



ELSEVIER

Contents lists available at ScienceDirect

Science of Computer Programming

www.elsevier.com/locate/scico



On the semantics and implementation of replicated data types

Fabio Gadducci^a, Hernán Melgratti^{b,c}, Christian Roldán^{b,*}^a Università di Pisa, Dipartimento di Informatica, Pisa, Italy^b Universidad de Buenos Aires, Facultad de Ciencias Exactas y Naturales, Departamento de Computación, Buenos Aires, Argentina^c CONICET-Universidad de Buenos Aires, Instituto de Investigación en Ciencias de la Computación (ICC), Buenos Aires, Argentina

ARTICLE INFO

Article history:

Received 2 November 2017

Received in revised form 7 May 2018

Accepted 20 June 2018

Available online xxxx

Keywords:

Replicated data types

Specification

Implementation correctness

ABSTRACT

Replicated data types (RDTs) concern the specification and implementation of data structures handled by replicated data stores, i.e., distributed data stores that maintain copies of the same data item on multiple devices. A distinctive feature of RDTs is that the behaviour of an operation depends on the state of the replica over which it performs, and hence, its result may differ from replica to replica. Abstractly, RDTs are specified in terms of two relations, *visibility* and *arbitration*. The former establishes whether an operation observes the effects of the execution of another operation, the latter is a total order on operations used to resolve conflicts between operations executed concurrently over different replicas. Traditionally, an operation of an RDT is specified as a function mapping a visibility and an arbitration into the expected result of the operation. This paper recasts such standard approaches into a denotational framework in which a data type is a function mapping visibility into admissible arbitrations. This characterisation provides a more abstract view of RDTs that (i) highlights some implicit assumptions shared in operational approaches to specification; (ii) accommodates underspecification and refinement; (iii) enables a direct characterisation of the correct implementations of an RDT in terms of a simulation relation between the states of a concrete implementation and of the abstract one determined by the specification.

© 2018 Published by Elsevier B.V.

1. Introduction

Distributed systems replicate their state over different nodes in order to satisfy several non-functional requirements, such as performance, availability, and reliability. It then becomes crucial to keep a consistent view of the replicated data. However, this is a challenging task because consistency is in conflict with two common requirements of distributed applications: *availability* (every request is eventually executed) and tolerance to network *partitions* (the system operates even in the presence of failures that prevent communication among components). In fact, it is impossible for a system to simultaneously achieve strong Consistency, Availability and Partition tolerance [1]. Since many domains cannot renounce availability or avoid network partitions, developers need to cope with weaker notions of consistency by allowing, e.g., replicas to (temporarily) exhibit some discrepancies, as long as they eventually converge to the same state.

This setting challenges the way in which data is specified: states, state transitions and return values should account for the different views that a data item may simultaneously have. Consider a data type `Register`: a memory cell that is read and updated by, respectively, operations `rd` and `wr`. In a replicated scenario, the value obtained when reading a register

* Corresponding author.

E-mail addresses: fabio.gadducci@unipi.it (F. Gadducci), hmelgra@dc.uba.ar (H. Melgratti), croidan@dc.uba.ar (C. Roldán).




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

after two concurrent updates $wr(0)$ and $wr(1)$ (i.e., updates taking place over different replicas) is affected by the way in which updates propagate among replicas: the result might be (i) undefined (when the read is performed over a third replica that has not received any of the updates), (ii) 0, or (iii) 1. Basically, the return value depends on the updates that are seen by that read operation. Choosing the return value is straightforward if a read sees just one update, less so if a read is performed over a replica that knows of both updates, since all replicas should consistently pick the same value among the available ones. A common strategy for registers is that the *last-write wins*: the last update is chosen when several concurrent updates are observed. This strategy implicitly assumes that all the events in a system can be arranged in a total order. Several recent approaches focus on the operational specification of replicated data types [2–9]. Usually, the specification describes the meaning of an operation in terms of two relations among events: *visibility*, which explains the causes for each result, and *arbitration*, which totally orders events. Consider the visibility relation V in Fig. 1a and the arbitrations A_1 and A_2 in Fig. 1b and Fig. 1c, respectively. The meaning of rd is such that $rd(V, A_1) = 1$ and $rd(V, A_2) = 0$. We remark that operational approaches require specifications to be *functional*: for every operation, visibility and arbitration, there is exactly one return value. In this way operational specifications commit to concrete policies for resolving conflicts.

This work aims at putting on firm grounds the operational approaches for RDTs by giving them a purely functional description. In our view, RDTs are functions that map visibility graphs (i.e., configurations) into sets of admissible arbitrations, i.e., all the executions that generate a particular configuration. In this setting, a configuration mapped to an empty set of admissible arbitrations stands for an unreachable configuration, i.e., a configuration that cannot be explained in terms of any arbitration. We rely on such an abstract view of RDTs to highlight some of the implicit assumptions shared by most of the operational approaches. In particular, we characterise operational approaches, such as [2,3], as those specifications that satisfy three properties: besides the evident requirement of being (locally) *functional* (i.e., deterministic and total), they must be *coherent* (i.e., larger states are explained as the composition of smaller ones) and *saturated* (e.g., an operation that has not been seen by any other operation can be arbitrated in any position, even before the events that it sees). We show this inclusion to be strict and discuss some interesting cases that do not fall in this class. Moreover, we show that our formulation elegantly accounts for underspecification and refinement, which are standard notions in data type specification.

The notion of implementation correctness, which is central to the theory of abstract data types [10–12], relates the expected behaviour of a family of operations as defined by a specification with the one that is provided by a more concrete realisation. In a replicated scenario, such concrete realisations consist of several replicas that keep their own local state and propagate changes asynchronously. On the one hand, we assume that the behaviour of an implementation is given in terms of two *labelled transition systems* (LTS): one that describes a single replica and another, which is obtained by composition, that accounts for the joint behaviour of several replicas. Technically, this is achieved by providing a composition operator over LTSs that reflects the adopted communication model. On the other hand, we note that our specifications have an implicit operational interpretation, which describes the expected behaviours of a single replica and of the composition of several replicas. Technically, each specification induces two LTSs: one, called *one-replica*, prescribes the behaviour of a single replica, and another, called *multi-replica*, defines the behaviour of multiple replicas. Then, implementation correctness is defined in terms of simulation relations between the LTSs associated with an implementation, i.e., a replica or a set of replicas, with the LTS corresponding to the specification, i.e., one-replica or multi-replica. We show that implementation correctness is preserved under standard parallel composition (synchronous or asynchronous buffered communication). Consequently, in order to show that an implementation is correct, we only need to show that a single replica is correct. We illustrate the approach with the implementation of a few well-known RDTs.

The paper has the following structure. Section 2 introduces the basic definitions concerning labelled directed acyclic graphs. Section 3 discusses our functional mechanism for the presentation of Replicated Data Types. Section 4 compares our proposal with the classical operational one [4]. Section 5 studies the correctness of the replicated data types implementations with respect to our specifications. Finally, in the closing section we draw some conclusions, discuss related works, and highlight further developments.

This paper is a revised and extended version of [13]. We enrich our previous work by providing an approach to assess whether an implementation of an RDT on top of several concurrent replicas is correct (the material in § 5 is completely new to this paper). In addition, we provide full proofs of already published results.

2. Labelled directed acyclic graphs

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




Fig. 2. Two simple LDAGs and two paths.

Definition 1 (Labelled directed acyclic graph). A Labelled Directed Acyclic Graph (LDAG) over a set of labels \mathcal{L} is a triple $G = \langle \mathcal{E}_G, \prec_G, \lambda_G \rangle$ such that \mathcal{E}_G is a set of events, $\prec_G \subseteq \mathcal{E}_G \times \mathcal{E}_G$ is a binary relation whose transitive closure is a strict partial order, and $\lambda_G : \mathcal{E}_G \rightarrow \mathcal{L}$ is a labelling function. An LDAG G is a path if \prec_G is a strict total order.

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

Example 1. Consider the set $\mathcal{L} = \{\langle rd, 0 \rangle, \langle rd, 1 \rangle, \langle wr(0), ok \rangle, \langle wr(1), ok \rangle\}$ of labels that describe the operations of a 1-bit register. Each label is a tuple $\langle op, rv \rangle$ where op denotes an operation and rv its return value. For homogeneity, we associate the return value ok to every write operation. Now, consider the LDAG over \mathcal{L} that is defined as $G_1 = \{\langle e_1, e_2, e_3 \rangle, \prec, \lambda\}$, where $\prec = \{\langle e_1, e_3 \rangle, \langle e_2, e_3 \rangle\}$ and λ is such that $\lambda(e_1) = \langle wr(0), ok \rangle$, $\lambda(e_2) = \langle wr(1), ok \rangle$, and $\lambda(e_3) = \langle rd, 0 \rangle$. A graphical representation of G_1 is provided in Fig. 2a. Note that we do not depict the events and just write instead the corresponding labels when no confusion arises. A representation of the LDAG G_2 , where \prec_{G_2} is empty, is in Fig. 2b. Neither G_1 nor G_2 is a path because they are not total orders. P_1 in Fig. 2c is an LDAG that is also a path. Hereinafter we use undirected arrows when depicting paths and avoid drawing transitions that are obtained by transitivity, as shown in Fig. 2d. All LDAGs in Fig. 2 belong to $\mathbb{G}(\mathcal{L})$, but only P_1 is in $\mathbb{P}(\mathcal{L})$.

2.1. LDAG operations

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

Definition 2 (Restriction and extension). Let $G = \langle \mathcal{E}, \prec, \lambda \rangle$ and $\mathcal{E}' \subseteq \mathcal{E}$. We define

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

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

Example 2. Consider the LDAGs G_1 and G_2 depicted in Fig. 2a and Fig. 2b, respectively. Then, $G_2 = G_1|_{\prec e_3}$ and G_1 is isomorphic (as a graph) to $G_2^{\langle rd, 0 \rangle}$. Indeed, G_1 can be obtained from G_2 by adding a new node labelled by $\langle rd, 0 \rangle$, in such a way that the new node is related with every node in G_2 via \prec .

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

Definition 3 (Product). Let $\mathcal{X} = \{\langle \mathcal{E}_i, \prec_i, \lambda_i \rangle\}_i$ be a set of paths such that $\forall e, i, j. e \in \mathcal{E}_i \cap \mathcal{E}_j$ implies $\lambda_i(e) = \lambda_j(e)$. The product of \mathcal{X} is

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

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




Fig. 3. Product between two paths.

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

Fig. 4. Counter specification.

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

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

3. Specifications

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

Definition 4 (Specification). A specification \mathcal{S} is a function $\mathcal{S} : \mathbb{G}(\mathcal{L}) \rightarrow 2^{\mathbb{P}(\mathcal{L})}$ such that $\mathcal{S}(\varepsilon) = \{\varepsilon\}$ and $\forall G. \mathcal{S}(G) \in 2^{\mathbb{P}(\mathcal{E}_G, \lambda_G)}$.

A specification \mathcal{S} maps an LDAG (i.e., a visibility relation) to a set of paths (i.e., its admissible arbitrations). Note that $P \in \mathcal{S}(G)$ is a path over \mathcal{E}_G , and hence a total order of the events in G . However, we do not require P to be a topological ordering of G , i.e., $\prec_G \subseteq \prec_P$ may not hold. Although some specification approaches consider only arbitrations that include visibility [6, 7], our definition accommodates also presentations, such as [3,4], in which arbitrations may not preserve visibility. Since our approach is independent of this choice, we adopted the most liberal presentation. We also remark that it could be the case that $\mathcal{S}(G) = \emptyset$, which means that \mathcal{S} forbids the configuration G (more details in Example 4 below). For technical convenience, we impose $\mathcal{S}(\varepsilon) = \{\varepsilon\}$ and disallow $\mathcal{S}(\varepsilon) = \emptyset$: a specification cannot forbid the empty configuration, which denotes the initial state of a data type.

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

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

$$\begin{array}{c}
 P \in \mathcal{S}_{Ctr}(G) \\
 \text{iff} \\
 \forall e \in \mathcal{E}_G. \forall k. \lambda(e) = \langle rd, k \rangle \text{ implies } k = \#\{e' \mid e' \prec_G e \text{ and } \lambda(e') = \langle inc, ok \rangle\}
 \end{array}$$

A visibility graph G has admissible arbitrations (i.e., $\mathcal{S}_{Ctr}(G) \neq \emptyset$) only when each event e in G labelled by `rd` has a return value k that matches the number of increments anteceding e in G . We illustrate two cases for the definition of \mathcal{S}_{Ctr} in Fig. 4. While the configuration in Fig. 4a has admissible arbitrations, the one in Fig. 4b has not, because the unique event labelled by `rd` returns 0 when it is actually preceded by an observed increment. In other words, an execution is not allowed to generate such a visibility graph. We remark that \mathcal{S}_{Ctr} imposes no constraint on the ordering \prec_P . In fact, a path $P \in \mathcal{S}_{Ctr}(G)$ does not need to be a topological ordering of G as, for instance, the rightmost path in the set of Fig. 4a.

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

$$\begin{aligned}
& P \in \mathcal{S}_{lwwR}(G) \\
& \text{iff} \\
\forall e \in \mathcal{E}_G. & \left\{ \begin{array}{l} \lambda(e) = \langle rd, \perp \rangle \text{ implies } \forall e' \prec_G e. \forall k. \lambda(e') \neq \langle wr(k), ok \rangle \\ \forall k. \lambda(e) = \langle rd, k \rangle \text{ implies } \exists e' \prec_G e. \lambda(e') = \langle wr(k), ok \rangle \text{ and} \\ \forall e'' \prec_G e. e' \prec_P e'' \text{ implies } \forall k'. \lambda(e'') \neq \langle wr(k'), ok \rangle \end{array} \right.
\end{aligned}$$

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

Example 6 (Generic Register). We now define a `Generic Register` that does not commit to a particular strategy for resolving conflicts.

$$\begin{aligned}
& P \in \mathcal{S}_{gR}(G) \\
& \text{iff} \\
\forall e \in \mathcal{E}_G. & \left\{ \begin{array}{l} \lambda(e) = \langle rd, \perp \rangle \text{ implies } \forall e' \prec_G e. \forall k. \lambda(e') \neq \langle wr(k), ok \rangle \\ \forall k. \lambda(e) = \langle rd, k \rangle \text{ implies } \exists e' \prec_G e. \lambda(e') = \langle wr(k), ok \rangle \text{ and} \\ \forall e''. \forall k''. \lambda(e'') = \langle rd, k'' \rangle \text{ and } - \prec_G e = - \prec_G e'' \text{ implies } k = k'' \end{array} \right.
\end{aligned}$$

As in Example 5, the return value of a read corresponds to a written value seen by that read, but the specification does not determine which value should be chosen. We require instead that all read operations with the same causes (i.e., $- \prec_G e = - \prec_G e'$) have the same result. Since this condition has to be satisfied by any admissible configuration G , it ensures convergence. Requiring convergence seems meaningful since we are identifying some minimal conditions ensuring the correctness of a specification without making any assumption on the chosen paths. Indeed, convergence can be proved for what we call *deterministic* specifications (i.e., specifications in which the return value of each operation is uniquely determined by the visibility and arbitration relations, as formally characterised in Section 4.1), whose instances are in e.g. Example 4 and Example 5. Thus, our proposal only apparently disagree with approaches like [3,4], where convergence is ensured automatically since the specifications they consider are implicitly deterministic (as shown formally in Section 4.2).

3.1. Refinement

Refinement is a standard approach in data type specification, which allows for a hierarchical organisation that goes from abstract descriptions to concrete implementations. The main benefit of refinement relies on the fact that applications can be developed and reasoned about in terms of abstract data types, which hide implementation details and leave some freedom for the implementation. Consider the specification \mathcal{S}_{gR} of the `Generic Register` introduced in Example 6, which only requires a policy for conflict resolution that ensures convergence. On the contrary, the specification \mathcal{S}_{lwwR} in Example 5 explicitly states that concurrent updates must be resolved by adopting the last-write-wins policy. Since the latter policy ensures convergence, we would like to think about \mathcal{S}_{lwwR} as a refinement of \mathcal{S}_{gR} . We characterise refinement in our setting as follows.

Definition 5 (Refinement). Let $\mathcal{S}_1, \mathcal{S}_2$ be specifications. We say that \mathcal{S}_1 refines \mathcal{S}_2 and write $\mathcal{S}_1 \sqsubseteq \mathcal{S}_2$ if $\forall G. \mathcal{S}_1(G) \subseteq \mathcal{S}_2(G)$.

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

Example 8. Consider the data type `Set`, which provides (among others) the operations `add`, `rem` and `lookup` for respectively adding, removing, and examining the elements within a set. Different alternatives have been proposed in the literature for resolving conflicts in the presence of concurrent additions and removals of elements (see [14] for a detailed discussion). We illustrate two possible alternatives by considering the execution scenario depicted in Fig. 5, for P a topological ordering of G . A reasonable semantics for `lookup` over G and P would fix the return value S , which should contain all the elements in the set, as one of the following two values: \emptyset or $\{1\}$. In fact, under the last-write-wins policy, the specification prescribes that `lookup` returns $\{1\}$ in this scenario. Differently, the strategy of $2P\text{-Set}$ ¹ establishes that the result is \emptyset .

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

¹ In $2P\text{-Set}$, the addition of elements that have been previously removed has no effect.




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

$P \in \mathcal{S}_{\text{Set}}(G)$ iff $\forall e \in \mathcal{E}_G. \forall S \in 2^{\mathbb{N}}. \lambda(e) = \langle \text{lookup}, S \rangle$ implies $B_e \subseteq S \subseteq A_e$ and $\text{Conv}_{e,S}$

where

$$\begin{aligned} A_e &= \{k \mid \exists e' \in \mathcal{E}_G. e' \prec_G e \text{ and } \lambda(e') = \langle \text{add}(k), \text{ok} \rangle\} \\ B_e &= A_e \setminus \{k \mid \exists e' \in \mathcal{E}_G. e' \prec_G e \text{ and } \lambda(e') = \langle \text{rem}(k), \text{ok} \rangle\} \\ \text{Conv}_{e,S} &= \forall e' \in \mathcal{E}_G. \forall S' \in 2^{\mathbb{N}}. \lambda(e') = \langle \text{lookup}, S' \rangle \text{ and } - \prec_G e = - \prec_G e' \text{ implies } S = S' \end{aligned}$$

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

Then, a concrete resolution policy such as $2P\text{-Set}$ can be specified as follows

$P \in \mathcal{S}_{2P\text{-Set}}(G)$ iff $\forall e \in \mathcal{E}_G. \forall S \in 2^{\mathbb{N}}. \lambda(e) = \langle \text{lookup}, S \rangle$ implies $S = B_e$

A different policy, called *Or-Set*, states that additions win against remove operations. This policy can be defined as follows

$P \in \mathcal{S}_{\text{Or-Set}}(G)$ iff $\forall e \in \mathcal{E}_G. \forall S \in 2^{\mathbb{N}}. \lambda(e) = \langle \text{lookup}, S \rangle$ implies $S = C_e$

where

$$C_e = \{k \mid \exists e' \in \mathcal{E}_G. e' \prec_G e \text{ and } \lambda(e') = \langle \text{add}(k), \text{ok} \rangle \text{ and } \forall e'' . e' \prec_G e'' \prec_G e \text{ implies } \lambda(e'') \neq \langle \text{rem}(k), \text{ok} \rangle\}.$$

It is immediate to note that $\mathcal{S}_{2P\text{-Set}}$ is a refinement of \mathcal{S}_{Set} . It can be also noticed that $B_e \subseteq C_e \subseteq A_e$. Consequently, $\mathcal{S}_{\text{Or-Set}}$ is also a refinement of \mathcal{S}_{Set} .

3.2. Classes of specifications

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

Definition 6 (*Past-coherent specification*). Let S be a specification. We say that S is past-coherent (briefly, coherent) if

$$\forall G. \mathcal{S}(G) = \bigotimes_{e \in \mathcal{E}_G} \mathcal{S}(G|_{-\prec^* e})$$

A past-coherent specification S is such that the arbitrations for any configuration G (i.e., the set of paths $\mathcal{S}(G)$) can be obtained by composing the arbitrations associated with all its sub-configurations $G|_{-\prec^* e}$.

Example 9. The specifications in Example 4, Example 5 and Example 6 are all coherent, because their definitions are in terms of restrictions of the LDAGs. This can be checked by application of Definition 6. Consider e.g. the specification of the data type *Counter* in Example 4. Hence, if $P \in \mathcal{S}_{\text{Ctr}}(G)$ then

$$\forall e \in \mathcal{E}_G. \forall k. \lambda(e) = \langle \text{rd}, k \rangle \text{ implies } k = \#\{e' \mid e' \prec_G e \text{ and } \lambda(e') = \langle \text{inc}, \text{ok} \rangle\}$$

Let $\bar{e} \in \mathcal{E}_G$ and $G_{\bar{e}} = G|_{-\prec^* \bar{e}}$. Clearly the property above holds $\forall e \in \mathcal{E}_{G_{\bar{e}}}$. Thus $P|_{-\prec^* \bar{e}} \in \mathcal{S}_{\text{Ctr}}(G_{\bar{e}}) = \mathcal{S}_{\text{Ctr}}(G|_{-\prec^* \bar{e}})$ and consequently $P \in \bigotimes_{\bar{e} \in \mathcal{E}_G} \mathcal{S}_{\text{Ctr}}(G|_{-\prec^* \bar{e}})$.

Conversely, let $P \in \bigotimes_{\bar{e} \in \mathcal{E}_G} \mathcal{S}_{\text{Ctr}}(G|_{-\prec^* \bar{e}})$ and $\bar{e} \in \mathcal{E}_G$. Then $P|_{-\prec^* \bar{e}} \in \mathcal{S}_{\text{Ctr}}(G_{\bar{e}})$, hence if $\lambda(\bar{e}) = \langle \text{rd}, k \rangle$ then $k = \#\{e' \mid e' \prec_{G_{\bar{e}}} \bar{e} \text{ and } \lambda(e') = \langle \text{inc}, \text{ok} \rangle\}$. By definition of $G_{\bar{e}}$, $k = \#\{e' \mid e' \prec_G \bar{e} \text{ and } \lambda(e') = \langle \text{inc}, \text{ok} \rangle\}$ and consequently $P \in \mathcal{S}_{\text{Ctr}}(G)$.

Now consider the specification \mathcal{S} defined such that the equalities in Fig. 6 hold. \mathcal{S} is not coherent because the arbitrations for the LDAG in Fig. 6b should contain all the interleavings for the paths associated with its sub-configurations, as depicted in Fig. 6a.

$$\begin{array}{l}
\mathcal{S}(\langle rd, 0 \rangle) = \{ \langle rd, 0 \rangle \} \\
\mathcal{S}(\langle inc, ok \rangle) = \{ \langle inc, ok \rangle \} \\
\text{(a)}
\end{array}
\qquad
\mathcal{S}(\langle rd, 0 \rangle \langle inc, ok \rangle) = \left\{ \begin{array}{c} \langle rd, 0 \rangle \\ | \\ \langle inc, ok \rangle \end{array} \right\}
\text{(b)}$$

Fig. 6. A non-coherent specification.

$$\begin{array}{l}
\mathcal{S}(\langle inc, ok \rangle) = \{ \langle inc, ok \rangle \} \\
\mathcal{S}(\langle rd, 0 \rangle) = \{ \langle rd, 0 \rangle \}
\end{array}
\qquad
\mathcal{S}\left(\begin{array}{c} \langle rd, 0 \rangle \\ \downarrow \\ \langle inc, ok \rangle \end{array}\right) = \left\{ \begin{array}{c} \langle rd, 0 \rangle \\ | \\ \langle inc, ok \rangle \end{array} \right\}$$

Fig. 7. A non-saturated specification.

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

Definition 7 (Path saturation). Let \mathcal{P} be a path and ℓ a label. We write $\text{sat}(\mathcal{P}, \ell)$ for the set of paths obtained by saturating \mathcal{P} with respect to ℓ , defined as follows

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

A path \mathcal{P} that is saturated with a label ℓ generates the set of all paths obtained by placing a new event labelled by ℓ in any position within \mathcal{P} . By analogy, a saturated specification thus extends a computation by adding a new operation that can be arbitrated in any position.

Definition 8 (Saturated specification). Let \mathcal{S} be a specification. We say that \mathcal{S} is saturated if

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

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

4. Replicated data type

In this section we show that our notion of specification can be considered as (and it is actually more general than) a model for the operational description of RDTs proposed in [3,4]. We start by recasting the original definition of RDT (as given in [4, Def. 4.5]) in terms of LDAGs. As hinted in the introduction, the meaning of each operation of an RDT is specified in terms of a context, written C , which is a pair $\langle G, \mathcal{P} \rangle$ such that $\mathcal{P} \in \mathbb{P}(\mathcal{E}_G, \lambda_G)$. We write $\mathbb{C}(\mathcal{L})$ for the set of contexts over \mathcal{L} , and fix a set \mathcal{O} of operations and a set \mathcal{V} of values. Then, the operational description of RDTs in [3,4] can be formulated as follows.

Definition 9 (Replicated data type). A Replicated Data Type (RDT) is a function $\mathcal{F} : \mathcal{O} \times \mathbb{C}(\mathcal{O}) \rightarrow \mathcal{V}$.

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

Example 11. The data type `Counter` introduced in Example 4 is formally specified in [3,4] as follows

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

Given a context $\langle G, \mathcal{P} \rangle$ in $\mathbb{C}(\mathcal{O} \times \mathcal{V})$, we may check whether the value associated with each operation matches the definition of a particular RDT. This notion is known as *return value consistency* [4, Def. 4.8]. In order to relate contexts with and without return values, we use the following notation: given $G \in \mathbb{G}(\mathcal{O} \times \mathcal{V})$, by $\bar{G} \in \mathbb{G}(\mathcal{O})$ we denote the LDAG obtained by projecting the labels of G in the obvious way, i.e., by removing the second component of every label.




Fig. 8. RVAL consistency for \mathcal{F}_{cr} .

Definition 10 (Return value consistent). Let \mathcal{F} be an RDT and $\langle G, P \rangle \in \mathbb{C}(\mathcal{O} \times \mathcal{V})$ a context. We say that \mathcal{F} is *Return Value Consistent* (RVAL) over G and P and write $\text{RVAL}(\mathcal{F}, G, P)$ if $\forall e \in \mathcal{E}_G. \lambda(e) = \langle \text{op}, v \rangle$ implies $\mathcal{F}(\text{op}, \bar{G}|_{\leftarrow e}, \bar{P}|_{\leftarrow e}) = v$.

Moreover, we define

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

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

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

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

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

4.1. Functional specifications

We now focus on the relation between our notion of specification, as introduced in Definition 4, and the operational description of RDTs, as introduced in [3,4] and formalised in Definition 9 in terms of LDAGs. Specifically, we characterise a proper subclass of specifications that precisely correspond to RDTs.

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

Definition 11 (Total specification). Let \mathcal{S} be a specification. We say that \mathcal{S} is total if

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

Intuitively, a specification is total when every operation of the data type can be performed in any state of the computation. Formally, this is stated by considering a context $\langle \bar{G}, \bar{P} \rangle$ as the representation of the state of a computation. We remark that differently from G and P , whose labels are in $\mathcal{O} \times \mathcal{V}$, \bar{G} and \bar{P} have no information about the return value of the operations (i.e., their labels are in \mathcal{O}). Hence, totality says that it is always possible to take an equivalent representation of the state (i.e., $\langle G_1, P_1 \rangle$) instead of $\langle G, P \rangle$ and extend it with an operation op . This is achieved by requiring $\bar{P} \in \overline{\mathcal{S}((G_1)_{\mathcal{E}}^{(\text{op}, v)})} \Big|_{\mathcal{E}_{G_1}}$

for some return value v . Thus $\mathcal{S}((G_1)_{\mathcal{E}}^{(\text{op}, v)})$ is not empty, hence, the specification allows the execution of op over G_1 .

We remark that a total specification does not prevent the definition of an operation that admits more than one return value in certain configurations, i.e., v in Definition 11 does not need to be unique. For instance, consider the *Generic Register* in Example 6, where operation rd may return any of the causally-independent, previously written values. Albeit being total, the specification for rd is not deterministic. On the contrary, a specification is deterministic if an operation executed over a configuration admits at most one return value, as formally stated below.

Definition 12 (Deterministic specification). Let \mathcal{S} be a specification. We say that \mathcal{S} is deterministic if

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

We say that \mathcal{S} is weakly deterministic if the property holds for $\mathcal{E} = \emptyset$.

A less restrictive notion for determinism could allow the result for an added operation to depend also on the given admissible path. We say that a specification \mathcal{S} is *value-deterministic* if

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

We say that a specification is *functional* if it is both deterministic and total.

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

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




Fig. 10. Generic register.

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

Example 14. It is straightforward to check that the specifications in Example 4 and Example 5 are deterministic. For Example 4 we reason as follows. The case for $\text{op} = \text{inc}$ follows immediately because the only possible return value is ok . When $\text{op} = \text{rd}$, from the definition of \mathcal{S}_{Ctr} (Example 4) we conclude that $\mathcal{S}_{\text{Ctr}}(G_{\mathcal{E}}^{(\text{rd}, v)}) \neq \emptyset$ only when $v = \#\{e \mid e \in \mathcal{E} \text{ and } \lambda(e) = \langle \text{inc}, \text{ok} \rangle\}$. Consequently, for any $v' \neq v$, $\mathcal{S}_{\text{Ctr}}(G_{\mathcal{E}}^{(\text{rd}, v')}) = \emptyset$ holds, hence, \mathcal{S}_{Ctr} is deterministic. For Example 5, we can reason analogously.

On the contrary, the specification of the Generic Register in Example 6 is not even value-deterministic. It suffices to consider a configuration G with two different written values, as shown in Fig. 10a. Consider now the two extensions $G^{(\text{rd}, 1)}$ and $G^{(\text{rd}, 2)}$ depicted in Fig. 10b and Fig. 10c and the two paths P_1 and P_2 in Fig. 10d and Fig. 10e. By the definition of \mathcal{S}_{GR} , we can conclude that $P_1 \in \mathcal{S}_{\text{GR}}(G^{(\text{rd}, 1)})$ and $P_2 \in \mathcal{S}_{\text{GR}}(G^{(\text{rd}, 2)})$. The path P in Fig. 10f corresponds to both $P_1|_{\mathcal{E}_G}$ and $P_2|_{\mathcal{E}_G}$. Consequently, $\mathcal{S}_{\text{GR}}(G^{(\text{rd}, 1)})|_{\mathcal{E}_G} \cap \mathcal{S}_{\text{GR}}(G^{(\text{rd}, 2)})|_{\mathcal{E}_G} \neq \emptyset$.

Similarly, Set in Example 8 is not deterministic.

The lemma below states a simple criterion for determinism.

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

$$\forall G_1, G_2. \overline{G_1} = \overline{G_2} \text{ implies } G_1 = G_2 \vee \overline{\mathcal{S}(G_1)} \cap \overline{\mathcal{S}(G_2)} = \emptyset$$

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

We then consider a last property that guarantees some sort of additional locality to the notion of coherence.

Definition 13 (Local specification). Let S be a specification. We say that S is local if

$$\forall G, \mathcal{E}. \exists G_1. \forall \text{op}, v. \overline{G}|_{\mathcal{E}} = \overline{G_1} \wedge \overline{\mathcal{S}(G_{\mathcal{E}}^{(\text{op}, v)})} \Big|_{-\preceq_{\top}} \subseteq \overline{\mathcal{S}(G_1^{(\text{op}, v)})}$$

We say that a specification is *locally functional* if it is both functional and local.

Intuitively, locality states that the admissible paths of an extended configuration are constrained: given a configuration G , its extension with an operation op with respect to the events \mathcal{E} can be explained by the sub-configuration of G that only contains those events, i.e., $\overline{G}|_{\mathcal{E}} = \overline{G_1}$. Such behaviour is given by $\mathcal{S}(G_1^{(\text{op}, v)})$.

As a side remark, it is noteworthy that locality has an impact on determinism.

Lemma 3. Let S be a local and weakly deterministic specification. Then it is deterministic.

4.2. Correspondence between RDTs and specifications

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

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

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

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

Lemma 4. For every RDT \mathcal{F} , $\mathbb{S}(\mathcal{F})$ is coherent, locally functional, and saturated.

The inverse mapping from specifications to RDTs is defined below.

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

$$\mathbb{F}(\mathcal{S})(\text{op}, \bar{G}, \bar{P}) = \nu \text{ if } \exists G_1. \bar{G} = \overline{G_1} \wedge \bar{P} \in \overline{\mathcal{S}(G_1^{\text{op}, \nu})} \Big|_{\mathcal{E}_{G_1}}$$

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

Lemma 5. For every coherent, functional, and saturated specification \mathcal{S} , $\mathbb{F}(\mathcal{S})$ is well-defined.

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

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

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

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

5. On the correct implementation of replicated data types

The previous section was devoted to the proof of correspondence between the novel notion of specification we introduced and a more standard proposal for modelling replicated data types. In this section we argue that our definition is flexible enough for reasoning about the correctness of possible implementations, in terms of the classical notion of *simulation*. More precisely, first we show how a specification naturally gives rise to a labelled transition system (LTS). Then, we consider state-based implementations of replicated data types [3].

5.1. From specifications to labelled transitions systems

We note that specifications have an implicit operational interpretation.

Definition 16 (One-replica LTS). Let \mathcal{S} be a specification. Then, the abstract *one-replica* LTS $\mathcal{T}_{\mathcal{S}}$ has pairs $\langle G, P \rangle$ as states, pairs $\langle \text{op}, \nu \rangle$ as labels, and triples $\langle G, P \rangle \xrightarrow{\ell} \langle G', P' \rangle$ as transitions such that

- $\langle G, P \rangle$ is a state whenever $P \in \mathcal{S}(G)$;
- $\langle G, P \rangle \xrightarrow{\ell} \langle G', P' \rangle$ is a transition whenever $G' = G^\ell$ and $P' \Big|_{\mathcal{E}_G} = P$.

A pair $\langle G, P \rangle$ such that $P \in \mathcal{S}(G)$ abstractly represents an admissible computation according to \mathcal{S} . Moreover, \mathcal{S} allows the extension of such computation with an event labelled by ℓ whenever the specification allows us to extend G and P with a fresh event labelled by ℓ .

Our next step is to show that a specification can be equipped with a notion of parallel composition. But first, we need to consider a suitable kind of morphism.

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

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

Definition 18 (*Past-reflecting morphisms*). An LDAG morphism $f : G_1 \rightarrow G_2$ from G_1 to G_2 is a function $f : \mathcal{E}_{G_1} \rightarrow \mathcal{E}_{G_2}$ such that $\lambda_{G_1} = f; \lambda_{G_2}$ and $e \prec_{G_1} e'$ implies $f(e) \prec_{G_2} f(e')$. An injective LDAG morphism $f : G_1 \rightarrow G_2$ is past-reflecting if

- $f(e) \prec_{G_2} f(e')$ implies $e \prec_{G_1} e'$;
- $\bigcup_{e \in \mathcal{E}_{G_1}} f(e)$ is downward closed.

Hence, past-reflecting injective morphisms $f : G_1 \rightarrow G_2$ are uniquely characterised as such by the properties of the image of \mathcal{E}_{G_1} with respect to G_2 .

Definition 19 (*Compatibility*). Let G_1, G_2 be two LDAGs. They are compatible if

- $e \in \mathcal{E}_{G_1} \cap \mathcal{E}_{G_2}$ implies $\lambda_1(e) = \lambda_2(e)$;
- the injective morphisms $G_i \rightarrow G_1 \cup G_2$ for $i \in \{1, 2\}$ are past-reflecting.

We write $G_1 \sqcup G_2$ to denote the union of compatible LDAGs.

Note that the operation \sqcup over LDAGs is idempotent, associative, and commutative. Also note that if G_1 and G_2 are compatible, then the product $P_1 \otimes P_2$ is correctly labelled for any pair $P_1 \in \mathcal{S}(G_1)$ and $P_2 \in \mathcal{S}(G_2)$ (i.e., events that belongs to both \mathcal{E}_{G_1} and \mathcal{E}_{G_2} have the same label) because each path is built out of the set of events (and the corresponding labels) of the associated LDAG. Consequently, we write $\langle G_1 \sqcup \dots \sqcup G_n, P \rangle$ for a replicated state consisting of n compatible replicas, with $\langle G_i, P|_{\mathcal{E}_{G_i}} \rangle$ denoting the state of the replica i .

Definition 20 (*Multi-replica LTS*). Let \mathcal{S} be a coherent specification. Then, the abstract multi-replica LTS $\mathcal{T}_{\mathcal{S}}$ is generated by the corresponding one-replica LTS and the additional rule below.

(COMP)

$$\frac{\langle G_1, P|_{\mathcal{E}_{G_1}} \rangle \xrightarrow{\ell} \langle G'_1, P'_1 \rangle \quad P' \in P \otimes P'_1}{\langle G_1 \sqcup G_2, P \rangle \xrightarrow{\ell} \langle G'_1 \sqcup G_2, P' \rangle}$$

Rule (COMP) describes the computations associated to a replicated state. Whenever a part of a replicated state (denoted by $\langle G_1, P|_{\mathcal{E}_{G_1}} \rangle$) evolves to $\langle G'_1, P'_1 \rangle$ by performing the action ℓ , then the whole system evolves to a state obtained by composing the part that has not changed (i.e., G_2) and the new state $\langle G'_1, P'_1 \rangle$.

Indeed, the states corresponding to the union of compatible LDAGs recalls parallel composition, and the soundness of the rule with respect to the target is witnessed by the proposition below, ensuring that the transitions of a composite state are in correspondence with the transitions of its components.

Lemma 6. *Let \mathcal{S} be a coherent specification, G_1, G_2 compatible LDAGs, and $P \in \mathcal{S}(G_1 \sqcup G_2)$ a path. If $\langle G_1, P|_{\mathcal{E}_{G_1}} \rangle \xrightarrow{\ell} \langle G'_1, P' \rangle$ then $P \otimes P' \subseteq \mathcal{S}(G'_1 \sqcup G_2)$.*

5.2. Implementing a specification

In this section we address the problem of identifying what a concrete implementation of a replicated data type is. We consider a setting in which a replicated data type is realised on top of a set of replicas, where each replica keeps a local version of the data and clients perform read and write operations over replicas. Roughly speaking, each request made by a client is handled by one of the replicas. In particular, a read request is answered according to the local state of the replica that handles that request. Analogously, any update operation is executed over the local state of a replica. The changes are then propagated asynchronously to the remaining replicas, which will update their own states later on. Since changes are propagated asynchronously and different replicas may be serving concurrent updates, each replica is responsible for resolving conflicts locally. We will focus on state-based implementation of replicated data types, as opposed to operation-based implementation [2]. In a state-based implementation, replicas propagate changes by communicating their own states.

Following the approach in [3], we describe an implementation of a data type in terms of the behaviour of a replica and we assume that all the replicas of an implementation behave the same. We will also consider LTSs as the operational model of replicas. We then let Σ (ranged over by σ, σ_0, \dots) denote the set of possible states of a replica and for a specification \mathcal{S} we use the following set of operations as labels

$$\mathcal{A}_{\mathcal{S}} = \mathcal{L} \cup (\{\text{send}, \text{rcv}\} \times \Sigma) \cup \{\tau\}$$

$$\begin{array}{ll}
\text{(READ)} & \text{(INC)} \\
\frac{k = \sum_{s \in \mathcal{R}} v(s)}{\langle r, v \rangle \xrightarrow{\text{rd}, k} \langle r, v \rangle} & \langle r, v \rangle \xrightarrow{\text{inc}, ok} \langle r, v[r \mapsto v(r) + 1] \rangle \\
\text{(SEND)} & \text{(RCV)} \\
\langle r, v \rangle \xrightarrow{\text{send}(r, v)} \langle r, v \rangle & \langle r, v \rangle \xrightarrow{\text{rcv}(r', v')} \langle r, \max\{v, v'\} \rangle
\end{array}$$

Fig. 11. Implementation of data type COUNTER.

The set \mathcal{A}_S of labels is then built-up from the operations of the data type (i.e., the elements in $\mathcal{L} = \mathcal{O} \times \mathcal{V}$), with additionally two special kinds of operations $\langle \text{send}, \sigma \rangle$ and $\langle \text{rcv}, \sigma \rangle$ that are used by the replicas to synchronise their states and finally the label τ for internal transitions.

The behaviour of a replica of a specification S is given by an LTS \mathcal{C}_S with labels in \mathcal{A}_S that is *total*, i.e., such that $\forall \sigma, \text{op}. \exists v. \sigma \xrightarrow{\text{op}, v}$. Totality ensures that all operations in the implementation of a replicated data type are non-blocking, i.e., a replica is able to perform any operation of the data type at any state.

A complete implementation of a replicated data type is then obtained by putting several replicas in parallel. As a formal counterpart, we rely on a composition operator that prescribes the way in which different replicas synchronise. Its definition depends on the chosen communication model (synchronous, asynchronous, broadcast, ...) and we defer its formal treatment to § 5.4.

Example 15. We consider the data type Counter in Example 4 and the state-based implementation presented in [3], which is defined below

- $\Sigma = \mathcal{R} \times (\mathcal{R} \mapsto \mathbb{N})$, where \mathcal{R} is the set of replicas' identifiers;
- $\mathcal{L} = \{\{\text{inc}, \text{ok}\}\} \cup \{\{\text{rd}\} \times \mathbb{N}\}$ is the set of data type operations;
- \rightarrow is given by the inference rules in Fig. 11.

The states of a replica are pairs of the form $\langle r, v \rangle$, where r is the replica's id and v is a mapping that keeps track of the known increments performed over all replicas, i.e., $v(r')$ is the number of increments performed over the replica r' .

We now comment on the inference rules in Fig. 11. Rule (READ) describes a replica that is handling a client's request for reading the counter. In such case, the replica returns the value k , which corresponds to the total number of increments known from all replicas. This transition does not change the state of the replica. Differently, the state changes when performing an increment, as described by rule (INC). The new state records the fact that r has performed another increment. The change has only local effect and can be propagated later on, by using rule (SEND). A replica also updates its local state when it receives a change propagated by another replica, as described by the rule (RCV). When receiving a message $\langle r', v' \rangle$, the replica updates its local mapping to $\max\{v, v'\}$, which is defined as follows

$$\forall s. \max\{v, v'\}(s) = \max\{v(s), v'(s)\}$$

The natural question is whether the behaviour of the replica in Example 15 correctly implements the specification in Example 4. We address this problem in two steps: first of all, we show that a single replica is a correct implementation of the data type (§ 5.3); then, we analyse the combined behaviour of several replicas (§ 5.4).

5.3. Linking a specification with the behaviour of a replica

In this section we provide a criterion to formally prove that the implementation of a replica is correct. More precisely, the correspondence between the behaviour of a replica and a specification is given as a (weak) simulation relation.

Definition 21 (Implementation correctness). Let S be a specification, \mathcal{T}_S the one-replica (multi-replica) abstract LTS, and \mathcal{C}_S an implementation. Then, an *implementation relation* \mathcal{I}_S is a relation between the states of \mathcal{C}_S and \mathcal{T}_S such that if $(\sigma, \langle G, P \rangle) \in \mathcal{I}_S$ then for any $\sigma', \sigma'', \text{op}$ and v

1. if $\sigma \xrightarrow{\text{op}, v} \sigma'$ then $\exists G', P'$ such that $\langle G, P \rangle \xrightarrow{\text{op}, v} \langle G', P' \rangle$ and $(\sigma', \langle G', P' \rangle) \in \mathcal{I}_S$;
2. if $\sigma \xrightarrow{\text{rcv}, \sigma'} \sigma''$ then $\exists G', P', P''$ such that $(\sigma', \langle G', P' \rangle) \in \mathcal{I}_S$ and $P'' \in P \otimes P'$ and $(\sigma'', \langle G \sqcup G', P'' \rangle) \in \mathcal{I}_S$;
3. if $\sigma \xrightarrow{\text{send}, \sigma'} \sigma$ then $\exists G', P'$ such that $(\sigma', \langle G', P' \rangle) \in \mathcal{I}_S$ and $G \sqcup G' = G$ and $P \otimes P' = \{P\}$;
4. if $\sigma \xrightarrow{\tau} \sigma'$ then $(\sigma', \langle G, P \rangle) \in \mathcal{I}_S$.

We write \sim_S^0 for the largest implementation relation with respect to the abstract one-replica LTS of the specification S , and \sim_S^m for the multi-replica one.

$$\begin{array}{c}
\text{(ADD)} \\
\frac{V' = V[r \mapsto V(r) + 1] \quad W' = W[(k, r) \mapsto V(r) + 1]}{\langle r, V, W \rangle \xrightarrow{\text{add}(k, \text{ok})} \langle r, V', W' \rangle} \\
\\
\begin{array}{cc}
\text{(REM)} & \text{(LOOKUP)} \\
\frac{W' = W[\forall s \in \mathcal{R}. (k, s) \mapsto 0]}{\langle r, V, W \rangle \xrightarrow{\text{rem}(k, \text{ok})} \langle r, V, W' \rangle} & \frac{S = \{k \mid \exists s \in \mathcal{R}. W(k, s) > 0\}}{\langle r, V, W \rangle \xrightarrow{\text{lookup}, S} \langle r, V, W \rangle} \\
\\
\text{(SEND)} & \text{(RECEIVE)} \\
\langle r, V, W \rangle \xrightarrow{\text{send}(r, V, W)} \langle r, V, W \rangle & \langle r, V, W \rangle \xrightarrow{\text{rcv}(r', V', W')} \langle r, \max\{V, V'\}, (V, W) \oplus (V', W') \rangle
\end{array}
\end{array}$$

Fig. 12. Implementation of data type OR-SET

Any pair $(\sigma, \langle G, P \rangle)$ in the relation $\mathcal{I}_{\mathcal{S}}$ establishes that the state σ is explained in terms of the visibility G and the arbitration P , which is admitted by the specification (i.e., $P \in \mathcal{S}(G)$). Moreover, there is a close correspondence between the evolution of σ and $\langle G, P \rangle$, which is stated by the items 1–4. Item 1 predicates on the transitions corresponding to the operations of the data type. Basically, if the replica performs the operation op that produces the result v , then the specification allows G and P to be extended with the corresponding event, which is captured by the transition $\langle G, P \rangle \xrightarrow{\text{op}, v} \langle G', P' \rangle$ formalised in Definition 16 and Definition 20.

Item 2 regards those transitions in which the replica receives updates propagated by other replicas. In a state-based implementation, the content of a message rcv is a state computed by another replica. For this reason, we require the received state σ' to be related to a configuration admitted by the specification, i.e., $(\sigma', \langle G', P' \rangle) \in \mathcal{I}_{\mathcal{S}}$. Moreover, $\langle G', P' \rangle$ must be consistent with the history already seen by the replica, i.e., the common history in $\langle G, P \rangle$ and $\langle G', P' \rangle$ must coincide. This is established by requiring that the union of the two visibilities, i.e., $G \sqcup G'$, and the merge of the two paths, i.e., $P'' \in P \otimes P'$, are defined. Under the above conditions, the state σ'' obtained as the combination of σ and σ' must correspond to the combination of $\langle G, P \rangle$ and $\langle G', P' \rangle$, i.e., $(\sigma'', \langle G \sqcup G', P'' \rangle) \in \mathcal{I}_{\mathcal{S}}$.

Item 3 states that a replica only propagates messages carrying information about its current state. Note that σ' may contain information about some events on the current state. This is formally stated by conditions $G \sqcup G' = G$ and $P \otimes P' = \{P\}$ (and it is going to be useful in § 5.4). Item 4 is self-explanatory.

In the following, we will refer to one-replica or multi-replica correctness of an implementation when the abstract LTS $\mathcal{T}_{\mathcal{S}}$ in the definition above is one-replica or multi-replica, respectively.

Example 16. We can show that the replica defined in Example 15 is a correct one-replica implementation of the data type Counter in Example 4 by showing that the following relation satisfies the conditions stated in Definition 21 (proof details are in Appendix B) with respect to the one-replica abstract LTS

$$\mathcal{I} = \{ \langle \langle r, v \rangle, \langle G, P \rangle \rangle \mid \text{there exists } f : \mathcal{E}_G \rightarrow \mathcal{R} \text{ such that } \forall r \in \mathcal{R}. v(r) = \#\{e \mid f(e) = r \text{ and } \lambda(e) = \langle \text{inc}, \text{ok} \rangle\} \}$$

Example 17. In this example we show that the optimised OR-set implementation in [3] is correct with respect to the specification in Example 8. The implementation of a replica is given by

- $\Sigma = \mathcal{R} \times (\mathcal{R} \mapsto \mathbb{N}) \times ((\mathcal{V} \times \mathcal{R}) \mapsto \mathbb{N})$;
- $\mathcal{L} = \{ \langle \text{add}(k), \text{ok} \rangle, \langle \text{rem}(k), \text{ok} \rangle \mid k \in \mathcal{V} \} \cup \{ \langle \text{lookup} \rangle \times 2^{\mathcal{V}} \}$;
- \rightarrow is given by the inference rules in Fig. 12.

States are triples $\langle r, v, w \rangle$, where $r \in \mathcal{R}$ is the replica's id. The mapping v associates each replica with a version number and means that r is up-to-date with the version $v(r)$ of the replica r . The mapping w indicates for each replica r' and element k the newest version of r' in which the element k has been added, if any (i.e., $w(k, r') = 0$ means that k either has not been added to r' or it has been added and then deleted). It is assumed that $w(k, r) \leq v(r)$ for all r and k .

We now discuss the rules in Fig. 12. Rule (ADD) describes the behaviour of a replica r that handles a request for adding the element k to the set. In this case, r changes its local state by (i) creating a new version, i.e., the entry $v(r)$ is updated to $v(r) + 1$ to reflect that there is a new version for r , and (ii) recording that the element k has been added in the newly created version of r , i.e., the entry $w(k, r)$ is updated to $v(r) + 1$. The removal of the element k is described by rule (REM). In this case, for all $s \in \mathcal{R}$, $w(k, s)$ is set to 0 to indicate the elimination of k . When performing a lookup, the replica r returns the set S of those elements that are present in at least one of the known versions of the replicas (Rule (LOOKUP)).

Rule (SEND) is straightforward. During synchronisation, which is described in rule (RECEIVE), the mappings v and w are updated with the information in the received message. It is assumed that the received mappings v' and w' are consistent, i.e., $w'(k, r_i) \leq v'(r_i)$ for all k . For v , the new state keeps the higher version for each replica, i.e., $\max\{v, v'\}$. The combination

$$\begin{array}{c}
\text{(PARL)} \\
\frac{\sigma_1 \xrightarrow{\ell} \sigma'_1}{\sigma_1 \parallel \sigma_2 \xrightarrow{\ell} \sigma'_1 \parallel \sigma_2}
\end{array}
\qquad
\begin{array}{c}
\text{(COMM)} \\
\frac{\sigma_1 \xrightarrow{\text{send}, \sigma} \sigma'_1 \quad \sigma_2 \xrightarrow{\text{rcv}, \sigma} \sigma'_2}{\sigma_1 \parallel \sigma_2 \xrightarrow{\ell} \sigma'_1 \parallel \sigma'_2}
\end{array}$$

Fig. 13. Synchronous communication of replicas (symmetric rules are omitted).

of \mathbb{W} and \mathbb{W}' is more involved because it handles the conflicts due to concurrent operations over the same element. There is a conflict between \mathbb{W} and \mathbb{W}' for (k, s) when one mapping indicates that k is present in the replica s and the other does not, i.e., $\mathbb{W}(k, s) > 0$ and $\mathbb{W}'(k, s) = 0$ or vice versa. In case of a conflict, the mappings \mathbb{V} and \mathbb{V}' are used to resolve it. The operator \oplus is defined as follows

$$(\mathbb{V}, \mathbb{W}) \oplus (\mathbb{V}', \mathbb{W}')(k, s) = \begin{cases} 0 & \text{IsREM}(\mathbb{W}, \mathbb{W}', \mathbb{V}, k, s) \vee \text{IsREM}(\mathbb{W}', \mathbb{W}, \mathbb{V}', k, s) \\ \max\{\mathbb{W}(k, s), \mathbb{W}'(k, s)\} & \text{otherwise} \end{cases}$$

where $\text{IsREM}(\mathbb{W}, \mathbb{W}', \mathbb{V}, k, s) = (\mathbb{W}(k, s) = 0 \wedge \mathbb{W}'(k, s) \leq \mathbb{V}(s))$ is the predicate characterising the fact that k has been removed from the replica s accordingly to the mappings \mathbb{W} , \mathbb{W}' , and \mathbb{V} .

We show that a replica is correct with respect to the specification in Example 8 by showing that the following relation satisfies the conditions in Definition 21

$$\mathcal{I} = \{ \langle (r, \mathbb{V}, \mathbb{W}), (G, P) \rangle \mid \text{there exists } f : \mathcal{E}_G \rightarrow \mathcal{R} \text{ such that} \\
\forall k. (\exists r \in \mathcal{R}. \mathbb{W}(k, r) > 0 \iff \exists S. S_{\text{Or-Set}}(G^{(\text{lookup}, S)}) \neq \emptyset \wedge k \in S) \}$$

Correctness and refinement. It is straightforward to notice that correctness is preserved by refinement, i.e., if S is a refinement of S' and an implementation I is correct with respect to S , then I is correct with respect to S' . This follows from the fact that $\langle G, P \rangle \xrightarrow{\text{op}, \mathbb{V}} \langle G', P' \rangle$ in \mathcal{S} implies $\langle G, P \rangle \xrightarrow{\text{op}, \mathbb{V}} \langle G', P' \rangle$ in \mathcal{S}' . Consequently, we can conclude that the implementation in Example 17 is also a correct implementation of the non-deterministic specification of S_{Set} introduced in Example 8.

5.4. On the behaviour of multiple replicas

We now address the problem of showing that the parallel composition of several correct replicas is actually a correct implementation of a data type. In this section we focus on coherent specifications and start by considering the standard synchronous communication model, which is defined as follows.

Definition 22 (*Synchronous implementation*). Let \mathcal{S} be a specification and $\mathcal{C}_{\mathcal{S}}$ an implementation. Then the *synchronous extension* $\mathcal{C}_{\mathcal{S}}^{\sigma}$ of $\mathcal{C}_{\mathcal{S}}$ is obtained by closing the set of states with a binary operation \parallel and extending the transition relation with the additional rules in Fig. 13 (where symmetric versions are omitted).

We now want to reach the conclusion that whenever we have an implementation $\mathcal{C}_{\mathcal{S}}$ that is one-replica correct, then the synchronous extension $\mathcal{C}_{\mathcal{S}}^{\sigma}$ is multi-replica correct. The intuition is that the state of the parallel composition of replicas is described in terms of the configurations associated with each of the components. The following result suffices for our purposes, since it states that the simulation for a coherent specification is closed under synchronous parallel composition, and it relies on Lemma 6.

Lemma 7. *Let \mathcal{S} be a coherent specification, $\mathcal{C}_{\mathcal{S}}$ an implementation that is one-replica correct, and σ_0, σ_1 two states of $\mathcal{C}_{\mathcal{S}}$. If $\sigma_i \sim_{\mathcal{S}}^o \langle G_i, P_i \rangle$ for $i \in \{1, 2\}$, then $\sigma_1 \parallel \sigma_2 \sim_{\mathcal{S}}^m \langle G_1 \sqcup G_2, P \rangle$ for any $P \in P_1 \otimes P_2$.*

Then, we obtain immediately the desired result.

Proposition 1. *Let \mathcal{S} be a coherent specification and $\mathcal{C}_{\mathcal{S}}$ an implementation that is one-replica correct. Then, $\mathcal{C}_{\mathcal{S}}^{\sigma}$ is multi-replica correct.*

Example 18. The result above allows us to conclude that the implementation consisting of the parallel composition of replicas behaving as described in Example 15 is correct with respect to the specification of \mathcal{S}_{Ctx} in Example 4, which is coherent (see Example 9). Similarly, we can conclude that the parallel composition of several replicas for OR-Set in Example 17 is a correct implementation.

We remark that different communication models can be accommodated analogously. We may consider the (admittedly simplistic) asynchronous implementation $\mathcal{C}_{\mathcal{S}}^{\sigma}$ obtained by adding a family of operators $|\mathbb{B}$ defined in Fig. 14, where \mathbb{B} denotes a FIFO buffer. Formally, we represent \mathbb{B} as a sequence $\sigma_1 \cdot \dots \cdot \sigma_n$ of states and write ε for the empty sequence. Rule (SENDL)

$$\begin{array}{ccc}
\text{(PARL)} & \text{(SENDL)} & \text{(RECEIVL)} \\
\frac{\sigma_1 \xrightarrow{\ell} \sigma'_1}{\sigma_1 \mid_B \sigma_2 \xrightarrow{\ell} \sigma'_1 \mid_B \sigma_2} & \frac{\sigma_1 \xrightarrow{\text{send}, \sigma} \sigma'_1}{\sigma_1 \mid_B \sigma_2 \xrightarrow{\ell} \sigma'_1 \mid_{B \cdot \sigma} \sigma_2} & \frac{\sigma_1 \xrightarrow{\text{rcv}, \sigma} \sigma'_1}{\sigma_1 \mid_{\sigma \cdot B} \sigma_2 \xrightarrow{\ell} \sigma'_1 \mid_B \sigma_2}
\end{array}$$

Fig. 14. Buffered communication of replicas (symmetric rules are omitted)

says that the state σ sent by a replica is added at the end of the buffer. Symmetrically, rule (RECEIVL) states that a replica consumes updates from the beginning of the buffer. Also in this case, we can prove that implementation correctness is preserved.

Lemma 8. *Let S be a coherent specification, C_S an implementation that is one-replica correct, and σ_0, σ_1 two states of C_S . If $\sigma_i \sim_S^0 \langle G_i, P_i \rangle$ for $i \in \{1, 2\}$, then $\sigma_1 \mid_{\varepsilon} \sigma_2 \sim_S^m \langle G_1 \sqcup G_2, P \rangle$ for any $P \in P_1 \otimes P_2$.*

Proposition 2. *Let S be a coherent specification and C_S be an implementation that is one-replica correct. Then, C_S^α is multi-replica correct.*

6. Related works

Weak consistency require to deal with conflict resolution [2,14,15] and convergence of replicas [8,16–18]. The design and implementation of data types that ensure convergence has been a very active area of research, notably, conflict-free and commutative replicated data type [2,14], which avoid implementing conflict resolution policies.

Different lines of work have addressed the problem of specifying and implementing replicated data types, considering also those that require policies for conflict resolution [3,17]. The approach in [3] has been addressed in detail in § 4. We remark that all of their specifications implicitly define a precise strategy to resolve conflicts. On the contrary, our specifications provide a more abstract view of RDTs, which do not commit to a particular strategy for conflict resolution. Implementation correctness is characterised in [3] in terms of *replication-aware simulations*. A replication-aware simulation is defined on top of abstract executions decorated with auxiliary information, such as time-stamps, which is defined *ad hoc* for each data type. Global correctness, i.e., the behaviour of several replicas, is derived from some agreement properties that impose a set of proof obligations that need to be checked. On the contrary, our characterisation of correctness relies on the abstract states induced by the specification. In our case, checking implementation correctness reduces to standard LTS simulation. Their approach can be instantiated to handle both state-based and operation-based implementations.

A different approach has been proposed in [17] to deal with optimistic replication systems. In this case, a specification associates an operation and a return value to the set of all possible executions (represented in terms of partial orders) that explain that particular return value for that operation. Moreover, they allow different operations to be associated with different partial orders. In this way, they deal with speculative systems. Furthermore, they reduce the problem of verifying eventual consistency to a model-checking and reachability problem.

A framework combining RDTs and transactions has been presented in [6]. As shown in [13], our specification style enables a categorical presentation of RDTs and the development of composition operators. Our results for parallel operators and implementation correctness suggests a way of dealing with RDT composition.

As far as the verification of (commutative) replicated data types is concerned, a framework in Isabelle/HOL has been proposed in [19], and other lines of work [8,9] have focused on the related problem of verifying properties of applications that use replicated data. In the long run, our goal is to exploit tools and techniques from the theory of simulation for precisely such purposes.

7. Conclusions and future works

We propose a denotational view of replicated data types. While most of the traditional approaches are operational in flavour [3,6,9], we strived for a specification formalism that could exploit the tools of algebraic specification theory. More precisely, we associate to each configuration (i.e., visibility) a set of admissible arbitrations. Differently from previous approaches, our presentation naturally accommodates non-deterministic specifications and enables abstract definitions allowing for different strategies in conflict resolution. Our formulation brings to light some properties held by mainstream specification formalisms: beside the obvious property of (local) functionality, they satisfy coherence and saturation. A coherent specification can neither prescribe an arbitration order between events that are unrelated by visibility nor allow for additional arbitrations over past events when a configuration is extended (i.e., a new top element is added to visibility). Instead, a saturated specification cannot impose any constraint to the arbitration of top events. Note that saturation does not hold when requiring that admissible arbitrations should be also topological orderings of visibility. Hence, the approaches in [3,4] generate specifications that are not saturated. We remark that this relation between visibility and arbitration translates in a quite different property in our setting, and this suggests that consistency models defined as relations between visibility

and arbitration (e.g., *monotonic* and *causal* consistency) could have alternative characterisations. We plan to explore these connections in future works.

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

We proposed an approach for checking that an implementation is correct with respect to a specification. Our characterisation relies on a notion of simulation and enjoys the property of being compositional with respect to standard communication models. This relieves us from checking that the behaviour of several replicas is correct. Our characterisation focuses only on state-based implementations. We are planning to extend the characterisation to operation-based implementations.

Acknowledgements

We thank the anonymous reviewers for their careful reading of our paper and their insightful comments. The first author has been partially supported by CONICET International Cooperation Grant 995/15. The second and third author have been partially supported by UBACyT project 2014–2017 20020130200092BA and CONICET project PIP 11220130100148CO.

Appendix A. Proof of the results in Section 4

Proof of Lemma 1. Let $P \in \bigotimes_{e \in \mathcal{E}_G} \text{PRVAL}(\mathcal{F}, G|_{-\prec^*e})$, that is, $\forall e \in \mathcal{E}_G. P|_{-\prec^*e} \in \text{PRVAL}(\mathcal{F}, G|_{-\prec^*e})$. By return value consistency (Definition 10), this is equivalent to

$$\forall e \in \mathcal{E}_G. \forall e' \in \mathcal{E}_{G|_{-\prec^*e}}. \lambda(e') = \langle \text{op}', v' \rangle \text{ implies } \mathcal{F}(\text{op}', (\bar{G}|_{-\prec^*e'})|_{-\prec e'}, (\bar{P}|_{-\prec^*e'})|_{-\prec e'}) = v'$$

Note that $(\bar{G}|_{-\prec^*e'})|_{-\prec e'}$ coincides with $\bar{G}|_{-\prec e'}$, and the same for \bar{P} . Then, the result follows because the formula above coincides with

$$\forall e \in \mathcal{E}_G. \lambda(e) = \langle \text{op}, v \rangle \text{ implies } \mathcal{F}(\text{op}, \bar{G}|_{-\prec e}, \bar{P}|_{-\prec e}) = v \quad \square$$

Proof of Lemma 2. Consider G_1, G_2 such that $\bar{G}_1 = \bar{G}_2$ and $G_1 \neq G_2$. We prove that $\overline{\mathcal{S}(G_1)} \cap \overline{\mathcal{S}(G_2)} = \emptyset$ follows. Since $G_1 \neq G_2$ there exists an event \bar{e} such that

$$G_1|_{-\prec+\bar{e}} = G_2|_{-\prec+\bar{e}} \text{ and } \lambda_1(\bar{e}) = \langle \text{op}, v_1 \rangle, \lambda_2(\bar{e}) = \langle \text{op}, v_2 \rangle \text{ for } v_1 \neq v_2$$

Let $G = G_i|_{-\prec+\bar{e}}$. By determinism, we have that $\overline{\mathcal{S}(G^{\lambda_1(\bar{e})})}|_{\mathcal{E}_G} \cap \overline{\mathcal{S}(G^{\lambda_2(\bar{e})})}|_{\mathcal{E}_G} = \emptyset$, and equivalently that $\overline{\mathcal{S}(G_1|_{-\prec^*\bar{e}})}|_{-\prec+\bar{e}} \cap \overline{\mathcal{S}(G_2|_{-\prec^*\bar{e}})}|_{-\prec+\bar{e}} = \emptyset$.

Now, assume that there exist $P_1 \in \mathcal{S}(G_1)$ such that $\bar{P}_1 = \bar{P}_2$, then by coherence $\forall e. P_1|_{-\prec^*e} \in \mathcal{S}(G_1|_{-\prec^*e})$. Consequently, $\forall e. P_1|_{-\prec^*e} \in \mathcal{S}(G_1|_{-\prec^*e})|_{-\prec+\bar{e}}$, which contradicts the statement above for $e = \bar{e}$. \square

Proof of Lemma 13. By locality, for any G and \mathcal{E} there exists G_1 such that for all op, v , we have $\bar{G}|_{\mathcal{E}} = \bar{G}_1 \wedge \overline{\mathcal{S}(G_{\mathcal{E}}^{(\text{op}, v)})}|_{-\prec_{\mathcal{T}}} \subseteq \overline{\mathcal{S}(G_1^{(\text{op}, v)})}$. Consequently, for any op and $v \neq v'$,

$$\begin{aligned} \overline{\mathcal{S}(G_{\mathcal{E}}^{(\text{op}, v)})}|_{-\prec_{\mathcal{T}}} &\subseteq \overline{\mathcal{S}(G_1^{(\text{op}, v)})} \\ \overline{\mathcal{S}(G_{\mathcal{E}}^{(\text{op}, v')})}|_{-\prec_{\mathcal{T}}} &\subseteq \overline{\mathcal{S}(G_1^{(\text{op}, v')})} \end{aligned}$$

If \mathcal{S} is not deterministic, then for some op and $v \neq v'$ we have $\overline{\mathcal{S}(G_{\mathcal{E}}^{(\text{op}, v)})}|_{\mathcal{E}_G} \cap \overline{\mathcal{S}(G_{\mathcal{E}}^{(\text{op}, v')})}|_{\mathcal{E}_G} \neq \emptyset$. Therefore $\overline{\mathcal{S}(G_{\mathcal{E}}^{(\text{op}, v)})}|_{-\prec_{\mathcal{T}}} \cap \overline{\mathcal{S}(G_{\mathcal{E}}^{(\text{op}, v')})}|_{-\prec_{\mathcal{T}}} \neq \emptyset$ and thus we have $\overline{\mathcal{S}(G_1^{(\text{op}, v)})} \cap \overline{\mathcal{S}(G_1^{(\text{op}, v')})} \neq \emptyset$, against weakly determinism for G_1 . \square

Proof of Lemma 4. We prove the four properties of $\mathbb{S}(\mathcal{F})$.

- Coherent. Immediate by Lemma 1.
- Saturated. We have to prove that

$$\forall (G, P), \mathcal{E}, \langle \text{op}, v \rangle. P \in \mathbb{S}(\mathcal{F})(G_{\mathcal{E}}^{(\text{op}, v)}) \Big|_{\mathcal{E}_G} \text{ implies } \text{sat}(P, \langle \text{op}, v \rangle) \subseteq \mathbb{S}(\mathcal{F})(G_{\mathcal{E}}^{(\text{op}, v)})$$

Since $P \in \mathbb{S}(\mathcal{F})(G_{\mathcal{E}}^{(\text{op}, v)}) \Big|_{\mathcal{E}_G}$, then there is $P_1 \in \text{sat}(P, \langle \text{op}, v \rangle) \cap \mathbb{S}(\mathcal{F})(G_{\mathcal{E}}^{(\text{op}, v)})$.

$$\begin{aligned} P_1 \in \mathbb{S}(\mathcal{F})(G_{\mathcal{E}}^{(\text{op}, v)}) &\stackrel{\text{Definition 14}}{=} P_1 \in \text{PRVAL}(\mathcal{F}, G_{\mathcal{E}}^{(\text{op}, v)}) \\ &\stackrel{\text{Definition 10}}{=} \text{RVAL}(\mathcal{F}, G_{\mathcal{E}}^{(\text{op}, v)}, P_1) \\ &\stackrel{\text{Definition 10}}{=} \text{RVAL}(\mathcal{F}, G_{\mathcal{E}}^{(\text{op}, v)}) \Big|_{\mathcal{E}_G}, P_1|_{\mathcal{E}_G} \wedge \mathcal{F}(\text{op}, \overline{G_{\mathcal{E}}^{(\text{op}, v)}}) \Big|_{\neg\prec\top}, \overline{P_1}|_{\neg\prec\top} = v \\ &= \text{RVAL}(\mathcal{F}, G, P) \wedge \mathcal{F}(\text{op}, \overline{G}|_{\mathcal{E}}, \overline{P_1}|_{\mathcal{E}}) = v \end{aligned}$$

Now, for every $P_2 \in \text{sat}(P, \langle \text{op}, v \rangle)$ note that $\overline{P_2}|_{\mathcal{E}} = \overline{P_1}|_{\mathcal{E}}$ and, consequently, $\mathcal{F}(\text{op}, \overline{G}|_{\mathcal{E}}, \overline{P_1}|_{\mathcal{E}}) = v$ holds. Hence, $P_2 \in \mathbb{S}(\mathcal{F})(G_{\mathcal{E}}^{(\text{op}, v)})$ holds and the result follows.

- Total. Let us assume that there exist G, P, \mathcal{E} , and op such that

$$\forall G_1, v. \overline{G} = \overline{G_1} \text{ implies } \overline{P} \notin \overline{\mathbb{S}(\mathcal{F})(G_1)_{\mathcal{E}}^{(\text{op}, v)}} \Big|_{\mathcal{E}_{G_1}} \quad (1)$$

Moreover, without loss of generality we can assume that G is a minimal LDAG that satisfies (1), i.e., that for all G' strictly contained in G and for all $P' \in \mathbb{S}(\mathcal{F})(G')$ we have

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

Now, let us consider G', \mathcal{E}' , and op' such that $\overline{G} = (\overline{G'})_{\mathcal{E}'}^{\text{op}'}$ and let $P' = P|_{\mathcal{E}'}$. Since we proved that $\mathbb{S}(\mathcal{F})$ is saturated we have

$$\exists (G'_1, P'_1), v'. \overline{G'} = \overline{G'_1} \wedge \overline{P'} = \overline{P'_1} \wedge \text{sat}(P'_1, \langle \text{op}', v' \rangle) \subseteq \mathbb{S}(\mathcal{F})(G'_1)_{\mathcal{E}'}^{(\text{op}', v')}$$

Then, let us take $G_1 = (G'_1)_{\mathcal{E}'}^{(\text{op}', v')}$ and $P_1 \in \text{sat}(P'_1, \langle \text{op}', v' \rangle)$ such that $\overline{P} = \overline{P_1}$ and $P_1 \in \mathbb{S}(\mathcal{F})(G_1)$. First, note that $\overline{G} = \overline{G_1}$ because $\overline{G} = (\overline{G'})_{\mathcal{E}'}^{\text{op}'}$, $\overline{G'} = \overline{G'_1}$, and $G_1 = (G'_1)_{\mathcal{E}'}^{(\text{op}', v')}$. We now show that for each \mathcal{E} and op there exists v such that $P_1^{(\text{op}, v)} \in \mathbb{S}(\mathcal{F})(G_1)_{\mathcal{E}}^{(\text{op}, v)}$, which is in contradiction with the assumption (1). Given \mathcal{E} and op , take v such that $\mathcal{F}(\text{op}, \overline{G_1}|_{\mathcal{E}}, \overline{P_1}|_{\mathcal{E}}) = v$ (such v exists because of the definition of \mathcal{F}). By following the same reasoning as for saturation we have

$$\begin{aligned} P_1^{(\text{op}, v)} \in \mathbb{S}(\mathcal{F})(G_1)_{\mathcal{E}}^{(\text{op}, v)} &\stackrel{\text{Definition 14}}{=} P_1^{(\text{op}, v)} \in \text{PRVAL}(\mathcal{F}, (G_1)_{\mathcal{E}}^{(\text{op}, v)}) \\ &\stackrel{\text{Definition 10}}{=} \dots \\ &= \text{RVAL}(\mathcal{F}, G_1, P_1) \wedge \mathcal{F}(\text{op}, \overline{G_1}|_{\mathcal{E}}, \overline{P_1}|_{\mathcal{E}}) = v \end{aligned}$$

Since $\text{RVAL}(\mathcal{F}, G_1, P_1)$ coincides with $P_1 \in \mathbb{S}(\mathcal{F})(G_1)$, the result follows.

- Deterministic. Let $G, \mathcal{E}, \text{op}, v_1$, and v_2 such that $v_1 \neq v_2$ and

$$\overline{\mathbb{S}(\mathcal{F})(G_{\mathcal{E}}^{(\text{op}, v_1)})} \Big|_{\mathcal{E}_G} \cap \overline{\mathbb{S}(\mathcal{F})(G_{\mathcal{E}}^{(\text{op}, v_2)})} \Big|_{\mathcal{E}_G} \neq \emptyset$$

Then, there exist paths P_1, P_2 such that $P_i \in \mathbb{S}(\mathcal{F})(G_{\mathcal{E}}^{(\text{op}, v_i)})$ and $\overline{P_1}|_{\mathcal{E}_G} = \overline{P_2}|_{\mathcal{E}_G}$. By the definition of $\mathbb{S}(\mathcal{F})$ (see Definition 14), the former is equivalent to $P_i \in \text{PRVAL}(\mathcal{F}, G_{\mathcal{E}}^{(\text{op}, v_i)})$ and then to $\text{RVAL}(\mathcal{F}, G_{\mathcal{E}}^{(\text{op}, v_i)}, P_i)$. By return value consistency (see Definition 10), we have that $\mathcal{F}(\text{op}, \overline{G_{\mathcal{E}}^{(\text{op}, v_i)}}) \Big|_{\neg\prec\top}, \overline{P_i}|_{\neg\prec\top} = v_i$, and since $\overline{G_{\mathcal{E}}^{(\text{op}, v_i)}} \Big|_{\neg\prec\top} = G|_{\mathcal{E}}$ and $\overline{P_i}|_{\neg\prec\top} = \overline{P_i}|_{\mathcal{E}}$, it results in a contradiction.

- Local. For every P such that $\overline{P} \in \overline{\mathbb{S}(\mathcal{F})(G_{\mathcal{E}}^{(\text{op}, v)})} \Big|_{\neg\prec\top}$ there exists $P_1 \in \mathbb{S}(\mathcal{F})(G_{\mathcal{E}}^{(\text{op}, v)})$ such that $\overline{P} = \overline{P_1}|_{\neg\prec\top}$. Then, by following the same reasoning as for saturation we have

$$\begin{aligned} P_1 \in \mathbb{S}(\mathcal{F})(G_{\mathcal{E}}^{(\text{op}, v)}) &\stackrel{\text{Definition 14}}{=} P_1 \in \text{PRVAL}(\mathcal{F}, G_{\mathcal{E}}^{(\text{op}, v)}) \\ &\stackrel{\text{Definition 10}}{=} \forall e \in \mathcal{E}_G. \lambda(e) = \langle \text{op}', v' \rangle \text{ implies } \mathcal{F}(\text{op}', \overline{G}|_{\neg\prec e}, \overline{P_1}|_{\neg\prec e}) \\ &= v' \wedge \mathcal{F}(\text{op}, \overline{G}|_{\mathcal{E}}, \overline{P_1}|_{\mathcal{E}}) = v \end{aligned}$$

Since \mathcal{F} is a function, for every G_1 such that $\bar{G}|_{\mathcal{E}} = \bar{G}_1$ the following holds

$$\forall e \in \mathcal{E}_{G_1}, \text{op}' . \exists v'' . \mathcal{F}(\text{op}', \bar{G}_1|_{-\prec e}, \bar{P}_1|_{-\prec e}) = v''$$

Consequently, by coherence there exists $P_2 \in \mathbb{S}(\mathcal{F})(G_1)$ such that $\bar{P}_1|_{\mathcal{E}_{G_1}} = \bar{P}_2$. Moreover, $\mathcal{F}(\text{op}, \bar{G}|_{\mathcal{E}}, \bar{P}_1|_{\mathcal{E}}) = v$ and by saturation this implies that $\text{sat}(P_2, (\text{op}, v)) \subseteq \mathbb{S}(\mathcal{F})(G_1^{(\text{op}, v)})$. Therefore, $\bar{P} \in \overline{\mathcal{S}(G_1^{(\text{op}, v)})}$. \square

Proof of Lemma 5. Since \mathcal{S} is total, there exists at least a value for every triple op, \bar{G} , and \bar{P} . Let us now assume that there are more than two values, i.e., that there exist op, G_i , and v_i such that

$$\bar{G} = \bar{G}_1 = \bar{G}_2 \wedge v_1 \neq v_2 \wedge \bar{P} \in \overline{\mathcal{S}(G_i^{(\text{op}, v_i)})}|_{\mathcal{E}_{G_i}}$$

Since \mathcal{S} is deterministic and $\bar{G}_1 = \bar{G}_2$, two cases may occur by Lemma 2

- $G_1 = G_2$. By hypothesis we have that $\bar{P} \in \overline{\mathcal{S}(G_1^{(\text{op}, v_i)})}|_{\mathcal{E}_{G_1}}$, hence

$$\overline{\mathcal{S}(G_1^{(\text{op}, v_1)})}|_{\mathcal{E}_{G_1}} \cap \overline{\mathcal{S}(G_1^{(\text{op}, v_2)})}|_{\mathcal{E}_{G_1}} \neq \emptyset$$

Again by being \mathcal{S} deterministic, it follows that $v_1 = v_2$.

- $\overline{\mathcal{S}(G_1)} \cap \overline{\mathcal{S}(G_2)} = \emptyset$. Since \mathcal{S} is coherent and saturated and by hypothesis $\bar{P} \in \overline{\mathcal{S}(G_i^{(\text{op}, v_i)})}|_{\mathcal{E}_{G_i}}$, we have that $\bar{P} \in \overline{\mathcal{S}(G_i)}$, which leads to a contradiction. \square

We state an auxiliary result that will be used in the proof of Theorem 1.

Lemma 9. Let \mathcal{S} be a coherent, functional, and saturated specification and $\langle G, P \rangle$ a context. If

$$\forall e \in \mathcal{E}_G. \lambda(e) = (\text{op}, v) \text{ implies } \exists G_1. \bar{G}|_{-\prec e} = \bar{G}_1 \wedge \bar{P}|_{-\prec e} \in \overline{\mathcal{S}(G_1^{(\text{op}, v)})}|_{\mathcal{E}_{G_1}} \quad (2)$$

then $\forall e \in \mathcal{E}_G. P|_{-\prec^* e} \in \mathcal{S}(G|_{-\prec^* e})$.

Moreover, let \mathcal{S} be a local and coherent specification. Then, the vice versa holds.

Proof. (\Leftarrow) Let \bar{e} such that $\lambda(\bar{e}) = (\text{op}, v)$. By hypothesis $P|_{-\prec^* \bar{e}} \in \mathcal{S}(G|_{-\prec^* \bar{e}})$, by coherence $P|_{-\prec^+ \bar{e}} \in \mathcal{S}(G|_{-\prec^+ \bar{e}})$, and by locality

$$\exists G_1. (\bar{G}|_{-\prec^+ \bar{e}})|_{-\prec \bar{e}} = \bar{G}_1 \wedge \overline{\mathcal{S}((G|_{-\prec^+ \bar{e}})_{-\prec \bar{e}}^{(\text{op}, v)})}|_{-\prec \bar{e}} \subseteq \overline{\mathcal{S}(G_1^{(\text{op}, v)})}$$

This is in turn equivalent to

$$\exists G_1. \bar{G}|_{-\prec \bar{e}} = \bar{G}_1 \wedge \overline{\mathcal{S}(G|_{-\prec^* \bar{e}})}|_{-\prec \bar{e}} \subseteq \overline{\mathcal{S}(G_1^{(\text{op}, v)})}$$

and again by the hypothesis $P|_{-\prec^* \bar{e}} \in \mathcal{S}(G|_{-\prec^* \bar{e}})$ the result follows.

(\Rightarrow) By contradiction, let us assume that there exist $\bar{e} \in \mathcal{E}_G$ such that the equation (2) holds but $P|_{-\prec^* \bar{e}} \notin \mathcal{S}(G|_{-\prec^* \bar{e}})$. Without loss of generality, assume e is minimal, i.e., for all e' such that $e' \prec^+ \bar{e}$ we have

$$P|_{-\prec^* e'} \in \mathcal{S}(G|_{-\prec^* e'})$$

By coherence we have

$$P|_{-\prec^+ \bar{e}} \in \mathcal{S}(G|_{-\prec^+ \bar{e}})$$

Assuming $\lambda(\bar{e}) = (\text{op}, v)$, by totality it follows

$$\exists G_1, v_1. \bar{G}_1 = \bar{G}|_{-\prec^+ \bar{e}} \wedge \bar{P}|_{-\prec^+ \bar{e}} \in \overline{\mathcal{S}((G_1)_{-\prec \bar{e}}^{(\text{op}, v_1)})}|_{-\prec^+ \bar{e}}$$

Now, by saturation and coherence we have that $\bar{P}|_{-\prec^+ \bar{e}} \in \overline{\mathcal{S}(G_1)}$, thus by determinism $G_1 = G|_{-\prec^+ \bar{e}}$, so that it holds

$$\exists v_1. \bar{P}|_{-\prec^+ \bar{e}} \in \overline{\mathcal{S}((G|_{-\prec^+ \bar{e}})_{-\prec \bar{e}}^{(\text{op}, v_1)})}|_{-\prec^+ \bar{e}}$$

Now, by saturation and (2)

$$\exists G_2. (\bar{G}|_{-\prec+\bar{e}})|_{-\prec\bar{e}} = \bar{G}_2 \wedge \bar{P}|_{-\prec\bar{e}} \in \overline{\mathcal{S}((G_2)^{(op,v_1)})}|_{-\prec\bar{e}}$$

However, applying (2) to G tells that

$$\exists G_3. \bar{G}|_{-\prec\bar{e}} = \bar{G}_3 \wedge \bar{P}|_{-\prec\bar{e}} \in \overline{\mathcal{S}((G_3)^{(op,v)})}|_{-\prec\bar{e}}$$

and by saturation and determinism $(G_2)^{(op,v_1)} = (G_3)^{(op,v)}$, so that $v = v_1$.

Thus now we have

$$\bar{P}|_{-\prec+\bar{e}} \in \overline{\mathcal{S}(G|_{-\prec*\bar{e}})}|_{-\prec+\bar{e}}$$

and by saturation it holds that

$$\bar{P}|_{-\prec*\bar{e}} \in \overline{\mathcal{S}(G|_{-\prec*\bar{e}})}$$

and since by hypothesis

$$P|_{-\prec+\bar{e}} \in \mathcal{S}(G|_{-\prec+\bar{e}})$$

we have that

$$P|_{-\prec*\bar{e}} \in \mathcal{S}(G|_{-\prec*\bar{e}})$$

which is in contradiction with the assumptions. \square

Proof of Theorem 1. We have to prove that for any G it holds $\mathcal{S}(G) = \mathbb{S}(\mathbb{F}(\mathcal{S}))(G)$.

$$\begin{aligned} P \in \mathbb{S}(\mathbb{F}(\mathcal{S}))(G) &\stackrel{\text{Definition 14}}{=} P \in \text{PRVAL}(\mathbb{F}(\mathcal{S}), G) \\ &\stackrel{\text{Definition 10}}{=} \text{RVAL}(\mathbb{F}(\mathcal{S}), G, P) \\ &\stackrel{\text{Definition 10}}{=} \forall e \in \mathcal{E}_G. \lambda(e) = \langle op, v \rangle \text{ implies } \mathbb{F}(\mathcal{S})(op, \bar{G}|_{-\prec e}, \bar{P}|_{-\prec e}) = v \\ &\stackrel{\text{Definition 15}}{=} \forall e \in \mathcal{E}_G. \lambda(e) = \langle op, v \rangle \text{ implies } \exists G_1. \bar{G}|_{-\prec e} = \bar{G}_1 \wedge \bar{P}|_{-\prec e} \in \overline{\mathcal{S}(G_1^{(op,v)})}|_{\mathcal{E}_{G_1}} \\ &\stackrel{\text{Lemma 9}}{=} \forall e \in \mathcal{E}_G. P|_{-\prec*e} \in \mathcal{S}(G|_{-\prec*e}) \\ &\stackrel{\text{Definition 3}}{=} P \in \bigotimes_{e \in \mathcal{E}_G} \mathcal{S}(G|_{-\prec*e}) \\ &\stackrel{\text{coh}}{=} P \in \mathcal{S}(G) \quad \square \end{aligned}$$

Proof of Theorem 2. We prove that $\mathcal{F}(op, \bar{G}, \bar{P}) = \mathbb{F}(\mathbb{S}(\mathcal{F}))(op, \bar{G}, \bar{P})$ for any $\langle \bar{G}, \bar{P} \rangle$ and op .

$$\begin{aligned} \mathbb{F}(\mathbb{S}(\mathcal{F}))(op, \bar{G}, \bar{P}) &= v \\ &\stackrel{\text{Definition 15}}{=} \exists (G_1, P_1). \bar{G} = \bar{G}_1 \wedge \bar{P} = \bar{P}_1 \wedge P_1 \in \mathbb{S}(\mathcal{F})(G_1^{(op,v)})|_{\mathcal{E}_{G_1}} \\ &\stackrel{\text{sat}}{\Rightarrow} \exists (G_1, P_1). \bar{G}_1 = \bar{G} \wedge \bar{P}_1 = \bar{P} \wedge P_1^{(op,v)} \in \mathbb{S}(\mathcal{F})(G_1^{(op,v)}) \\ &\stackrel{\text{Definition 15}}{=} \dots \wedge P_1^{(op,v)} \in \text{PRVAL}(\mathcal{F}, G_1^{(op,v)}) \\ &\stackrel{\text{Definition 10}}{=} \dots \wedge \text{RVAL}(\mathcal{F}, G_1^{(op,v)}, P_1^{(op,v)}) \\ &\stackrel{\text{Definition 10}}{=} \dots \wedge \text{RVAL}(\mathcal{F}, G_1, P_1) \wedge \mathcal{F}(op, \overline{(G_1)^{(op,v)}}|_{-\prec_T}, \overline{(P_1)^{(op,v)}}|_{-\prec_T}) = v \\ &\stackrel{\text{coh}}{=} \dots \wedge \mathcal{F}(op, \bar{G}_1, \bar{P}_1) = v \\ &\Rightarrow \mathcal{F}(op, \bar{G}, \bar{P}) = v \end{aligned}$$

Then the result holds, since \mathcal{F} and $\mathbb{F}(\mathbb{S}(\mathcal{F}))$ are total functions. \square

Appendix B. Proof of the results in Section 5

Proof of Lemma 6. It follows from the fact that \mathcal{S} is coherent. \square

Proof of Example 16. By case analysis on the derivation of $\sigma \xrightarrow{\alpha} \sigma'$. Consider $\sigma = \langle r, v \rangle$ and $(\sigma, \langle G, P \rangle) \in \mathcal{I}$. Consequently, there exists $f : \mathcal{E}_G \rightarrow \mathcal{R}$ such that

$$\forall r \in \mathcal{R}. v(r) = \#\{e \mid f(e) = r \text{ and } \lambda(e) = \langle inc, ok \rangle\} \quad (3)$$

- (READ) Hence, $\sigma = \langle r_j, v \rangle \xrightarrow{rd, k} \langle r_j, v \rangle = \sigma'$ and $k = \sum_{r \in \mathcal{R}} v(r)$. From (3), $k = \sum_{r \in \mathcal{R}} v(r) = \#\{e \mid e \in \mathcal{E}_G \wedge \lambda(e) = \langle inc, ok \rangle\}$. Define $G' = G^{(rd, k)}$ and $f' : \mathcal{E}_{G'} \rightarrow \mathcal{R}$ as follows

$$f'(e) = \begin{cases} f(e) & \text{if } e \in \mathcal{E}_G \\ r_j & \text{if } e = \top \end{cases}$$

It is immediate to check that

$$\forall r \in \mathcal{R}. v(r) = \#\{e \mid f'(e) = r \text{ and } \lambda(e) = \langle inc, ok \rangle\}$$

Hence,

$$k = \sum_{r \in \mathcal{R}} v(r) = \#\{e \mid e \in \mathcal{E}_{G'} \wedge \lambda(e) = \langle inc, ok \rangle\}$$

Consequently, $\text{sat}(P, (rd, k)) \subseteq \mathcal{S}_{Ctr}(G')$. Then, for any $P' \in \text{sat}(P, (rd, k))$, we have $(\langle r_j, v \rangle, \langle G^{(rd, k)}, P' \rangle) \in \mathcal{I}$.

- (INC) Then, $\sigma = \langle r_j, v \rangle \xrightarrow{inc, ok} \langle r_j, v[r_j \mapsto v(r_j) + 1] \rangle = \sigma'$. Take $G' = G^{(inc, ok)}$. It is immediate to check that $\text{sat}(P, \langle inc, ok \rangle) \subseteq \mathcal{S}_{Ctr}(G')$. Then, define $f' : \mathcal{E}_{G'} \rightarrow \mathcal{R}$ as follows

$$f'(e) = \begin{cases} f(e) & \text{if } e \in \mathcal{E}_G \\ r_j & \text{if } e = \top \end{cases}$$

Hence, $\#\{e \mid f'(e) = r \text{ and } \lambda(e) = \langle inc, ok \rangle\} = v(r)$ for all $r \neq r_j$, and

$$\#\{e \mid f'(e) = r_j \text{ and } \lambda(e) = \langle inc, ok \rangle\} = v(r_j) + 1$$

Consequently, for any $P' \in \text{sat}(P, \langle inc, ok \rangle)$ the following holds

$$(\langle r_j, v[r_j \mapsto v(r_j) + 1] \rangle, \langle G^{(inc, ok)}, P' \rangle) \in \mathcal{I}$$

- (RECEIVE) Then, $\sigma = \langle r, v \rangle \xrightarrow{rcv, (r', v')} \langle r, v' \rangle = \sigma'$ such that $v' = \max\{v, v'\}$. Define $G' = \langle \mathcal{F}, \prec', \lambda' \rangle$ and $f' : \mathcal{E}_{G'} \rightarrow \mathcal{R}$ such that $\forall r \in \mathcal{R}$
 - $v'(r) = \#\{e \mid f'(e) = r \text{ and } \lambda'(e) = \langle inc, ok \rangle\}$
 - $\min\{v(r), v'(r)\} = \#\{e \mid f(e) = r\} \cap \#\{e \mid f'(e) = r\}$
 - $\prec_{G'} \upharpoonright_{\mathcal{E}_G \cap \mathcal{E}_{G'}} = \prec_G \upharpoonright_{\mathcal{E}_G \cap \mathcal{E}_{G'}}$,
 - $\lambda_{G'} \upharpoonright_{\mathcal{E}_G \cap \mathcal{E}_{G'}} = \lambda_G \upharpoonright_{\mathcal{E}_G \cap \mathcal{E}_{G'}}$.

It is immediate to note that there exists P' such that $(\langle r', v' \rangle, \langle G', P' \rangle) \in \mathcal{I}$. Moreover, $G \sqcup G'$ is defined under the above conditions and for any $P' \in \mathcal{S}_{Ctr}(G')$ we have $P \otimes P' \subseteq \mathcal{S}_{Ctr}(G \sqcup G')$. Hence, it suffices to take $P'' \in P \otimes P'$.

Define $f'' : \mathcal{E}_{G \sqcup G'} \rightarrow \mathcal{R}$ as

$$f''(e) = \begin{cases} f(e) & \text{if } e \in \mathcal{E}_G \\ f'(e) = r_j & \text{otherwise} \end{cases}$$

From the definition of G' and f'' it follows that for all $r \in \mathcal{R}$

$$\#\{e \mid e \in \mathcal{E}_{G \sqcup G'} \text{ and } \lambda(e) = \langle inc, ok \rangle\} = \max\{v(r), v'(r)\} = v'(r)$$

Consequently, $(\langle r, v' \rangle, \langle G \sqcup G', P'' \rangle) \in \mathcal{I}$.

- (SEND) Then, $\sigma = \langle r, v \rangle \xrightarrow{send, (r, v)} \langle r, v \rangle = \sigma'$. It follows immediately, because $(\langle r, v \rangle, \langle G, P \rangle) \in \mathcal{I}$ holds by hypothesis. \square

Proof of Example 17. By case analysis on the applied rule for the derivation of $\sigma \xrightarrow{\alpha} \sigma'$. Consider $\sigma = \langle r_j, V, W \rangle$ and $(\sigma, \langle G, P \rangle) \in \mathcal{I}$. Hence, there exists $f : \mathcal{E}_G \rightarrow \mathcal{R}$ and

$$\forall k. (\exists r. W(k, r) > 0 \iff \exists S. \mathcal{S}_{Or-Set}(G^{(lookup, S)}) \neq \emptyset \wedge k \in S) \quad (4)$$

- (LOOKUP) Hence, $\sigma \xrightarrow{lookup, S} \sigma' = \sigma$. Define $G' = G^{(lookup, S)}$ and take $f' : \mathcal{E}_{G'} \rightarrow \mathcal{R}$ defined as follows

$$f'(e) = \begin{cases} f(e) & \text{if } e \in \mathcal{E}_G \\ r_j & \text{if } e = \top \end{cases}$$

Note that

$$\forall k. (\exists r. W(k, r) > 0 \iff \exists S. \mathcal{S}_{Or-Set}(G'^{(lookup, S)}) \neq \emptyset \wedge k \in S)$$

follows from the definition of \mathcal{S}_{Or-Set} and (4). Moreover, we conclude that $\text{sat}(P, \langle lookup, S \rangle) \subseteq \mathcal{S}_{Or-Set}(G'^{(lookup, S)})$. Therefore, $(\sigma', \langle G', P' \rangle) \in \mathcal{I}$ for any $P' \in \text{sat}(P, \langle lookup, S \rangle)$.

- (ADD) Then, $\sigma \xrightarrow{\text{add}(k'), \text{ok}} \langle r_j, \forall [r_j \mapsto V(r_j) + 1], W[(k', r_j) \mapsto V(r_j) + 1] \rangle = \sigma'$. Define $G' = G^{(\text{add}(k'), \text{ok})}$ and take $f' : \mathcal{E}_{G'} \rightarrow \mathcal{R}$ defined as follows

$$f'(e) = \begin{cases} f(e) & \text{if } e \in \mathcal{E}_G \\ r_j & \text{if } e = \top \end{cases}$$

Now, we show that the following holds

$$\forall k (\exists r. W[(k', r_j) \mapsto V(r_j) + 1](k, r) > 0 \iff \exists S. \mathcal{S}_{Or\text{-}Set}(G'^{(\text{lookup}, S)}) \neq \emptyset \wedge k \in S) \quad (5)$$

For all $k \neq k'$, the condition holds by (4)

$$W[(k', r_j) \mapsto V(r_j) + 1](k, r) > 0 \iff \exists S. \mathcal{S}_{Or\text{-}Set}(G^{(\text{lookup}, S')}) \neq \emptyset \wedge k \in S$$

Then, for $k = k'$ it is enough to take $r = r_j$ to conclude that the double implication holds.

By the definition of $\mathcal{S}_{Or\text{-}Set}$, we have that $\mathcal{S}_{Or\text{-}Set}(G^{(\text{lookup}, S')}) \neq \emptyset$ implies $\mathcal{S}_{Or\text{-}Set}(G'^{(\text{lookup}, \text{SU}\{k'\})}) \neq \emptyset$. Additionally, $\text{sat}(\mathcal{P}, \langle \text{add}(k), \text{ok} \rangle) \subseteq \mathcal{S}_{Or\text{-}Set}(G')$ because $\mathcal{P} \in \mathcal{S}_{Or\text{-}Set}(G)$. Then, for any $P' \in \text{sat}(\mathcal{P}, \langle \text{add}(k), \text{ok} \rangle)$, we have $(\sigma', (G', P')) \in \mathcal{I}$.

- (REM) Then, $\sigma \xrightarrow{\text{rem}(k'), \text{ok}} \langle r_j, V, W' \rangle = \sigma'$, and $\forall s. W'(k', s) = 0$, and $\forall s, k'' \neq k'. W'(k'', s) = W(k'', s)$. Define $G' = G^{(\text{rem}(k'), \text{ok})}$ and take $f' : \mathcal{E}_{G'} \rightarrow \mathcal{R}$ defined as follows

$$f'(e) = \begin{cases} f(e) & \text{if } e \in \mathcal{E}_G \\ r_j & \text{if } e = \top \end{cases}$$

As in the previous case, we conclude that the following holds for all $k \neq k'$

$$\exists r. W'(k, r) > 0 \iff \exists S. \mathcal{S}_{Or\text{-}Set}(G'^{(\text{lookup}, S)}) \neq \emptyset \wedge k \in S$$

from (4) and the fact that $W'(k'', r) = W(k'', r)$ holds for all $k'' \neq k'$ and r . By $\mathcal{S}_{Or\text{-}Set}$, $\mathcal{S}_{Or\text{-}Set}(G'^{(\text{lookup}, S)}) \neq \emptyset$ implies $\{k'\} \notin S$. Then, the proof is completed by noting that $\forall r. W'(k', r) = 0$.

- (SEND) Then, $\sigma \xrightarrow{\text{send}, \sigma} \sigma$. It follows because $(\sigma, (G, \mathcal{P})) \in \mathcal{I}$ by hypothesis.
- (RECEIVE) Then, $\sigma \xrightarrow{\text{receive}(r_k, V', W')} \langle r_j, \max\{V, V'\}, (V, W) \oplus (V', W') \rangle = \sigma'$. Define $G' = \langle \mathcal{E}_{G'}, \prec_{G'}, \lambda_{G'} \rangle$ where
 - $\mathcal{E}_{G'} = \bigsqcup_{r \in \mathcal{R}, k \in \mathbb{N}} \mathcal{F}_r^k$ (where \sqcup stands for disjoint union) and

$$\mathcal{F}_r^k = \begin{cases} \mathcal{E}_r^k & \text{if } W'(k, r) \neq 0 \wedge W'(k, r) \leq V(r) & \text{(i)} \\ \mathcal{E}_r^k \sqcup \{a_r^k\} & \text{if } W'(k, r) \neq 0 \wedge W'(k, r) > V(r) & \text{(ii)} \\ \mathcal{E}_r^k \sqcup \{r_r^k\} & \text{if } W'(k, r) = 0 \wedge W(k, r) \leq V'(r) & \text{(iii)} \\ \emptyset & \text{if } W'(k, r) = 0 \wedge W(k, r) > V'(r) & \text{(iv)} \end{cases}$$

where $\mathcal{E}_r^k = \{e \mid e \in \mathcal{E}_G \wedge f(e) = r \wedge \lambda_G(e) = \langle \text{add}(k), \text{ok} \rangle\}$

- $\lambda_{G'}$ is defined such that

$$\lambda_{G'}(e) = \begin{cases} \langle \text{add}(k), \text{ok} \rangle & \text{if } e \in \mathcal{F}_r^k \wedge e \neq r_r^k \\ \langle \text{rem}(k), \text{ok} \rangle & \text{if } e = r_r^k \end{cases}$$

- $\prec_{G'}$ defined such that

$$\prec_{G'} \upharpoonright \mathcal{F}_r^k = \begin{cases} \prec \upharpoonright \mathcal{E}_r^k & \text{if } \mathcal{F}_r^k = \mathcal{E}_r^k \\ \prec \upharpoonright \mathcal{E}_r^k & \text{if } \mathcal{F}_r^k = \mathcal{E}_r^k \sqcup \{a_r^k\} \\ \prec \upharpoonright \mathcal{E}_r^k \cup (\mathcal{E}_r^k \times \{r_r^k\}) & \text{if } \mathcal{F}_r^k = \mathcal{E}_r^k \sqcup \{r_r^k\} \\ \emptyset & \text{if } \mathcal{F}_r^k = \emptyset \end{cases}$$

Now we check that $((r_k, V', W'), (G', P')) \in \mathcal{I}$. By definition of $\mathcal{S}_{Or\text{-}Set}$ in Example 8, we may conclude that $\mathcal{S}_{Or\text{-}Set}(G'^{(\text{lookup}, S')}) \neq \emptyset$ whenever

$$S' = \{k \mid \exists e' \in \mathcal{E}_{G'}, \exists k \in \mathbb{N}. \lambda_{G'}(e') = \langle \text{add}(k), \text{ok} \rangle \text{ and } \forall e''. e' \prec_{G'} e'' \text{ implies } \lambda_{G'}(e'') \neq \langle \text{rem}(k), \text{ok} \rangle\}$$

From the definition of G' , it is immediate that $k \in S'$ iff $\exists r. W'(k, r) > 0$ (the cases (i) and (ii) on the definition of \mathcal{F}_r^k apply). Then, for any $P' \in \mathcal{S}_{Or\text{-}Set}(G')$, we have $((r_k, V', W'), (G', P')) \in \mathcal{I}$ by taking $f'' : \mathcal{E}_{G'} \rightarrow \mathcal{R}$ defined such that $f''(e) = r$ if there exists $k \in \mathbb{N}$ s.t. $e \in \mathcal{F}_r^k$.

Note that $G \sqcup G'$ is well-defined because G' is defined such that $\lambda_{G'} \upharpoonright \mathcal{E}_G \cap \mathcal{E}_{G'} = \lambda_G \upharpoonright \mathcal{E}_G \cap \mathcal{E}_{G'}$ and $\prec_{G'} \upharpoonright \mathcal{E}_G \cap \mathcal{E}_{G'} = \prec_G \upharpoonright \mathcal{E}_G \cap \mathcal{E}_{G'}$. Moreover, $\mathcal{P} \otimes P' \neq \emptyset$ by the definition of G' . It remains to prove that $(\sigma', (G \sqcup G', P'')) \in \mathcal{I}$ for any $P'' \in \mathcal{P} \otimes P'$.

Now define $f' : \mathcal{E}_{G \sqcup G'} \rightarrow \mathcal{R}$ such that

$$f'(e) = \begin{cases} f(e) & \text{if } e \in \mathcal{E}_G \\ f''(e) & \text{if } e \in \mathcal{E}_{G'} \end{cases}$$

Note that f' is well-defined because $\{\mathcal{F}_r^k\}_{r \in \mathcal{R}, k \in \mathbb{N}}$ is a partition of $\mathcal{E}_{G'}$, and the definition of G' ensures that $\mathcal{E}_r^k \cap \mathcal{F}_s^{k'} = \emptyset$ for all k, k' and $r \neq s$. Finally, we show that for all k , the following holds

$$\exists r. (\mathcal{V}, \mathcal{W}) \oplus (\mathcal{V}', \mathcal{W}')(k, r) > 0 \iff \exists S. \mathcal{S}((G \sqcup G')^{(\text{lookup}, S)}) \neq \emptyset \wedge k \in S$$

For \Rightarrow), assume that $(\mathcal{V}, \mathcal{W}) \oplus (\mathcal{V}', \mathcal{W}')(k, r) > 0$ for some r . By the definition of \oplus , $\neg(\text{IsREM}(\mathcal{W}, \mathcal{W}', \mathcal{V}, k, r) \vee \text{IsREM}(\mathcal{W}', \mathcal{W}, \mathcal{V}', k, r))$. Consequently,

$$\neg ((\mathcal{W}(k, r) = 0 \wedge \mathcal{W}'(k, r) \leq \mathcal{V}(r)) \vee (\mathcal{W}'(k, r) = 0 \wedge \mathcal{W}(k, r) \leq \mathcal{V}'(r)))$$

By rearranging terms,

$$\begin{aligned} & (\mathcal{W}(k, r) \neq 0 \wedge \mathcal{W}'(k, r) \neq 0) \vee (\mathcal{W}(k, r) \neq 0 \wedge \mathcal{W}(k, r) > \mathcal{V}'(r)) \vee \\ & (\mathcal{W}'(k, r) > \mathcal{V}(r) \wedge \mathcal{W}'(k, r) \neq 0) \vee (\mathcal{W}'(k, r) > \mathcal{V}(r) \wedge \mathcal{W}(k, r) > \mathcal{V}'(r)) \end{aligned}$$

We proceed by case analysis:

- $\mathcal{W}(k, r) \neq 0 \wedge \mathcal{W}'(k, r) \neq 0$. Since $((r, \mathcal{V}, \mathcal{W}), \langle G, P \rangle) \in \mathcal{I}$, $\mathcal{W}(k, r) > 0$ implies $\exists S. \mathcal{S}_{Or-Set}(G^{(\text{lookup}, S)}) \neq \emptyset \wedge k \in S$. By the definition of \mathcal{S}_{Or-Set} , $\exists e' \in \mathcal{E}_G$ such that $\lambda_G(e') = \langle \text{add}(k), \text{ok} \rangle$ and $\forall e''. e' <_G e''$ implies $\lambda_G(e'') \neq \langle \text{rem}(k), \text{ok} \rangle$. Moreover, since $\mathcal{W}'(k, r) \neq 0$, the definition of G' is such that the cases (i) and (ii) in \mathcal{F}_r^k only apply. Consequently, $\forall e''. e' <_{G \sqcup G'} e''$ implies $\lambda_{G \sqcup G'}(e'') \neq \langle \text{rem}(k), \text{ok} \rangle$. Therefore, $\mathcal{S}_{Or-Set}((G \sqcup G')^{(\text{lookup}, S)}) \neq \emptyset$ implies $k \in S$.
- $\mathcal{W}(k, r) \neq 0 \wedge \mathcal{W}(k, r) > \mathcal{V}'(r)$. If $\mathcal{W}'(k, r) \neq 0$, then the proof follows as in the previous case. Otherwise ($\mathcal{W}'(k, r) = 0$), the only possible case in the definition of \mathcal{F}_r^k is (iv) and we reason analogously to the previous case to conclude that $\mathcal{S}_{Or-Set}((G \sqcup G')^{(\text{lookup}, S)}) \neq \emptyset$ implies $k \in S$.
- $\mathcal{W}'(k, r) > \mathcal{V}(r) \wedge \mathcal{W}'(k, r) \neq 0$. Hence, only the case (ii) in the definition of \mathcal{F}_r^k applies. Consequently, for a_r^k it holds that $\forall e''. a_r^k <_{G \sqcup G'} e''$ implies $\lambda_{G \sqcup G'}(e'') \neq \langle \text{rem}(k), \text{ok} \rangle$. Therefore, $\mathcal{S}_{Or-Set}(G^{(\text{lookup}, S)}) \neq \emptyset$ implies $k \in S$.
- $\mathcal{W}'(k, r) > \mathcal{V}(r) \wedge \mathcal{W}(k, r) > \mathcal{V}'(r)$. Now note that $\mathcal{W}'(k, r) > \mathcal{V}(r)$ implies $\mathcal{W}'(k, r) \neq 0$ and $\mathcal{W}(k, r) > \mathcal{V}'(r)$ implies $\mathcal{W}(k, r) \neq 0$. Then, the proof follows as in case a).

Now, we consider the case in which $(\mathcal{V}, \mathcal{W}) \oplus (\mathcal{V}', \mathcal{W}')(k, r) = 0$ for all r . As before, we conclude that for all r , the following holds:

$$(\mathcal{W}(k, r) = 0 \wedge \mathcal{W}'(k, r) \leq \mathcal{V}(r)) \vee (\mathcal{W}'(k, r) = 0 \wedge \mathcal{W}(k, r) \leq \mathcal{V}'(r))$$

We proceed by case analysis for all r

- $\mathcal{W}(k, r) = 0 \wedge \mathcal{W}'(k, r) \leq \mathcal{V}(r)$. From $\mathcal{W}(k, r) = 0$ and $((r, \mathcal{V}, \mathcal{W}), \langle G, P \rangle) \in \mathcal{I}$ we deduce that there is not $e' \in \mathcal{E}_G$ such that $\lambda(e') = \langle \text{add}(k), \text{ok} \rangle$ and $\forall e''. e' <_G e''$ implies $\lambda(e'') \neq \langle \text{rem}(k), \text{ok} \rangle$. Moreover, the applicable cases in the definition of \mathcal{F}_r^k are (i), (iii), and (iv). Consequently, there is not $e' \in \mathcal{E}_r^k \cup \mathcal{F}_r^k$ such that $\lambda_{G \sqcup G'}(e') = \langle \text{add}(k), \text{ok} \rangle$ and for all e'' if $e' <_{G \sqcup G'} e''$ then $\lambda_{G \sqcup G'}(e'') \neq \langle \text{rem}(k), \text{ok} \rangle$.
- $\mathcal{W}'(k, r) = 0 \wedge \mathcal{W}(k, r) \leq \mathcal{V}'(r)$. This case corresponds to (iii) in the definition of \mathcal{F}_r^k . Hence, for all $e' \in \mathcal{E}_r^k \cup \mathcal{F}_r^k$ such that $\lambda(e') = \langle \text{add}(k), \text{ok} \rangle$ we have that $e' <_{G \sqcup G'} r^k$. Consequently, there is not $e' \in \mathcal{E}_r^k \cup \mathcal{F}_r^k$ such that $\lambda_{G \sqcup G'}(e') = \langle \text{add}(k), \text{ok} \rangle$ and for all e'' if $e' <_{G \sqcup G'} e''$ then $\lambda_{G \sqcup G'}(e'') \neq \langle \text{rem}(k), \text{ok} \rangle$.

Since the above two cases hold for all k and r , we conclude that

$$\forall r. (\mathcal{V}, \mathcal{W}) \oplus (\mathcal{V}', \mathcal{W}')(k, r) = 0 \Rightarrow \forall S. \mathcal{S}((G \sqcup G')^{(\text{lookup}, S)}) = \emptyset \vee k \notin S$$

For \Leftarrow), assume that there exists S and k such that $\mathcal{S}((G \sqcup G')^{(\text{lookup}, S)}) \neq \emptyset \wedge k \in S$. Then, there exists e' such that $\lambda_{G \sqcup G'}(e') = \langle \text{add}(k), \text{ok} \rangle$ and $\forall e''. e' <_{G \sqcup G'} e''$ implies $\lambda_{G \sqcup G'}(e'') \neq \langle \text{rem}(k), \text{ok} \rangle$. There are two cases:

- $e' \in \mathcal{E}_G$. By definition of $G \sqcup G'$, $\lambda_G(e') = \langle \text{add}(k), \text{ok} \rangle$ and $\forall e''. e' <_G e''$ implies $\lambda_G(e'') \neq \langle \text{rem}(k), \text{ok} \rangle$. Since, $((r, \mathcal{V}, \mathcal{W}), \langle G, P \rangle) \in \mathcal{I}$, then there exists r such that $\mathcal{W}(k, r) > 0$. Consequently, $\neg \text{IsREM}(\mathcal{W}, \mathcal{W}', \mathcal{V}, k, r)$ holds. Since $\forall e''. e' <_{G \sqcup G'} e''$ implies $\lambda_{G \sqcup G'}(e'') \neq \langle \text{rem}(k), \text{ok} \rangle$, we conclude that \mathcal{F}_r^k is such that either (i), (ii) or (iv) holds. For (i) or (ii), note that $\mathcal{W}'(k, s) \neq 0$ and, hence, $\neg \text{IsREM}(\mathcal{W}', \mathcal{W}, \mathcal{V}', k, r)$ holds. For (iii), $\mathcal{W}(k, r) \leq \mathcal{V}'(r)$ implies $\neg \text{IsREM}(\mathcal{W}', \mathcal{W}, \mathcal{V}', k, r)$. Consequently, $\neg \text{IsREM}(\mathcal{W}, \mathcal{W}', \mathcal{V}, k, r)$ and $\neg \text{IsREM}(\mathcal{W}', \mathcal{W}, \mathcal{V}', k, r)$. Hence,

$$(\mathcal{V}, \mathcal{W}) \oplus (\mathcal{V}', \mathcal{W}')(k, r) = \max\{\mathcal{W}(k, r), \mathcal{W}'(k, r)\}$$

Since $\mathcal{W}(k, r) > 0$, $(\mathcal{V}, \mathcal{W}) \oplus (\mathcal{V}', \mathcal{W}')(k, r) > 0$ holds.

- $e' \notin \mathcal{E}_G$. Then, $e' \notin \mathcal{E}_{G'}$. The only possibility is $e' = a_r^k$ for some \mathcal{F}_r^k (case (ii)). Hence, $\mathcal{W}'(k, r) \neq 0$ and $\mathcal{W}'(k, r) > \mathcal{V}(r)$. The case follows by noting that $\mathcal{W}'(k, r) \neq 0$ implies $\neg \text{IsREM}(\mathcal{W}', \mathcal{W}, \mathcal{V}', k, r)$ and $\mathcal{W}'(k, r) > \mathcal{V}(r)$ implies $\neg \text{IsREM}(\mathcal{W}, \mathcal{W}', \mathcal{V}, k, r)$.

Now, consider that there exists k such that for all S , if $\mathcal{S}((G \sqcup G')^{(\text{lookup}, S)}) \neq \emptyset$ then $k \notin S$. Consequently, for all $e' \in \mathcal{E}_{G \sqcup G'}$ if $\lambda_{G \sqcup G'}(e') = \langle \text{add}(k), \text{ok} \rangle$ then there exists e'' such that $e' \prec_{G \sqcup G'} e''$ and $\lambda_{G \sqcup G'}(e'') = \langle \text{rem}(k), \text{ok} \rangle$. Then, \mathcal{F}_r^k was obtained by using either (i), (iii) or (iv). In case (iii) is used, then $w'(k, r) = 0$ and $w(k, r) \leq v'(r)$. Therefore, $\text{isREM}(w', w, v', k, r)$ holds and $(v, w) \oplus (v', w')(k, r) = 0$. For cases (i) and (iv), we first note that $e'' \in \mathcal{E}_G$. Hence, for all S , if $\mathcal{S}(G^{(\text{lookup}, S)}) \neq \emptyset$ then $k \notin S$. Since $((r, v, w), (G, P)) \in \mathcal{I}$, we have $w(k, r) = 0$ for all r . Additionally, (i) implies $w'(k, r) \leq v(r)$ and, hence, $\text{isREM}(w, w', v, k, r)$ holds. Consequently, $(v, w) \oplus (v', w')(k, r) = 0$. Case (iv) also implies $w'(k, r) = 0$. Hence, $\max\{w(k, r), w'(k, r)\} = 0$ and $(v, w) \oplus (v', w')(k, r) = 0$. Therefore, $((r, \max\{v, v'\}), (w, v) \oplus (w', v')), (G \sqcup G', P') \in \mathcal{I}$ holds. \square

Proof of Lemma 7. We show that the following relation satisfies Definition 21.

$$\mathcal{I} = \{(\sigma_1 \parallel \sigma_2, \langle G_1 \sqcup G_2, P \rangle) \mid \sigma_1 \sim_S^0 \langle G_1, P_1 \rangle \text{ and } \sigma_2 \sim_S^0 \langle G_2, P_2 \rangle \text{ and } P \in P_1 \otimes P_2\}$$

We proceed by case analysis on the transitions of $\sigma_1 \parallel \sigma_2$.

- (PARL). Since $\sigma_1 \sim_S^0 \langle G_1, P_1 \rangle$, we have that $\langle G_1, P \rangle_{\mathcal{E}_{G_1}} \xrightarrow{\ell} \langle G'_1, P'_1 \rangle$. Then, the case follows by Lemma 6.
- (COMM). Since $\sigma_1 \xrightarrow{\text{send}, \sigma} \sigma'_1$ and $\sigma_1 \sim_S^0 \langle G_1, P_1 \rangle$, we have (i) $\sigma'_1 = \sigma_1$, (ii) $\sigma \sim_S^0 \langle G', P' \rangle$, (iii) $G_1 \sqcup G' = G_1$, and (iv) $P_1 \otimes P' = P_1$. Additionally, $\sigma_2 \xrightarrow{\text{rcv}, \sigma} \sigma'_2$ implies there exist G'', P'' such that (v) $\sigma \sim_S^0 \langle G'', P'' \rangle$, (vi) $P'_2 \in P_2 \otimes P''$ and (vii) $\sigma'_2 \sim_S^0 \langle G_2 \sqcup G'', P'_2 \rangle$. From (ii), (v) and (vii) we conclude that $\sigma'_2 \sim_S^0 \langle G_2 \sqcup G', P'_2 \rangle$. Moreover, $G_1 \sqcup (G_2 \sqcup G') = G_1 \sqcup G_2$ because of (iii). Hence, $(\sigma'_1 \parallel \sigma'_2, \langle G_1 \sqcup G_2, P \rangle) \in \mathcal{I}$. \square

Proof of Prop. 1. It follows as in the previous lemma by showing that \sim_S^m is preserved by parallel composition \parallel of states. \square

Proof of Lemma 8. Proof follows analogously to Lemma 7 by showing that the following relation satisfies Definition 21.

$$\mathcal{I} = \{(\sigma_1 \mid_B \sigma_2, \langle G_1 \sqcup G_2, P \rangle) \mid \sigma_1 \sim_S^0 \langle G_1, P_1 \rangle \text{ and } \sigma_2 \sim_S^0 \langle G_2, P_2 \rangle \text{ and } P \in P_1 \otimes P_2 \\ \text{and } \forall \sigma \in B. \exists G. (\sigma \sim_S^0 \langle G, P \rangle \text{ and } (G_1 \sqcup G_2) \sqcup G \text{ defined})\} \quad \square$$

Proof of Prop. 2. It follows as in the previous lemma by showing that \sim_S^m is preserved by parallel composition \mid_B of states. \square

References

- [1] S. Gilbert, N. Lynch, Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services, SIGACT News 33 (2) (2002) 51–59.
- [2] M. Shapiro, N.M. Preguiça, C. Baquero, M. Zawirski, Conflict-free replicated data types, in: X. Défago, F. Petit, V. Villain (Eds.), SSS 2011, in: Lect. Notes Comput. Sci., vol. 6976, Springer, 2011, pp. 386–400.
- [3] S. Burckhardt, A. Gotsman, H. Yang, M. Zawirski, Replicated data types: specification, verification, optimality, in: S. Jagannathan, P. Sewell (Eds.), POPL 2014, ACM, 2014, pp. 271–284.
- [4] S. Burckhardt, Principles of eventual consistency, Found. Trends Program. Lang. 1 (1–2) (2014) 1–150.
- [5] S. Burckhardt, A. Gotsman, H. Yang, Understanding eventual consistency, Tech. Rep. MSR-TR-2013-39, Microsoft Research, 2013.
- [6] A. Gotsman, H. Yang, Composite replicated data types, in: J. Vitek (Ed.), ESOP 2015, in: Lect. Notes Comput. Sci., vol. 9032, Springer, 2015, pp. 585–609.
- [7] A. Cerone, G. Bernardi, A. Gotsman, A framework for transactional consistency models with atomic visibility, in: L. Aceto, D. de Frutos-Escrig (Eds.), CONCUR 2015, in: LIPIcs, vol. 42, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2015, pp. 58–71.
- [8] K.C. Sivaramakrishnan, G. Kaki, S. Jagannathan, Declarative programming over eventually consistent data stores, in: D. Grove, S. Blackburn (Eds.), PLDI 2015, ACM, 2015, pp. 413–424.
- [9] A. Gotsman, H. Yang, C. Ferreira, M. Najafzadeh, M. Shapiro, Cause I'm strong enough: reasoning about consistency choices in distributed systems, in: R. Bodík, R. Majumdar (Eds.), POPL 2016, ACM, 2016, pp. 371–384.
- [10] M. Wirsing, Algebraic specification, in: Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics (B), Elsevier, 1990, pp. 675–788.
- [11] J.V. Guttag, E. Horowitz, D.R. Musser, The design of data type specifications, in: R.T. Yeh, C.V. Ramamoorthy (Eds.), ICSE 1976, IEEE Computer Society Press, 1976, pp. 414–420.
- [12] H.-D. Ehrich, On the theory of specification, implementation, and parametrization of abstract data types, J. ACM 29 (1) (1982) 206–227.
- [13] F. Gadducci, H.C. Melgratti, C. Roldán, A denotational view of replicated data types, in: J. Jacquet, M. Massink (Eds.), COORDINATION 2017, in: Lect. Notes Comput. Sci., vol. 10319, Springer, 2017, pp. 138–156.
- [14] M. Shapiro, N. Preguiça, C. Baquero, M. Zawirski, A comprehensive study of convergent and commutative replicated data types, Tech. Rep. RR-7506, Inria-Centre Paris-Rocquencourt, 2011.
- [15] D.B. Terry, M. Theimer, K. Petersen, A.J. Demers, M. Spreitzer, C. Hauser, Managing update conflicts in Bayou, a weakly connected replicated storage system, in: M.B. Jones (Ed.), SOSP 1995, ACM, 1995, pp. 172–183.
- [16] R. Jagadeesan, J. Riely, From sequential specifications to eventual consistency, in: M.M. Halldórsson, K. Iwama, N. Kobayashi, B. Speckmann (Eds.), ICALP 2015, in: Lect. Notes Comput. Sci., vol. 9135, Springer, 2015, pp. 247–259.
- [17] A. Bouajjani, C. Enea, J. Hamza, Verifying eventual consistency of optimistic replication systems, in: S. Jagannathan, P. Sewell (Eds.), POPL 2014, ACM, 2014, pp. 285–296.
- [18] K. von Gleissenthall, A. Rybalchenko, An epistemic perspective on consistency of concurrent computations, in: P.R. D'Argenio, H.C. Melgratti (Eds.), CONCUR 2013, in: Lect. Notes Comput. Sci., vol. 8052, Springer, 2013, pp. 212–226.
- [19] P. Zeller, A. Bieniusa, A. Poetzsch-Heffter, Formal specification and verification of CRDTs, in: E. Ábrahám, C. Palamidessi (Eds.), FORTE 2014, in: Lect. Notes Comput. Sci., vol. 8461, Springer, 2014, pp. 33–48.