directly acyclic graphs category). **PDag** is the full sub-category of **PRel**
118 whose objects are directed acyclic graphs.

119 In other terms, objects are relations whose transitive closure is a *strict* partial order.

120 ► **Remark 6**. The category whose arrows are morphisms is not that interesting, categorically
121 speaking, because, e.g., it does not admit pushouts, not even along monos. The one with
122 pr-morphisms is much more so, still remaining computationally simple.

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

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

126 ► **Definition 8** (Paths Category). **Path** is the full sub-category of **Rel** whose objects are
127 paths.

128 Note that the sub-category of just pr-morphisms is not so relevant, since there exists a
129 pr-morphism between two paths if and only if one path is a prefix of the other.

130 ► **Proposition 9** (Properties of **Path**). The inclusion functor $\mathbf{Path} \rightarrow \mathbf{Rel}$ reflects finite
131 colimits.

132 As for relations, we consider suitable comma categories in order to capture labelled paths
133 and graphs. In particular, we use the forgetful functors $\mathbf{U}_{\text{rp}} : \mathbf{Path} \rightarrow \mathbf{Set}$ and $\mathbf{U}_{\text{pd}} : \mathbf{PDag} \rightarrow$
134 **Set**: for a set of labels \mathcal{L} we denote $\mathbf{PDag}(\mathcal{L}) = \mathbf{U}_{\text{rp}} \downarrow \mathcal{L}$ and $\mathbf{Path}(\mathcal{L}) = \mathbf{U}_{\text{pd}} \downarrow \mathcal{L}$. Once
135 more, the relevant categorical structure is preserved and computed as in **Rel**.

136 4 Replicated Data Type Specification

137 We briefly recall the set-theoretical model of replicated data types (RDT), introduced in [7].
138 Our main result is its categorical characterisation, which is given in the following sections.

139 First, some notation. We denote a graph as $\langle \mathcal{E}, \prec, \lambda \rangle$ and a path as $\langle \mathcal{E}, \leq, \lambda \rangle$, in order to
140 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
141 $\mathbf{G}|_{\mathcal{E}'}$, the obvious restriction (and the same for a path **P**).

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

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

$$145 \quad \bigotimes \mathcal{X} = \{ \mathbf{P} \mid \mathbf{P} \text{ is a path over } \bigcup_i \mathcal{E}_i \text{ and } \mathbf{P}|_{\mathcal{E}_i} \in \mathcal{X} \}$$

146 Intuitively, the product of paths is analogous to the synchronous product of transition
147 systems, in which common elements are identified and the remaining ones can be freely
148 interleaved, as long as the original orders are respected. A set of sets of paths $\mathcal{X}_1, \mathcal{X}_2, \dots$ is
149 compatible if $\bigcup_i \mathcal{X}_i$ is so. In such case we can define the product $\bigotimes_i \mathcal{X}_i$ as $\bigotimes \bigcup_i \mathcal{X}_i$.

150 Now, let us further denote with $\mathbb{G}(\mathcal{L})$ and $\mathbb{P}(\mathcal{L})$ the sets of graphs and paths, respectively,
151 labelled over \mathcal{L} and with ϵ the empty graph. Also, when the set of labels \mathcal{L} is chosen, we let
152 $\mathbb{G}(\mathcal{E}, \lambda)$ and $\mathbb{P}(\mathcal{E}, \lambda)$ the sets of graphs and paths, respectively, whose elements are those in \mathcal{E}
153 and are labelled by $\lambda : \mathcal{E} \rightarrow \mathcal{L}$.

154 ► **Definition 11** (Specifications). A specification \mathcal{S} is a function $\mathcal{S} : \mathbb{G}(\mathcal{L}) \rightarrow 2^{\mathbb{P}(\mathcal{L})}$ such that
 155 $\mathcal{S}(\epsilon) = \{\epsilon\}$ and $\forall \mathbf{G}. \mathcal{S}(\mathbf{G}) \in 2^{\mathbb{P}(\mathcal{E}_{\mathbf{G}}, \lambda_{\mathbf{G}})}$.

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

159 As shown in [7], Definition 11 offers an alternative characterisation of RDTs [4] for a
 160 suitable choice of the set of labels. In particular, an RDT boils down to a specification labelled
 161 over pairs $\langle \text{operation}, \text{value} \rangle$ that is *saturated* and *past-coherent*. The former property is a
 162 technical one: roughly, if \mathbf{G}' is an extension of \mathbf{G} with an fresh event e , then the admissible
 163 arbitrations that a saturated specification \mathcal{S} assigns to \mathbf{G}' (i.e., the set of paths $\mathcal{S}(\mathbf{G}')$) are
 164 the admissible arbitrations of \mathbf{G} saturated with respect to e , i.e., all the paths that extends a
 165 path in $\mathcal{S}(\mathbf{G})$ with e inserted at an arbitrary position. Coherence instead is fundamental and
 166 expresses that admissible arbitrations of a visibility graph can be obtained by composing the
 167 admissible arbitrations of smaller visibilities.

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

$$170 \quad \forall \mathbf{G} \neq \epsilon. \mathcal{S}(\mathbf{G}) = \bigotimes_{e \in \mathcal{E}_{\mathbf{G}}} \mathcal{S}(\mathbf{G}|_{[e]})$$

171 Explicitly, in a coherent specification \mathcal{S} the arbitrations of a configuration \mathbf{G} (i.e., the set
 172 of paths $\mathcal{S}(\mathbf{G})$) are the composition of the arbitrations associated with its sub-graphs $\mathbf{G}|_{[e]}$.

173 Next example illustrates a saturated and coherent specification for the **Counter** RDT.

174 ► **Example 13** (Counter). Fix the following set of labels: $\mathcal{L} = \{\langle \text{inc}, \text{ok} \rangle\} \cup (\{\text{rd}\} \times \mathbb{N})$.
 175 Then, the specification of the RDT **Counter** is given by the function \mathcal{S}_{Ctr} defined such that

$$176 \quad \begin{aligned} P \in \mathcal{S}_{Ctr}(\mathbf{G}) \\ \text{iff} \\ \forall e \in \mathcal{E}_{\mathbf{G}}. \forall k. \lambda(e) = \langle \text{rd}, k \rangle \text{ implies } k = \#\{e' \mid e' \prec_{\mathbf{G}} e \text{ and } \lambda(e') = \langle \text{inc}, \text{ok} \rangle\} \end{aligned}$$

177 Intuitively, a visibility graph \mathbf{G} is mapped to a non-empty set of arbitrations (i.e., $\mathcal{S}_{Ctr}(\mathbf{G}) \neq$
 178 \emptyset) only when each event e in \mathbf{G} associated with a read operation has a return value k that
 179 matches the number of increments preceding e in \mathbf{G} . We remark that this specification is
 180 *coherent* and *saturated*. Saturation follows immediately because the definition of \mathcal{S}_{Ctr} does not
 181 impose any constraint on the ordering of events for the arbitrations P in $\mathcal{S}_{Ctr}(\mathbf{G})$. Coherence
 182 can be shown as follows. By definition of \mathcal{S}_{Ctr} , $P \in \mathcal{S}_{Ctr}(\mathbf{G})$ implies $P|_{[e]} \in \mathcal{S}_{Ctr}(\mathbf{G}|_{[e]})$ for
 183 all $e \in \mathcal{E}_{\mathbf{G}}$. Consequently, $P \in \bigotimes_{e \in \mathcal{E}_{\mathbf{G}}} \mathcal{S}(\mathbf{G}|_{[e]})$. On the contrary, take $P \in \bigotimes_{e \in \mathcal{E}_{\mathbf{G}}} \mathcal{S}(\mathbf{G}|_{[e]})$.
 184 Then, $e \in \mathcal{E}_{\mathbf{G}}$ implies $e \in \mathcal{E}_P$. Moreover, $e \in \mathcal{E}_P$ implies $\lambda(e) = \langle \text{rd}, k \rangle$ iff $k = \#\{e' \mid e' \prec_{\mathbf{G}}$
 185 $e \text{ and } \lambda(e') = \langle \text{inc}, \text{ok} \rangle\}$. Hence, $P \in \mathcal{S}_{Ctr}(\mathbf{G})$. Therefore, the equality in Definition 12 holds.

186 5 The model category

187 In order to provide a categorical characterisation of coherent specifications, we must first
 188 define precisely the model category. So far, we know that its objects have to be sets of
 189 compatible paths. We fix a set of labels \mathcal{L} , and we start looking at morphisms.

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

$$192 \quad \text{sat}(P, \mathbf{f}) = \{Q \mid Q \in \mathbb{P}(\mathcal{E}, \lambda) \text{ and } \mathbf{f} \text{ induces a path morphism } \mathbf{f} : P \rightarrow Q\}$$

193 The notion of saturation is extended to sets of paths $\mathcal{X} \subseteq \mathbb{P}(\mathcal{E}, \lambda)$ as $\bigcup_{P \in \mathcal{X}} \text{sat}(P, \mathbf{f})$.

23:6 A Categorical Account for the Specification of Replicated Data Type

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

195 ► **Example 15.** Consider the (injective, label-preserving) function \mathbf{f} mapping two events
196 with labels $\{\mathbf{a}, \mathbf{b}\}$ to three events with labels $\{\mathbf{a}, \mathbf{b}, \mathbf{c}\}$. Then we have

$$197 \quad \text{sat} \left(\begin{array}{c} \mathbf{a} \\ | \\ \mathbf{b} \end{array}, \mathbf{f} \right) = \left\{ \begin{array}{c} \mathbf{a} \quad \mathbf{a} \quad \mathbf{c} \\ | \quad | \quad | \\ \mathbf{b}, \mathbf{c}, \mathbf{a} \\ | \quad | \quad | \\ \mathbf{c} \quad \mathbf{b} \quad \mathbf{b} \end{array} \right\}$$

198 Intuitively, saturation adds \mathbf{c} – and in general events not in the image of \mathbf{f} – to the original
199 path in all possible ways, preserving the order of original events.

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

201 ► **Definition 16 (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
202 path-set morphism (shortly, ps-morphism) $\mathbf{f} : \mathcal{X}_1 \rightarrow \mathcal{X}_2$ is a function $\mathbf{f} : (\mathcal{E}_1, \lambda_1) \rightarrow (\mathcal{E}_2, \lambda_2)$
203 preserving labels such that

$$204 \quad \mathcal{X}_2 \subseteq \text{sat}(\mathcal{X}_1, \mathbf{f})$$

205 Intuitively, there is a ps-morphism from the set of paths \mathcal{X}_1 to the set of path \mathcal{X}_2 if any
206 path in \mathcal{X}_2 can be obtained by adding events to some path in \mathcal{X}_1 . This notion captures the
207 idea that arbitrations of larger visibilities are obtained as extensions of smaller visibilities.

208 ► **Example 17.** Consider the following three sets and the function \mathbf{f} from Example 15.

$$209 \quad \mathcal{X}_1 = \left\{ \begin{array}{c} \mathbf{a} \\ | \\ \mathbf{b} \end{array} \right\} \quad \mathcal{X}_2 = \left\{ \begin{array}{c} \mathbf{a} \quad \mathbf{a} \\ | \quad | \\ \mathbf{b}, \mathbf{c} \\ | \quad | \\ \mathbf{c} \quad \mathbf{b} \end{array} \right\} \quad \mathcal{X}_3 = \left\{ \begin{array}{c} \mathbf{a} \quad \mathbf{b} \\ | \quad | \\ \mathbf{b}, \mathbf{c} \\ | \quad | \\ \mathbf{c} \quad \mathbf{a} \end{array} \right\}$$

210 \mathbf{f} induces a ps-morphism $\mathbf{f} : \mathcal{X}_1 \rightarrow \mathcal{X}_2$ because $\mathcal{X}_2 \subseteq \text{sat}(\mathcal{X}_1, \mathbf{f})$ ($\text{sat}(\mathcal{X}_1, \mathbf{f})$ is depicted in
211 Example 15). On the contrary, there is no ps-morphism from \mathcal{X}_1 to \mathcal{X}_3 because the second
212 path of \mathcal{X}_3 cannot be obtained by extending some path of \mathcal{X}_1 with an event labelled by \mathbf{c} .

213 ► **Definition 18 (Retraction).** Let \mathbf{Q} be a path and $\mathbf{f} : \mathcal{E} \rightarrow \mathcal{E}_{\mathbf{Q}}$ a function. The retraction of \mathbf{Q}
214 along \mathbf{f} is defined as

$$215 \quad \text{ret}(\mathbf{Q}, \mathbf{f}) = \{\mathbf{P} \mid \mathbf{P} \in \mathbb{P}(\mathcal{E}, \lambda) \text{ and } \mathbf{f} \text{ induces a path morphism } \mathbf{f} : \mathbf{P} \rightarrow \mathbf{Q}\}$$

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

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

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

220 ► **Lemma 19.** 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
221 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.

222 ► **Lemma 20.** Let $\mathcal{X}_2 \subseteq \mathbb{P}(\mathcal{E}_2, \lambda_2)$ be a set of paths and $\mathbf{f} : \mathcal{E}_1 \rightarrow \mathcal{E}_2$ a function. Then
223 $\mathcal{X}_2 \subseteq \text{sat}(\text{ret}(\mathcal{X}_2, \mathbf{f}), \mathbf{f})$.

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

225 ▶ **Example 21.** Consider the ps-morphism

$$226 \quad f: \left\{ \begin{array}{c} a \\ | \\ b \end{array} \right\} \rightarrow \left\{ \begin{array}{c} a \\ | \\ b \\ | \\ c \end{array} \right\}$$

227 whose underlying function is f from Example 15. This is *not* saturated. In fact, we have

$$228 \quad \left\{ \begin{array}{c} a \\ | \\ b \\ | \\ c \end{array} \right\} \neq \text{sat}(\text{ret}\left(\left\{ \begin{array}{c} a \\ | \\ b \\ | \\ c \end{array} \right\}, f\right), f) = \text{sat}\left(\left\{ \begin{array}{c} a \\ | \\ b \end{array} \right\}, f\right) = \left\{ \begin{array}{c} a & a & c \\ | & | & | \\ b & c & a \\ | & | & | \\ c & b & b \end{array} \right\}$$

229 ▶ **Definition 22** (Sets of Paths Category). We define $\mathbf{SPath}(\mathcal{L})$ as the category whose objects
230 are sets of paths $\mathcal{X} \subseteq \mathbb{P}(\mathcal{E}, \lambda)$ and morphisms are ps-morphisms.

231 ▶ **Proposition 23** (Properties of \mathbf{SPath}). The category $\mathbf{SPath}(\mathcal{L})$ has finite colimits along
232 monos and binary pullbacks.

233 **Proof.** (Strict) initial object. The choice is $\langle \emptyset, \{\epsilon\}, \emptyset \rangle$, with $\epsilon \in \mathbb{P}(\emptyset, \emptyset)$ the empty path. Let
234 $\mathcal{X} \subseteq \mathbb{P}(\mathcal{E}, \lambda)$ and $!: \emptyset \rightarrow \mathcal{E}$ the unique function. We have a function $!: (\emptyset, \emptyset) \rightarrow (\mathcal{E}, \lambda)$ such
235 that $\mathcal{X} \subseteq \text{sat}(\{\epsilon\}, !) = \mathbb{P}(\mathcal{E}, \lambda)$.

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

239 *Binary Pullbacks.* Let $\mathcal{X}, \mathcal{X}_1$, and \mathcal{X}_2 be sets of paths and $f_i: \mathcal{X}_i \rightarrow \mathcal{X}$ ps-morphisms.
240 Consider the underlying functions $f_i: \mathcal{E}_i \rightarrow \mathcal{E}$ and their pullback $f'_i: \mathcal{E}_1 \times_{\mathcal{E}} \mathcal{E}_2 \rightarrow \mathcal{E}$ in the
241 category of sets. This induces a pullback $f'_i: \text{ret}(\mathcal{X}_1, f'_1) \cup \text{ret}(\mathcal{X}_2, f'_2) \rightarrow \mathcal{X}_i$ in $\mathbf{SPath}(\mathcal{L})$.
242 ◀

243 The characterisation of pushouts might not work, should not be on a span of injective
244 functions. To help intuition, we now instantiate the constructions above to suitable inclusions.

245 ▶ **Lemma 24.** Let $f_i: \mathcal{X} \rightarrow \mathcal{X}_i$ be ps-morphisms such that the underlying functions $f_i: \mathcal{E} \rightarrow \mathcal{E}_i$
246 are inclusions and $\mathcal{E} = \mathcal{E}_1 \cap \mathcal{E}_2$. Then the pushout is given by $f'_i: \mathcal{X}_i \rightarrow \mathcal{X}_1 \otimes \mathcal{X}_2$.

247 **Proof.** By definition $\mathcal{X}_1 \otimes \mathcal{X}_2 = \{P \mid P \text{ is a path over } \bigcup_i \mathcal{E}_i \text{ and } P|_{\mathcal{E}_i} \in \mathcal{X}_i\}$. Note also that
248 $\text{sat}(\mathcal{X}_i, f'_i) = \bigcup_{Q \in \mathcal{X}_i} \{P \mid P \in \mathbb{P}(\bigcup_i \mathcal{E}_i, \bigcup_i \lambda_i) \text{ and } f'_i \text{ induces a path morphism } f'_i: P \rightarrow Q\}$.
249 Since f'_i is an inclusion, the latter condition equals to $P|_{\mathcal{E}_i} = Q$, thus the property holds. ◀

250 ▶ **Example 25.** Consider the following ps-morphisms

$$251 \quad f_1: \left\{ \begin{array}{cc} a & b \\ | & | \\ b & a \end{array} \right\} \rightarrow \left\{ \begin{array}{c} a \\ | \\ b \\ | \\ c \end{array} \right\} \quad f_2: \left\{ \begin{array}{cc} a & b \\ | & | \\ b & a \end{array} \right\} \rightarrow \left\{ \begin{array}{cc} a & b \\ | & | \\ b & a \\ | & | \\ d & d \end{array} \right\}$$

252 then, the pushout is given by the following two morphisms

$$253 \quad g_1: \left\{ \begin{array}{c} a \\ | \\ b \\ | \\ c \end{array} \right\} \rightarrow \left\{ \begin{array}{cc} a & b \\ | & | \\ b & a \\ | & | \\ c & d \\ | & | \\ d & c \end{array} \right\} \quad g_2: \left\{ \begin{array}{cc} a & b \\ | & | \\ b & a \\ | & | \\ d & d \end{array} \right\} \rightarrow \left\{ \begin{array}{cc} a & b \\ | & | \\ b & a \\ | & | \\ c & d \\ | & | \\ d & c \end{array} \right\}$$

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

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

254 6 Operators for Visibility

255 We now introduce a family of operations that will be handy for our categorical characterisation.
 256 First, we provide a new operation on visibility relations.

257 ► **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}
 258 over \mathcal{E}' with ℓ as the graph $\mathbf{G}_{\mathcal{E}'}^\ell = \langle \mathcal{E}_\top, \prec \cup (\mathcal{E}' \times \{\top\}), \lambda[\top \mapsto \ell] \rangle$.*

259 Intuitively, $\mathbf{G}_{\mathcal{E}'}^\ell$ is obtained by adding to the visibility relation \mathbf{G} one additional event which
 260 sees the events in \mathcal{E}' . We will just write \mathbf{G}^ℓ whenever \mathcal{E}' is the set of top elements of \mathbf{G} – i.e.,
 261 the additional event may see *all* the events of \mathbf{G} – and we call it *top extensions*. Note how
 262 top extensions can be lifted to endofunctors (and actually, monads) on $\mathbf{PDag}(\mathcal{L})$. Extension
 263 allows us to characterise *saturated* specifications.

264 ► **Definition 27** (Saturated specification). *Let \mathcal{S} be a specification. It is saturated if for all
 265 graphs \mathbf{G} the inclusion $\mathbf{f} : \mathcal{E}_{\mathbf{G}} \rightarrow \mathcal{E}_{\mathbf{G}^\ell}$ is saturated with respect to $\mathcal{S}(\mathbf{G}_{\mathcal{E}}^\ell)$ (see Lemma 20), that is*

$$266 \quad \forall \mathbf{G}. \mathcal{S}(\mathbf{G}_{\mathcal{E}}^\ell) = \text{sat}(\text{ret}(\mathcal{S}(\mathbf{G}_{\mathcal{E}}^\ell), \mathbf{f}), \mathbf{f})$$

267 We now show that all graphs are generated from suitable top extensions via pushout
 268 contructions. We consider *tree* extensions $\mathbf{T} \rightarrow \mathbf{T}^\ell$ for a tree \mathbf{T} , i.e., a graph such that each
 269 event has a unique successor. Intuitively, trees represent the simplest visibility relations, and
 270 can be seen as “generators” for $\mathbf{PDag}(\mathcal{L})$. We first show that trees are freely generated via
 271 pushouts and tree extensions.

272 ► **Lemma 28.** *The sub-category of $\mathbf{PIDag}(\mathcal{L})$ of trees is freely generated from the empty
 273 tree via coproduct and tree extensions.*

274 Now we can show that trees and monic arrows between them generate the whole $\mathbf{PIDag}(\mathcal{L})$
 275 via pushouts.

276 ► **Lemma 29.** *Every monic arrow $\mathbf{f} : \mathbf{G}' \rightarrow \mathbf{G}$ of $\mathbf{PIDag}(\mathcal{L})$ is given by a pushout in $\mathbf{PDag}(\mathcal{L})$
 277 of the form*

$$278 \quad \begin{array}{ccc} \mathbf{T}' & \xhookrightarrow{\mathbf{f}'} & \mathbf{T} \\ \downarrow & & \downarrow \\ \mathbf{G}' & \xhookrightarrow{\mathbf{f}} & \mathbf{G} \end{array}$$

279 where $\mathbf{f}' : \mathbf{T}' \rightarrow \mathbf{T}$ is a monic arrow between trees.

280 **Proof.** Given a graph \mathbf{G} , we proceed by induction on the set \mathcal{E} of the events of \mathbf{G} .

281 For the base case let us now consider a graph $\mathbf{G} = \mathbf{G}|_{\{e\}}$ for a (necessarily unique) $e \in \mathcal{E}$.
 282 Note that we can find an epic pr-morphism $\mathbf{f} : \mathbf{T} \rightarrow \mathbf{G}$, for a tree \mathbf{T} . This induces another

283 epic pr-morphism $T|_{\mathcal{E}_\tau \setminus \{f^{-1}(e)\}} \rightarrow G|_{\mathcal{E}_\tau \setminus \{e\}}$. Since $(T|_{\mathcal{E}_\tau \setminus \{f^{-1}(e)\}})^\ell$ is isomorphic to T , G is now
 284 obtained as the obvious pushout.

285 The inductive step is immediate. In fact, note that $G = \bigcup_{e \in \mathcal{E}} G|_{[e]}$ for \mathcal{E} the set of top
 286 elements and let $e_1 \in \mathcal{E}$ and $\mathcal{E}_1 = \mathcal{E} \setminus \{e_1\}$. Then, $G = G|_{\mathcal{E}_1} \cup G|_{[e_1]}$ is the obvious pushouts
 287 of two inclusions.

288

289 7 A categorical correspondence

290 It is now time to move towards our categorical characterisation of specifications. In this
 291 section we will show that coherent specifications induce functors preserving relevant structure
 292 (soundness) and, viceversa, that a certain class of functors induce coherent specifications
 293 (completeness). Finally, we show that these functors are “mutually inverse”.

294 We first provide a simple technical result for coherent specifications.

295 ► **Lemma 30.** *Let \mathcal{S} be a coherent specification and $\mathcal{E} \subseteq \mathcal{E}_G$. If $\mathcal{E} = \bigcup_{e \in \mathcal{E}} [e]$ (that is, if \mathcal{E} is
 296 downward closed in G), then $\mathcal{S}(G)|_{\mathcal{E}} \subseteq \mathcal{S}(G|_{\mathcal{E}})$.*

297 **Proof.** Since \mathcal{E} is downward closed, for all $e \in \mathcal{E}$ we have that $(G|_{\mathcal{E}})|_{[e]} = G|_{[e]}$. Now,
 298 by the latter and by coherence we have that $\mathcal{S}(G)|_{\mathcal{E}} = (\bigotimes_{e \in \mathcal{E}_G} \mathcal{S}(G|_{[e]}))|_{\mathcal{E}}$ and $\mathcal{S}(G|_{\mathcal{E}}) =$
 299 $\bigotimes_{e \in \mathcal{E}} \mathcal{S}(G|_{[e]})$. Note that $(\bigotimes_{e \in \mathcal{E}_G} \mathcal{S}(G|_{[e]}))|_{\mathcal{E}} \subseteq \bigotimes_{e \in \mathcal{E}} \mathcal{S}(G|_{[e]})$, because a path P can always
 300 be restricted to a suitable path on fewer events (the viceversa in general does not hold). This
 301 concludes the proof. ◀

302 Our second step is to further curb the arrows in our syntax category to *monic* ones. Intuitively,
 303 we are only interested in what happens if we add further events to visibility relations. Note
 304 that a morphism in $\mathbf{PDag}(\mathcal{L})$ is a mono if and only if the underlying function is injective. We
 305 thus consider the sub-category $\mathbf{PIDag}(\mathcal{L})$ of direct acyclic graphs and monic pr-morphisms.

306 We now give our soundness results. We assume that specifications are *iso-coherent*, i.e.,
 307 they map isomorphic graphs to isomorphic sets of paths (along the same isomorphism on
 308 events).

309 ► **Proposition 31** (functors induced by specifications). *An iso-coherent specification \mathcal{S} induces
 310 a functor $\mathbb{M}(\mathcal{S}) : \mathbf{PIDag}(\mathcal{L}) \rightarrow \mathbf{SPath}(\mathcal{L})$.*

311 **Proof.** We define $\mathbb{M}(\mathcal{S})(G) = \mathcal{S}(G)$ and $\mathbb{M}(\mathcal{S})(f)$ as the ps-morphism with underlying injective
 312 function $f : (\mathcal{E}_G, \lambda_G) \hookrightarrow (\mathcal{E}_{G'}, \lambda_{G'})$. The proof boils down to show that f really is a ps-morphism
 313 from $\mathcal{S}(G)$ into $\mathcal{S}(G')$, i.e., $\mathcal{S}(G') \subseteq \mathbf{sat}(\mathcal{S}(G), f)$ and, since we are considering specifications
 314 preserving isos, we can restrict our attention to the case where f is an inclusion.

315 Since f is a pr-morphism, $\bigcup_{e \in \mathcal{E}_G} f(e)$ is downward-closed in G' and thus by Lemma 30 we
 316 have $\mathcal{S}(G')|_{\mathcal{E}_G} \subseteq \mathcal{S}(G'|_{\mathcal{E}_G}) = \mathcal{S}(G)$, the latter equality by iso-coherence. Now, consider a path
 317 $P \in \mathcal{S}(G')$. Since $P|_{\mathcal{E}_G} \in \mathcal{S}(G)$, we have $P \in \mathbf{sat}(\mathcal{S}(G), f)$, because saturation adds missing
 318 events – namely those in $\mathcal{E}_{G'} \setminus \mathcal{E}_G$ – to $P|_{\mathcal{E}_G}$ in all possible ways. Therefore we can conclude
 319 $\mathcal{S}(G') \subseteq \mathbf{sat}(\mathcal{S}(G), f)$.

320

321 A simple corollary instantiates the result to saturated specifications. So, let $\mathbf{SSPath}(\mathcal{L})$
 322 be the sub-category of $\mathbf{SPath}(\mathcal{L})$ of saturated monos.

323 ► **Corollary 32** (functors induced by saturated specifications). *An iso-coherent, saturated
 324 specification \mathcal{S} induces a functor $\mathbb{S}(\mathcal{S}) : \mathbf{PIDag}(\mathcal{L}) \rightarrow \mathbf{SSPath}(\mathcal{L})$.*

23:10 A Categorical Account for the Specification of Replicated Data Type

325 As is the case for the category of sets and injective functions, $\mathbf{PIDag}(\mathcal{L})$ lacks pushouts.
 326 However, we have an easy way out via the inclusion functor into $\mathbf{PDag}(\mathcal{L})$.

327 ► **Lemma 33** (Mono of \mathbf{PDag}). *Pushouts in \mathbf{PDag} preserve monos.*

328 Thus in the following we say that a functor $\mathbb{F} : \mathbf{PIDag}(\mathcal{L}) \rightarrow \mathbf{SPath}(\mathcal{L})$ weakly preserves
 329 binary pushout (and in fact, finite colimits) if any commuting square in $\mathbf{PIDag}(\mathcal{L})$ that is a
 330 pushout (via the inclusion functor) in $\mathbf{PDag}(\mathcal{L})$ is mapped by \mathbb{F} to a pushout in $\mathbf{SPath}(\mathcal{L})$.

331 ► **Theorem 34.** *Let \mathcal{S} be an iso-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.*

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

335 We can now move to the completeness part.

336 ► **Theorem 35.** *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 an iso-coherent specification $\mathcal{S}(\mathbb{F})$.*

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

$$\begin{array}{ccc}
 \mathbf{G}|_{[e_1] \cap [e_2]} & \hookrightarrow & \mathbf{G}|_{[e_2]} \\
 \downarrow & & \downarrow \\
 \mathbf{G}|_{[e_1]} & \hookrightarrow & \mathbf{G}|_{[e_1] \cup [e_2]}
 \end{array}$$

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

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

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

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

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

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

349 Iso-coherence follows from \mathbb{F} being a functor, hence preserving isos. ◀

350 Furthermore, the two constructions are inverse to each other.

351 ► **Proposition 36.** *We have $\mathbb{M}(\mathcal{S}(\mathbb{F})) \simeq \mathbb{F}$.*

352 **Proof.** For notational convenience, we denote $\mathbb{M}(\mathcal{S}(\mathbb{F}))$ by \mathbb{M}' . We will show the existence
 353 of a natural isomorphism $\varphi: \mathbb{M}' \Rightarrow \mathbb{F}$. By definition, we have $\mathbb{M}'(\mathbb{G}) = \mathcal{S}(\mathbb{F})(\mathbb{G}) = \mathbb{F}(\mathbb{G})$,
 354 therefore we can define $\varphi_{\mathbb{G}} = \text{Id}_{\mathbb{F}(\mathbb{G})}$. We need to prove that it is natural, which in this
 355 case amounts to show $\mathbb{M}'(\mathbf{f}) = \mathbb{F}(\mathbf{f})$, for $\mathbf{f}: \mathbb{G} \rightarrow \mathbb{G}'$ in $\mathbf{PIDag}(\mathcal{L})$. This follows from
 356 $\mathbb{M}'(\mathbf{f})$ and $\mathbb{F}(\mathbf{f})$ having the same underlying function between events, namely the inclusion
 357 $(\mathcal{E}_{\mathbb{F}(\mathbb{G})}, \lambda_{\mathbb{F}(\mathbb{G})}) \rightarrow (\mathcal{E}_{\mathbb{F}(\mathbb{G}')} , \lambda_{\mathbb{F}(\mathbb{G}')})$. \blacktriangleleft

358 We can sharpen the result above by removing the on-the-nose requirement $\mathbb{F}(\mathbb{G}) \subseteq \mathbb{P}(\mathcal{E}_{\mathbb{G}}, \lambda_{\mathbb{G}})$.
 359 To this end, we need to further constraint the class of functors. First, we consider the effect
 360 of top extension on sets of paths.

361 **► Definition 37.** Let $\mathcal{X} \subseteq \mathbb{P}(\mathcal{E}, \lambda)$ be a set of paths and $\ell \in \mathcal{L}$ a label. Its top extension is
 362 defined as $\{\mathbb{P}^{\ell} \mid \mathbb{P} \in \mathcal{X}\}$. Its saturated top extension is defined as $\text{sat}(\mathbb{P}, \mathbf{f})$ for $\mathbf{f}: (\mathcal{E}, \lambda) \rightarrow$
 363 $(\mathcal{E}_{\top}, \lambda[\top \mapsto \ell])$ the obvious inclusion.

364 We say that \mathbb{F} preserves top extensions if it maps top extensions (of dags) to top extensions (of
 365 paths). We can now state two additional instances of Theorem 35. We call *topological* those
 366 specifications such that $\mathcal{S}(\mathbb{G}) \subseteq \{\mathbb{P} \mid \prec_{\mathbb{G}} \subseteq \leq_{\mathbb{P}}\}$. In other words, a topological specification
 367 maps a dag \mathbb{G} to paths that are topological sorts of \mathbb{G} .

368 **► Proposition 38.** Let $\mathbb{F}: \mathbf{PIDag}(\mathcal{L}) \rightarrow \mathbf{SPath}(\mathcal{L})$ that preserves top extensions. If \mathbb{F}
 369 weakly preserves finite colimits and preserves binary pullbacks, it induces an iso-coherent
 370 topological specification $\mathcal{S}(\mathbb{F})$.

371 Finally, we consider the sub-category of $\mathbf{SSPath}(\mathcal{L})$ of saturated ps-morphisms,

372 **► Proposition 39.** Let $\mathbb{F}: \mathbf{PIDag}(\mathcal{L}) \rightarrow \mathbf{SSPath}(\mathcal{L})$ be a functor that preserves top exten-
 373 sions. If \mathbb{F} weakly preserves finite colimits and preserves binary pullbacks, it induces an
 374 iso-coherent saturated specification $\mathcal{S}(\mathbb{F})$.

375 8 Conclusions

376 In this paper we have provided a functorial characterisation of RDTs specifications. Our
 377 starting point is the denotational approach proposed in [6], in which RDTs specifications are
 378 associated with those functions mapping visibility graphs into sets of admissible arbitrations
 379 that are also saturated and coherent. In this work, we consider the category $\mathbf{PDag}(\mathcal{L})$ that
 380 has labelled, acyclic graphs as objects and pr-morphisms as arrows for representing visibility
 381 graphs. We equip $\mathbf{PDag}(\mathcal{L})$ with operators that model the evolution of visibility graphs and
 382 we show that monic arrows in $\mathbf{PDag}(\mathcal{L})$ can be obtained as pushouts. We call $\mathbf{PIDag}(\mathcal{L})$
 383 the full-subcategory of acyclic graphs and monic pr-morphisms. For arbitrations, we take
 384 $\mathbf{SPath}(\mathcal{L})$, which is the category of sets of labelled, total orders and ps-morphisms. Then,
 385 we show that each coherent specification mapping isomorphic graphs into isomorphic set of
 386 paths (i.e., iso-coherent) induces a functor $\mathbb{M}(\mathcal{S}): \mathbf{PIDag}(\mathcal{L}) \rightarrow \mathbf{SPath}(\mathcal{L})$. Conversely, we
 387 prove that a functor $\mathbb{F}: \mathbf{PIDag}(\mathcal{L}) \rightarrow \mathbf{SPath}(\mathcal{L})$ that preserves finite colimits and binary
 388 pullbacks induces an iso-coherent specification $\mathcal{S}(\mathbb{F})$. Moreover, $\mathbb{M}(\mathcal{S})$ and $\mathcal{S}(\mathbb{F})$ are shown
 389 to be inverse of each other.

390 We believe that our characterisation of RDTs provides an ideal setting for the development
 391 of techniques for handling RDT composition. Our long term goal is to equip RDT specific-
 392 ations with a set of operators that enable us to specify and reason about complex RDTs
 393 compositionally, i.e., in terms of constituent parts. We aim to provide a uniform formal
 394 treatment of compositional approaches such as those proposed in [1, 10, 12].

395 — **References** —

- 396 **1** Baquero, C., Almeida, P.S., Cunha, A., Ferreira, C., et al.: Composition in state-based
397 replicated data types. *Bulletin of EATCS* **3**(123) (2017)
- 398 **2** Bouajjani, A., Enea, C., Hamza, J.: Verifying eventual consistency of optimistic replication
399 systems. In: *ACM SIGPLAN Notices*. vol. 49(1), pp. 285–296. ACM (2014)
- 400 **3** Burckhardt, S., Gotsman, A., Yang, H.: Understanding eventual consistency. Tech. Rep.
401 MSR-TR-2013-39, Microsoft Research (2013)
- 402 **4** Burckhardt, S., Gotsman, A., Yang, H., Zawirski, M.: Replicated data types: specification,
403 verification, optimality. In: Jagannathan, S., Sewell, P. (eds.) *POPL 2014*. pp. 271–284. ACM
404 (2014)
- 405 **5** Cerone, A., Bernardi, G., Gotsman, A.: A framework for transactional consistency models
406 with atomic visibility. In: Aceto, L., de Frutos-Escrig, D. (eds.) *CONCUR 2015*. LIPIcs, vol. 42.
407 Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2015)
- 408 **6** Gadducci, F., Melgratti, H., Roldán, C.: A denotational view of replicated data types. In:
409 *International Conference on Coordination Languages and Models*. pp. 138–156. Springer (2017)
- 410 **7** Gadducci, F., Melgratti, H., Roldán, C.: On the semantics and implementation of replicated
411 data types. *Science of Computer Programming* **167**, 91–113 (2018)
- 412 **8** Gilbert, S., Lynch, N.: Brewer’s conjecture and the feasibility of consistent, available, partition-
413 tolerant web services. *SIGACT News* **33**(2), 51–59 (Jun 2002)
- 414 **9** Gotsman, A., Burckhardt, S.: Consistency models with global operation sequencing and their
415 composition. In: *LIPIcs-Leibniz International Proceedings in Informatics*. vol. 91. Schloss
416 Dagstuhl-Leibniz-Zentrum fuer Informatik (2017)
- 417 **10** Gotsman, A., Yang, H.: Composite replicated data types. In: Vitek, J. (ed.) *ESOP 2015*.
418 LNCS, vol. 9032, pp. 585–609. Springer (2015)
- 419 **11** Kaki, G., Earanky, K., Sivaramakrishnan, K., Jagannathan, S.: Safe replication through
420 bounded concurrency verification. *Proceedings of the ACM on Programming Languages*
421 **2**(OOPSLA), 164 (2018)
- 422 **12** Leijnse, A., Almeida, P.S., Baquero, C.: Higher-order patterns in replicated data types. In:
423 *Proceedings of the 6th Workshop on Principles and Practice of Consistency for Distributed*
424 *Data*. p. 5. ACM (2019)
- 425 **13** Shapiro, M., Preguiça, N., Baquero, C., Zawirski, M.: Conflict-free replicated data types. In:
426 Défago, X., Petit, F., Villain, V. (eds.) *SSS 2011*. LNCS, vol. 6976, pp. 386–400. Springer
427 (2011)
- 428 **14** Sivaramakrishnan, K.C., Kaki, G., Jagannathan, S.: Declarative programming over eventually
429 consistent data stores. In: Grove, D., Blackburn, S. (eds.) *PLDI 2015*. pp. 413–424. ACM
430 (2015)

431 **A** **Proof of Lem. 2.**

432 **Proof.**

433 For \Leftarrow), assume that f is a pr-morphism, then

1. By definition of pr-morphism

$$f(s) \gamma f(s') \text{ implies } \exists \bar{s} \in E. \bar{s} \rho s' \wedge f(s) = f(\bar{s})$$

434 Since f is injective, $\bar{e} = e$ holds, and hence $e \rho e'$.

2. Let $\mathcal{T} = \bigcup_{e \in E} f(e)$. We want to show that

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

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

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

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

436 For \Rightarrow), assume that 1) and 2) hold. Take $e \in E$ and $t \in T$. If $t \gamma f(e)$, then there exists
437 $e' \in E$ such that $t = f(e')$ because of (2). By (1) holds, $f(e') \gamma f(e)$ implies $e' \rho e$. \blacktriangleleft