

1 Coinductive Reasoning about CRDT Emulation

2 Nathan Liittschwager¹, Stelios Tsampas²,
3 Jonathan Castello¹, and Lindsey Kuper¹

4 ¹ University of California, Santa Cruz, USA

5 ² Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany

6 **Abstract.** Conflict-free replicated data types (CRDTs) are distributed
7 data structures designed for fault tolerance and high availability. CRDTs
8 have historically been taxonomized into operation-based (or op-based)
9 CRDTs and state-based CRDTs. The notion that state-based and op-
10 based CRDTs are equivalent is often appealed to in the literature. In
11 particular, verification techniques and results for one kind of CRDT are
12 often said to be applicable to the other kind, thanks to this equivalence.
13 However, while there are general algorithms for constructing a
14 state-based CRDT from a given op-based CRDT and vice versa, what it
15 means for one kind of CRDT to emulate another has never been made
16 fully precise. In this paper, we model CRDT systems as transition system
17 coalgebras, and argue that emulation can be understood formally in
18 terms of various *weak simulations* between the coalgebras of the original
19 and emulating CRDT systems. As a simple corollary, we deduce which
20 properties are preserved by the emulation algorithms, thus closing a gap
21 in the CRDT literature.

22 **Keywords:** CRDTs, coalgebras

23 1 Introduction

24 In distributed data storage systems, *data replication* is a ubiquitous mechanism
25 for guarding against machine failures and ensuring that data is physically close
26 to far-flung clients. Informally, replication takes an object and copies it over n
27 sites, or *nodes*, with each *replica* acting as an independent copy of the original
28 object. With replication comes the challenge of ensuring that replicas remain
29 (more or less) consistent with one another, especially in the face of inevitable
30 network partitions and clients who demand “always-on” access to data. The ideal
31 of consistency for such a replicated system is *linearizability* [18], under which
32 clients cannot tell whether they are interfacing with the original object, or the
33 replicated system. Unfortunately, linearizability is impractical to implement —
34 indeed, systems that prioritize *high availability* of data must necessarily do so at
35 the expense of linearizability [13, 14].

36 The quest for an optimal point in the availability/consistency trade-off space
37 has led to the development of *conflict-free replicated data types* (CRDTs) [43,
38 35, 34], which are data structures designed for replication and high availability.

39 CRDTs sacrifice linearizability in favor of a weaker *strong convergence* prop-
40 erty [43], which says that replicas that have received and applied the same *set*
41 of updates will agree in state, regardless of the order in which those updates
42 were received and applied. When coupled with a guarantee of *eventual delivery*
43 of updates to replicas, strong convergence ensures that replicas will *eventually*
44 come to agree. In the last decade, CRDTs have been an active area of research
45 in programming languages and verification communities, with considerable at-
46 tention paid to the formal specification and verification (of various properties,
47 but especially strong convergence) of CRDT designs [9, 45, 15, 12, 25, 31, 24, 32,
48 33], with recent work moving toward automated verification [30, 10] and even
49 synthesis of correct-by-construction CRDTs [22].

50 In their pioneering work on CRDTs, Shapiro et al. [43] taxonomize CRDTs
51 into *state-based* CRDTs, in which replicas apply updates locally and periodically
52 broadcast their local state to other replicas over the network, and *operation-based*
53 (or *op-based*) CRDTs, in which every state-updating operation is broadcast and
54 applied at each replica.³ In state-based CRDTs, the states that a replica can take
55 on must be elements of a join-semilattice, and a replica receiving an update from
56 a remote replica will apply the update by taking the least upper bound (join)
57 of its local state and the received update. Op-based CRDTs, on the other hand,
58 only require concurrent operations to commute, but rely on stronger ordering
59 guarantees (in particular, *causal broadcast* [8, 7]) from the underlying network
60 transport mechanism. Both the state-based and op-based approaches result in
61 strong convergence, the defining characteristic of CRDTs.

62 Most work on CRDT specification and verification focuses on either state-
63 based [45, 12, 31, 44, 33, 22] or op-based [15, 30, 25, 24, 32] CRDTs exclusively.
64 The justification for this choice is that state-based and op-based CRDTs can
65 *emulate* each other. Shapiro et al. [43] give general algorithms by which one may
66 construct a state-based CRDT out of a given op-based CRDT, and vice versa.
67 However, Shapiro et al. stop short of formally defining a notion of emulation
68 and proving that their construction satisfies it. Yet the notion that state-based
69 and op-based CRDTs can emulate each other is frequently appealed to in the
70 literature. For instance, Nagar and Jagannathan [30], in their work on verification
71 of op-based CRDTs, write that “our technique naturally extends to state-based
72 CRDTs since they can be emulated by an op-based model,” and Laddad et
73 al. [22], in their work on synthesis of state-based CRDTs, write that they “can
74 always be translated to op-based CRDTs if necessary.” Hence this emulation
75 notion is “load-bearing” and therefore deserving of being made precise.

76 In this paper, we seek to close this gap in the CRDT literature and formalize
77 the notion of emulation. To that end, we give a simple transition system model
78 of CRDT systems, modeling the network and interactions between replicas. This
79 model can be understood from the perspective of *universal coalgebra* [37, 21].
80 Given an endofunctor $F: \mathcal{C} \rightarrow \mathcal{C}$ in some category \mathcal{C} , F -coalgebras are pairs

³ More recently, Almeida et al. [3] introduce *delta state* CRDTs, an optimization of traditional state-based CRDTs in which only state changes, rather than entire states, must be disseminated over the network.

81 (X, h) of an object $X \in \mathcal{C}$ and a morphism $h: X \rightarrow FX$. The idea is that the
82 endofunctor F encapsulates the generic behavior of a system. This simple ab-
83 straction has proven to be remarkably general, capable of modelling systems such
84 as deterministic and non-deterministic automata, Mealy machines [36], prob-
85 abilistic systems [5] and higher-order languages [16] by varying the choice of
86 category \mathcal{C} and endofunctor F , yet powerful enough to be the foundation of a
87 unifying theory of bisimulation and coinduction. It is through this unifying the-
88 ory of coalgebra that we can reason coinductively about emulation of CRDTs.

89 Our contributions are (1) to demonstrate that the operational semantics of a
90 CRDT system can be modeled by an appropriate transition system (and, there-
91 fore, by a coalgebra), and (2) to formalize precisely what is meant by *emulation*
92 of CRDTs in terms of *weak simulations* between coalgebras. Our results give
93 researchers working on CRDTs a rigorous way to think about equivalence of
94 state-based and op-based CRDTs: in particular, properties that transfer from
95 one kind of CRDT to the other are precisely those properties that are preserved
96 by weak simulation relations.

97 The rest of this paper is organized as follows. After giving background on
98 CRDTs (Section 2), we present a semantics for a system of CRDT replicas and
99 make the straightforward observation that our semantics is a particular kind
100 of coalgebra (Section 3). We then review the formal notion of simulations in
101 our setting, and argue how it can be adapted to characterize *emulation* (Sec-
102 tion 4). We justify this choice of simulation by reasoning that emulation should
103 be about relating the observable behaviors of the original system and the emu-
104 lating system. We have kept the emulation algorithms of Shapiro et al. intact,
105 albeit with minor modifications. This results in two emulation *mappings*, one in
106 each direction of emulation. Each emulation mapping corresponds to *two* simu-
107 lation arguments, which we argue coinductively. We conclude with a discussion
108 of related work (Section 5).

109 2 Background on CRDTs

110 In this section, we recall the definitions of op-based and state-based CRDTs,
111 based on those of Shapiro et al. [43]. CRDTs are networks of communicating
112 replicas, but the definitions of Shapiro et al. specify only the *interface* that each
113 individual replica exposes to the network, as this fully determines the semantics,
114 i.e., the behavior, of the CRDT network as a whole. Here, what we refer to as
115 a *system* or a *network* is a collection of non-byzantine processes, which we call
116 *replicas*, that communicate with each other through an underlying asynchronous
117 communication protocol. Later, in Section 3, we will recast these definitions in
118 terms of coalgebras.

119 CRDTs are constructed in such a way as to guarantee *strong eventual consis-*
120 *tency* [43, §2.2] of replicas, the most salient aspect which is of *strong convergence*,
121 which says that replicas that have received the same (unordered) set of updates
122 have equivalent state.

123 **Assumption 1.** In what follows, we fix a set A of *commands* and a set B of
124 *observables*. We also assume that CRDT systems consist of $n \in \mathbb{N}$ replicas.

125 2.1 Op-based CRDTs

126 The definition below formalizes operation-based CRDTs, originally introduced
127 by Shapiro et al. [43, §2.4]. We refer to them simply as *op-based CRDTs*.

128 **Definition 2** (Op-based CRDT). An *op-based CRDT* is a tuple $(S, s_0, M, \mathbf{u}, \mathbf{t}, \mathbf{e}, \mathbf{q})$,
129 consisting of:

- 130 – A set S of *local states*.
- 131 – An initial state $s_0 : S$.
- 132 – A set M of *messages*.
- 133 – An *update* map $\mathbf{u} : S \times A \rightarrow S$.
- 134 – A *prepare-update* map $\mathbf{t} : S \times A \rightarrow M$.
- 135 – An *effect-update* map $\mathbf{e} : S \times M \rightarrow S$.
- 136 – A *query* map $\mathbf{q} : S \rightarrow B$.

137 The core principle behind op-based CRDTs is that replicas locally (and inde-
138 pendently of the other replicas) execute commands, which are then propagated
139 to the rest of the network via broadcast messages. The mechanism works as fol-
140 lows: suppose a given replica i , with a local state of $s_i \in S$, executes a command
141 $a \in A$, thus transitioning to the state $s'_i = \mathbf{u}(s_i, a)$. Replica i then prepares a
142 message $m = \mathbf{t}(s_i, a)$, which is broadcast to all other replicas. The broadcast
143 happens asynchronously, meaning messages are received by other replicas in an
144 indeterminate order. Finally, the replicas may “consume” the message m via the
145 effect-update map \mathbf{e} , e.g., replica j applies $\mathbf{e}(s_j, m)$. In this case, we say that the
146 message m is *delivered* at replica j . The query function \mathbf{q} exposes the observable
147 part of a replica’s local state to the external world.

148 **Notation 3.** Often we will consider applying *sequences* of messages. The fol-
149 lowing notation is thus useful. Let M^* denote the set of strings of messages
150 (where $(-)^*$ Kleene star). Let ε be the empty string and \cdot the usual string con-
151 catenation. $\forall w \in M^*$, define $\mathbf{e}_w : S \rightarrow S$ inductively by setting $\mathbf{e}_\varepsilon(s) = s$, and
152 $\mathbf{e}_{m \cdot w}(s) = \mathbf{e}_w(\mathbf{e}(s, m))$ for all $m \in M$.

153 **Remark 4** (The role of causality in op-based CRDTs). An important detail
154 about op-based CRDTs is that they must be implemented on top of a *causal*
155 *broadcast* mechanism. Causal broadcast ensures that when a broadcast message
156 m is delivered at a replica j , any message sent before m (in the specific sense
157 of Lamport’s *happens-before* partial order [23]) will have already been delivered
158 at j . (Messages not ordered by the happens-before relation, on the other hand,
159 may be delivered in arbitrary order.) Protocols for causal message delivery are
160 well known in the distributed systems literature [6, 39, 8, 7].

161 We express the causal broadcast requirement as follows: if M is the set of all
162 messages in a given execution of a CRDT system, then (M, \prec) is a partial order,

163 and \prec is called a *causality relation*, which encodes Lamport’s happens-before
 164 relation. Intuitively, if $m' \prec m$, then the sending of message m' is a potential
 165 “cause” for the sending of message m . Furthermore, we say that m and m' are
 166 *concurrent* (written $m \parallel m'$) if $\neg(m \prec m') \wedge \neg(m' \prec m)$. As an axiom, op-based
 167 CRDTs require that the effects of concurrent messages *commute*. That is, in
 168 Notation 3,

$$\forall s \in S. m \parallel m' \implies \mathbf{e}_{m \cdot m'}(s) = \mathbf{e}_{m' \cdot m}(s) \quad (1)$$

169 **Example 5.** A typical example of an op-based CRDT is the *grow-only set* or
 170 *G-Set*. Fix a set of objects Σ , and define the local states as elements $s \in \mathcal{P}(\Sigma)$.
 171 Initially, $s_0 = \emptyset$ and there is only one type of command: $\mathbf{add}(e)$ where $e \in \Sigma$.
 172 Then we define

- 173 – *prepare-update*: $\mathbf{t}(s, \mathbf{add}(e)) = (\mathbf{add}, e)$
- 174 – *effect-update*: $\mathbf{e}(s, (\mathbf{add}, e)) = s \cup \{e\}$
- 175 – *update*: $\mathbf{u}(s, \mathbf{add}(e)) = \mathbf{e}(s, \mathbf{t}(s, \mathbf{add}(e)))$

176 Queries enable clients to check if elements are in the set, meaning $B = 2^\Sigma$
 177 and $q : \mathcal{P}(\Sigma) \cong 2^\Sigma$. Furthermore, all updates commute, so causal delivery is
 178 in fact not necessary for this particular CRDT. However, more sophisticated
 179 op-based CRDTs such as *add-remove sets*, do require causal broadcast to ensure
 180 strong convergence.

181 2.2 State-based CRDTs

182 Next, we define state-based CRDTs [43, §2.3], in which local replica states form
 183 a join-semilattice.

184 **Definition 6** (State-based CRDT). A *state-based CRDT* is a tuple $((S, \sqcup), s_0, \mathbf{u}, \mathbf{q})$,
 185 consisting of:

- 186 – A join-semilattice (S, \sqcup) .
- 187 – An initial state $s_0 : S$.
- 188 – An *update* map $\mathbf{u} : S \rightarrow S^A$ that is furthermore *inflationary*, i.e.

$$s \leq \mathbf{u}(s, a) \text{ for all } s \in S \text{ and } a \in A.$$

- 189 – A *query* map $\mathbf{q} : S \rightarrow B$.

190 As is standard, we write $x \leq y$ whenever $x \sqcup y = y$.

191 A state-based CRDT operates as follows: as in op-based CRDTs, replicas
 192 may locally and independently execute commands $a \in A$. The synchronization
 193 mechanism, however, is radically different; whereas an op-based CRDT system
 194 uses a causal broadcast mechanism, in a state-based CRDT system any replica i
 195 may initiate a synchronization exchange with another replica j . In this exchange,
 196 replica i sends its entire local state s_i to replica j through the network. Upon
 197 arrival of the message (with payload s_i), replica j joins s_i with its current state
 198 s_j , i.e. $s'_j = s_i \sqcup s_j$. In that case, we say that (the message with payload) s_i is
 199 *delivered* at replica j .

200 **Example 7.** One kind of state-based CRDT with many applications in dis-
 201 tributed systems is a *Lamport clock* [23]. The join-semilattice is $(\mathbb{N}, \text{max} : \mathbb{N} \times$
 202 $\mathbb{N} \rightarrow \mathbb{N})$, where *max* takes its usual meaning. The initial state is $\text{clock}_0 = 0$. The
 203 command is simply *tick*, and the *update* is defined $u(\text{clock}, \text{tick}) = \text{clock} + 1$.
 204 Queries compare a given clock to the current *clock* that is the replica’s internal
 205 state. Then n replicas are seen as n distributed logical clocks, and may be used
 206 to track the order of events in a distributed system.

207 3 CRDTs, coalgebraically

208 In this section, we formalize the behavior of CRDTs as *non-deterministic tran-*
 209 *sition systems* on vectors of replica states. The following definition of a replica
 210 state covers both state-based and op-based CRDTs.

211 **Notation 8.** We write $\mathcal{M} : \mathbf{Set} \rightarrow \mathbf{Set}$ for the multiset functor.

212 **Definition 9** (Replica states). Let M and S be sets, representing resp. *messages*
 213 and *local states*. A *replica state* is a pair $(s, \sigma) \in S \times \mathcal{M}(M)$.

214 A replica state (s, σ) represents the runtime information of a single replica,
 215 which consists of the local state s and its *message buffer* σ . A vector of replica
 216 states thus represents the runtime information of a CRDT system; a network of
 217 replicas. The following definition captures the nature of a CRDT system and is
 218 independent of whether it is state-based or op-based.

219 **Definition 10.** Global State and Configurations Let \mathcal{E} be a set of events and
 220 $S \times \mathcal{M}(M)$ the set of replica states. We say a tuple $(x_i)_{i \in n} \in (S \times \mathcal{M}(M))^n$ is
 221 a *global state* of n replicas and $\langle \alpha, (x_i)_{i \in n} \rangle \in \mathcal{E} \times (S \times \mathcal{M}(M))^n$ is an *augmented*
 222 *global state* or *configuration*.

223 **Definition 11** (CRDT system). A *CRDT system* (on n replicas) is a pair
 224 (R, obs) consisting of a query map $\text{obs} : S \rightarrow B$ and a transition relation

$$R \subseteq \mathcal{E} \times (S \times \mathcal{M}(M))^n \times \mathcal{E} \times (S \times \mathcal{M}(M))^n.$$

225 **Remark 12.** CRDT systems form coalgebras in the category \mathbf{Set} of sets and
 226 total functions. In particular, a CRDT system (R, obs) is equivalently the coal-
 227 gebra

$$(h, \vec{\text{obs}}) : X \rightarrow \mathcal{P}(X) \times B^n$$

228 where $X = \mathcal{E} \times (S \times \mathcal{M}(M))^n$, $\vec{\text{obs}} = (\text{proj}_2 \circ \text{obs})^n$ and $\vec{y} \in h(\vec{x})$ if $R(\vec{x}, \vec{y})$.

229 From the coalgebraic perspective, the state space of CRDT systems cor-
 230 responds to an n -sized vector of replica states, tagged with event-labels; the
 231 observables are collected by the map $\text{obs} : S \rightarrow B$, acting on the local state of
 232 each replica. Recall that our CRDTs are always equipped with such a map.

233 Next, for both op-based and state-based CRDTs, we describe the semantics
 234 of CRDT systems, including how a local replica changes state, and how replicas
 235 interact with each other.

236 **3.1 Op-based CRDT system semantics**

237 **Assumption 13.** For the purposes of Section 3.1, we assume an op-based CRDT
 238 $(S, s_0, M, u, \tau, e, q)$.

239 We model the behavior of individual replicas in an op-based CRDT using a
 240 transition relation $\longrightarrow_{\text{op}}$ on replica states. For the op-based CRDT in Assump-
 241 tion 13, we pick the set of messages in replica states to be M . Specifically for
 242 op-based CRDTs, commands $a \in A$ also generate a *message* $m = \tau(s, a) \in M$ to
 243 be broadcast to other replicas. We denote a step that generates message m via
 244 command a with the notation $x_i \longrightarrow_{\text{op}} x'_i \uparrow (a, m)$. The replica semantics of the
 245 op-based CRDT is given by the following rules:

$$\frac{s' = u(s, a) \quad m = \tau(s, a)}{(s, \sigma) \longrightarrow_{\text{op}} (s', \sigma) \uparrow (a, m)} \text{ [OpLUpd]} \quad (2)$$

$$\frac{\exists m \in \sigma \quad \text{deliverable}(s, m) = \top \quad s' = e(s, m)}{(s, \sigma) \longrightarrow_{\text{op}} (s', \sigma \setminus \{m\})} \text{ [OpLRecv]}$$

246 In rule **OpLStep**, the replica executes a command $a \in A$, generating message
 247 m and changing the replica's state accordingly. In **OpLRecv**, the replica in ques-
 248 tion *delivers* a message $m \in \sigma$, i.e. it executes the underlying command included
 249 in the message ($\mathbf{emap}(s, m)$) and removes the message from its message buffer.
 250 We defer discussion of the $\text{deliverable}(s, m) = \top$ until Remark 16.

251 With the above replica semantics in mind, the *global* semantics of an op-
 252 based CRDT system models communication between replicas, with computation
 253 modeled by a transition relation on states augmented with *events* or *actions*,
 254 denoting what sort of computation step took place. More formally, if x, x' are
 255 global states and α, α' are events, then write $\langle \alpha, x \rangle \rightsquigarrow_{\text{op}} \langle \alpha', x' \rangle$ to mean that x
 256 transitions to x' via event α' .

257 To be precise, we now define the events in op-based CRDT systems.

258 **Definition 14.** For $i, j \in n$, $m \in M$ and $a \in A$,

$$\mathcal{E}_{\text{op}} \ni \alpha ::= \top \mid \mathbf{upd}^i(a, m) \mid \mathbf{bc}^i(m) \mid \mathbf{dlvr}^{j \leftarrow i}(m) \quad (3)$$

259 where \top is a special *null* event (for initial states), $\mathbf{upd}^i(a, m)$ denotes a local
 260 update at replica i with command a and emitting message m , $\mathbf{bc}^i(m)$ denotes
 261 the sending of m from replicas i to all other other replicas $j \in n$, and $\mathbf{dlvr}^{j \leftarrow i}(m)$
 262 denotes the delivery of message m at replica j , where the sender is replica i .

263 Write $\alpha \notin \mathbf{update}$ to mean α is not an update event. Then we define the
 264 global transition relation $\rightsquigarrow_{\text{op}}$ as follows:

$$\frac{\alpha \notin \mathbf{update} \quad x_i \longrightarrow_{\text{op}} x'_i \uparrow (a, m)}{\langle \alpha, (x_1, \dots, x_n) \rangle \rightsquigarrow_{\text{op}} \langle \mathbf{upd}^i(a, m), (x_1, \dots, x'_i, \dots, x_n) \rangle} \text{ [OpUpdate]}$$

$$\frac{(x'_1, \dots, x'_n) = \mathbf{bcast}_m^i(x_1, \dots, x_n)}{\langle \mathbf{upd}^i(a, m), (x_1, \dots, x_n) \rangle \rightsquigarrow_{\text{op}} \langle \mathbf{bc}^i(m), (x'_1, \dots, x'_n) \rangle} \text{ [OpBroadcast]} \quad (4)$$

$$\frac{\alpha \notin \mathbf{update} \quad x_i \longrightarrow_{\text{op}} x'_i}{\langle \alpha, (x_1, \dots, x_n) \rangle \rightsquigarrow_{\text{op}} \langle \mathbf{dlvr}^{j \leftarrow i}(m), (x_1, \dots, x'_i, \dots, x_n) \rangle} \text{ [OpRecv]}$$

265 In our model, transitions taken by each replica can be seen as atomic, local
 266 computations, and in a CRDT network, replicas perform local computations in
 267 no particular order. As such, rule **OpRecv** acknowledges that if a replica can
 268 take a step individually (by delivering a message via **OpLRecv**), it can also do
 269 so inside the network. Local updates at a replica induce a cascade of global
 270 transitions: when a replica updates its state via command a , it also does so inside
 271 the network. At the same time, the replica prepares a message m to be emitted
 272 to the network, and then does so immediately. This is captured by the **OpUpdate**
 273 and **OpBroadcast** rules in the global semantics. Note that the premise $\alpha \notin \text{send}$
 274 in **OpUpdate** and **OpRecv** requires that broadcast events always follow an update
 275 event. A broadcast $\text{bcast}_m^i(x_1, \dots, x_n)$ returns the global state (x'_1, \dots, x'_n) , where
 276 message m is placed in the message buffer of each replica j such that $j \neq i$.
 277 We stress that this communication is asynchronous: a message being placed in
 278 some replica's message buffer denotes that the message has simply been *sent*.
 279 Asynchrony is modeled by the fact that there can be an arbitrary number of
 280 steps between the sending/reception and the *delivery* (i.e., consumption) of the
 281 message.

282 **Definition 15.** Let $(x_i)_{i \in n}$ be a global state and α an event. We say a (possibly
 283 empty) sequence $\langle \alpha_i, z^i \rangle_{i \in k}$ of configurations is a *trace* of $\langle \alpha, (x_i)_{i \in n} \rangle$ if

$$(\langle \alpha, (x_i)_{i \in n} \rangle \rightsquigarrow_{\text{op}} \langle \alpha_1, z^1 \rangle) \wedge (\forall i \geq 1. \langle \alpha_i, z^i \rangle \rightsquigarrow_{\text{op}} \langle \alpha_{i+1}, z^{i+1} \rangle).$$

284 If all the events $(\alpha_i)_{i \in k}$ have a certain form (e.g., $\alpha_i = \text{dlvr}^{j \leftarrow l_i}(m_i)$) we may
 285 speak of *traces of events*.

286 **Remark 16.** We now turn to a more precise discussion of the deliverable premise
 287 in the **OpLRecv** rule. Recall from Remark 4 that op-based CRDTs are built on
 288 top of causal broadcast, and that the semantics of a causal broadcast mechanism
 289 induces a causality relation \prec where (i.) (M, \prec) is a partial order, and (ii.) causal
 290 broadcast guarantees that messages delivered at each replica are delivered in an
 291 order consistent with \prec .

292 We now show how such a relation \prec might be obtained in our semantics.
 293 Let $\langle \alpha_i, z^i \rangle_{i \in k}$ be a trace of initial state $\langle \top, (s^0, \emptyset)_{i \in n} \rangle$ (Definition 15), and take
 294 the set of events $E = \bigcup_{i \in k} \{\alpha_i\}$, which may be partitioned as $E = \bigcup_{i \in n} E_i$ by
 295 binning events α_i based on where they occurred (e.g., E_j includes all events of
 296 the form $\text{upd}^j(a)$, $\text{bc}^j(m)$ and $\text{dlvr}^{j \leftarrow i'}(m)$).

297 We then define a standard [23] happens-before relation \prec_{hb} as the smallest
 298 relation on events E satisfying

- 299 – $p < q$ and α_p and α_q are both events on the same replica (e.g., $\alpha_p, \alpha_q \in E_j$)
- 300 – $\alpha_p = \text{bc}^i(m)$ and $\alpha_q = \text{dlvr}^{j \leftarrow i}(m)$; and
- 301 – $\exists \alpha_r$ event s.t. $\alpha_p \prec_{\text{hb}} \alpha_r$ and $\alpha_r \prec_{\text{hb}} \alpha_q$.

302 The relation \prec_{hb} is a partial order relation and induces the partial order relation
 303 \prec on messages: $m' \prec m \iff \text{bc}^i(m') \prec_{\text{hb}} \text{bc}^j(m)$. We can understand the
 304 partial order (M, \prec) induced by the causal broadcast mechanism as essentially
 305 the abstraction of this construction.

306 A standard implementation strategy for such a mechanism [6, 7, 8] is to
 307 have the sender of a message augment the message with causal metadata (for
 308 instance, a *vector clock* [27, 11, 40]) that summarizes information about the
 309 causally preceding messages, and have each replica additionally maintain causal
 310 metadata as part of its state.

311 It is then the job of the causal broadcast mechanism to enforce *causal de-*
 312 *livery*: each $\text{dlvr}^{j \leftarrow i}(m)$ event can only occur if all messages $m' \prec m$ in the
 313 execution have already been delivered to replica j . This is done by inspecting
 314 the message's causal metadata and comparing it with the causal metadata in
 315 the replica's state to determine whether the message can be safely delivered or
 316 needs to be buffered for later delivery. We leave the details of the causal deliv-
 317 ery enforcement mechanism abstract, hiding all causal metadata, and instead
 318 capture its semantics using the deliverable relation and the partial order (M, \prec) .

319 In particular, $\text{deliverable}(s, m) = \top$ in the premise of OpLRecv means that
 320 all messages m' such that $m' \prec m$ have already been delivered at the replica in
 321 question.

322 **Remark 17.** The replica and global semantics of the op-based CRDT, coupled
 323 with the query map $q : S \rightarrow B$, forms the CRDT system $(\rightsquigarrow_{\text{op}}, q)$ in the sense
 324 of Definition 11.

325 3.2 State-based CRDT semantics

326 **Assumption 18.** For the purposes of Section 3.2, we assume a state-based
 327 CRDT $((S, \sqcup), s_0, u, q)$.

328 We proceed analogously to the op-based CRDT case. For the state-based
 329 CRDT $((S, \sqcup), s_0, u, q)$, we pick the set of messages in replica states to be S , as
 330 we are sending entire internal states to other replicas. The replica semantics is
 331 given by the following rules:

$$\begin{aligned}
 & \frac{a \in A \quad s' = u(s, a)}{(s, \sigma) \longrightarrow_{\text{st}} (s', \sigma)} \text{ [StLUpd]} \\
 & \frac{\exists s' \in \sigma}{(s, \sigma) \longrightarrow_{\text{st}} (s \sqcup s', \sigma \setminus \{s'\})} \text{ [StLRecv]}
 \end{aligned} \tag{5}$$

332 Compared to the op-based case, the semantics are quite similar at a local
 333 level, only now replicas are not required to broadcast as a consequence of an up-
 334 date. Delivering a message from another replica is just a join operation. Sending
 335 a state to another replica is given as a separate step in the global semantics,
 336 which are given below over the following events.

337 **Definition 19.** For $i, j \in n$, and $s \in S$, and $a \in A$,

$$\mathcal{E}_{\text{st}} \ni \alpha = \top \mid \text{upd}^i(a) \mid \text{send}^{i \rightarrow j}(s) \mid \text{dlvr}^{j \leftarrow i}(s) \tag{6}$$

338 where the events are analogous to those in Definition 14, with the caveat that now
 339 send events are point-to-point rather than broadcasts. Now, for events $\alpha \in \mathcal{E}_{\text{st}}$,
 340 we have the rules:

$$\begin{array}{c}
\frac{x_i = (s, \sigma) \quad x_i \longrightarrow_{\text{st}} x'_i \quad x'_i = (\mathbf{u}(s, a), \sigma)}{\langle \alpha, (x_1, \dots, x_i, \dots, x_n) \rangle \rightsquigarrow_{\text{st}} \langle \mathbf{upd}^i(a), (x_1, \dots, x'_i, \dots, x_n) \rangle} \text{ [StUpdate]} \\
\frac{x_i = (s_i, \sigma_i) \quad x_j = (s_j, \sigma_j) \quad x'_j = (s_j, \sigma_j \cup \{s_i\})}{\langle \alpha, (x_1, \dots, x_j, \dots, x_n) \rangle \rightsquigarrow_{\text{st}} \langle \mathbf{send}^{i \rightarrow j}(s_i), (x_1, \dots, x'_j, \dots, x_n) \rangle} \text{ [StSend]} \\
\frac{x_j = (s_j, \sigma) \quad x_j \longrightarrow_{\text{st}} x'_j \quad x'_j = (s_j \sqcup s, \sigma \setminus \{s\})}{\langle \alpha, (x_1, \dots, x_j, \dots, x_n) \rangle \rightsquigarrow_{\text{st}} \langle \mathbf{dlvr}^{j \leftarrow i}(s), (x_1, \dots, x'_j, \dots, x_n) \rangle} \text{ [StRecv]}
\end{array} \quad (7)$$

341 The global semantics for state-based CRDT systems are interpreted similarly
342 to op-based CRDT systems. The critical difference is that while op-based systems
343 are required to broadcast messages as a result of a local update, state-based
344 systems need no requirement. The rules **StUpdate** and **StRecv** show that that
345 replicas which take a step may also take a step inside the network. The rule
346 **StSend** models a spontaneous synchronization connection initiated by replica i
347 towards replica j with $i \neq j$. As we can see from the premise, the internal state
348 of x_i is placed in the message buffer of replica j . Again, the communication is
349 asynchronous, as replica j may consume this message at an arbitrary point in
350 the future.

351 **Remark 20.** As with op-based CRDTs, the replica and global semantics of the
352 state-based CRDT, coupled with the query map $\mathbf{q} : S \rightarrow B$, forms the CRDT
353 system $(\rightsquigarrow_{\text{st}}, \mathbf{q})$ in the sense of Definition 11.

354 **Remark 21.** The reason for augmenting the state with events, rather than using
355 a labeled transition (i.e., using behavior functor $\mathcal{P}(\mathcal{E} \times -)$) is because the event
356 types \mathcal{E}_{op} and \mathcal{E}_{st} in the op-based rules in Equation (4) and state-based rules in
357 Equation (7) are not identical up to relabeling. In the weak simulation arguments
358 we will present in Section 4, a single labeled event must be simulated by *multiple*
359 other labeled events. To this end, we augment the states with events, thus keeping
360 the behavior functor \mathcal{P} in both systems. Passing to the following intermediate
361 transition system for $(\rightsquigarrow_{\text{st}}, \mathbf{q}_{\text{st}})$ -systems is useful for relating events \mathcal{E}_{op} to those
362 in \mathcal{E}_{st} .

363 **Definition 22.** Set $\mathbf{q}' = \mathbf{q}_{\text{st}}$ and let \mathcal{E} be the set of events given by the syntax

$$\mathcal{E} \ni \alpha := \top \mid \mathbf{upd}^i(a) \mid \mathbf{bc}^i(s) \mid \mathbf{dlvr}^{j \leftarrow i}(s)$$

364 where $s \in S$. Then construct a transition system (\mapsto, \mathbf{q}'') by defining $\forall \alpha$ events,
365 $\forall z, z' \in (S \times \mathcal{M}(S))^n$ global states,

- 366 1. $\langle \alpha, z \rangle \mapsto \langle \mathbf{upd}^i(a), z' \rangle \iff \exists \alpha'. \langle \alpha', z \rangle \rightsquigarrow_{\text{st}} \langle \mathbf{upd}^i(a), z' \rangle$
- 367 2. $\langle \alpha, z \rangle \mapsto \langle \mathbf{dlvr}^{j \leftarrow i}(s), z' \rangle \iff \exists \alpha'. \langle \alpha', z \rangle \rightsquigarrow_{\text{st}} \langle \mathbf{dlvr}^{j \leftarrow i}(s), z' \rangle$.
- 368 3. $\langle \alpha, z \rangle \mapsto \langle \mathbf{bc}^i(s), z' \rangle \iff \exists \alpha', \alpha_1, \dots, \alpha_{n-1}$ events, $\exists z^1, \dots, z^{n-1}$ states
369 s.t.

$$(\langle \alpha', z \rangle \rightsquigarrow_{\text{st}} \langle \alpha_1, z^1 \rangle) \wedge (\forall j \geq 1. \langle \alpha_j, z^j \rangle \rightsquigarrow_{\text{st}} \langle \alpha_{j+1}, z^{j+1} \rangle)$$

370 where $z^{n-1} = z'$ and each $\alpha_j = \mathbf{send}^{i \rightarrow j}(s)$ s.t. $j \neq i$ ranges over n .

371 4 Emulation of CRDTs, coinductively

372 Shapiro et al. [43] argue that it is possible for an op-based CRDT to be *emulated*
373 by a state-based CRDT and vice versa. To that end, they provide two trans-
374 formations, one that constructs an op-based CRDT given a state-based CRDT,
375 and one that constructs a state-based CRDT given an op-based CRDT. The
376 translations are given in a precise manner, while the claim that op-based and
377 state-based CRDTs emulate each other is left informal. In particular, there is
378 no formal definition of “CRDT emulation”; rather, Shapiro et al. argue that
379 their constructions preserve the *strong eventual consistency* (SEC) property of
380 the original CRDTs. However, SEC preservation is insufficient for behavioral
381 equivalence, since, for instance, a trivial CRDT that does nothing is also SEC.

382 To close this gap, we begin this section by first introducing our notion of
383 CRDT emulation based on coalgebraic (*weak*) *simulation* (Section 4.1). Next, in
384 Section 4.2, we show that Shapiro et al.’s translation of state-based CRDTs to
385 op-based CRDTs is indeed an emulation in our sense, in that the resulting op-
386 based system weakly simulates the state-based system and vice versa. We present
387 the opposite direction in Section 4.3. It will become apparent that the two weak
388 simulations are non-trivial and underline the differences between state-based and
389 op-based CRDTs.

390 4.1 Emulation as coalgebraic simulation

391 We argue that the suitable notion to capture the relationship between op-based
392 and state-based CRDTs is (*weak*) *simulation*. We move on to define simulations
393 at the level of coalgebras for the endofunctor $B = L \times \mathcal{P}(-) : \mathbf{Set} \rightarrow \mathbf{Set}$, as our
394 CRDT systems are all instances of B -coalgebras (recall Remark 12, Remark 17
395 and Remark 20).

396 **Definition 23** (Simulations). Let $(X, \langle \varepsilon, h \rangle)$ and $(Y, \langle \zeta, g \rangle)$ be coalgebras for
397 endofunctor $B = L \times \mathcal{P}(-)$, for some set of observables L . A *simulation of*
398 (X, h) and (Y, g) is a relation $R \subseteq X \times Y$ such that for all pairs of states $x \in X$,
399 $y \in Y$ with $R(x, y)$, the following hold:

- 400 1. $\varepsilon(x) = \zeta(y)$.
- 401 2. If $x' \in h(x)$, then there exists $y' \in g(y)$ such that $R(x', y')$.

402 If $R(x, y)$, we say that y *simulates* x . Simulations are closed under arbitrary
403 unions: if R and Q are simulations, so is $R \cup Q$. The greatest simulation thus
404 exists and is the union of all simulations; this relation is known as *similarity*,
405 written as \lesssim .

406 **Definition 24** (Weak simulation). Let $(X, \langle \varepsilon, h \rangle)$ and $(Y, \langle \zeta, g \rangle)$ be coalgebras
407 for endofunctor $B = L \times \mathcal{P}(-)$, for some set of observables L , and let $g^* : Y \rightarrow$
408 $\mathcal{P}(Y)$ be the *reflexive, transitive closure* of $g : Y \rightarrow \mathcal{P}(Y)$. That is, g^* is the
409 least map $Y \rightarrow \mathcal{P}(Y)$ (w.r.t. pointwise inclusion), satisfying (i) $\forall y. g(y) \in g^*(y)$,
410 (ii) $\forall y. y \in g^*(y)$ and (iii) $\forall y, y', y''. y' \in g^*(y) \wedge y'' \in g^*(y') \implies y'' \in g^*(y)$. A
411 *weak simulation of* (X, h) and (Y, g) is a simulation of (X, h) and (Y, g^*) .

412 **Example 25.** Let $(\longrightarrow_1, \text{obs}_1)$ and $(\longrightarrow_2, \text{obs}_2)$ be two CRDT systems according
 413 to Definition 11, and let $(\vec{\text{obs}}_1, \longrightarrow_1) : X \rightarrow L \times \mathcal{P}X$ and $(\vec{\text{obs}}_2, \longrightarrow_2) : Y \rightarrow$
 414 $L \times \mathcal{P}Y$ be the induced coalgebras from Remark 12. Instantiating Definition 24
 415 to $(\vec{\text{obs}}_1, \longrightarrow_1)$ and $(\vec{\text{obs}}_2, \longrightarrow_2)$, a relation $R \subseteq X \times Y$ is weak simulation if the
 416 following is true for all $x \in X, y \in Y$ with $R(x, y)$:

$$\text{obs}_1(x) = \text{obs}_2(y) \wedge (\forall x. x \longrightarrow_1 x' \implies \exists y'. y \longrightarrow_2^* y' \wedge R(x', y')).$$

417 This says that two states x, y are related if they produce the same observables
 418 and the system \longrightarrow_2 can “match” the transitions of x , in potentially many steps.

419 **Notation 26.** Our coalgebras are transition systems where the state-space has
 420 been augmented with events, e.g., $\langle \alpha, x \rangle \in \mathcal{E}_1 \times X$ and $\langle \beta, y \rangle \in \mathcal{E}_2 \times Y$. We say
 421 a pair $(\approx, Q) \subseteq (\mathcal{E}_1 \times \mathcal{E}_2) \times (X \times Y)$ is a (weak) simulation if the relation R
 422 defined by $(\langle \alpha, x \rangle, \langle \beta, y \rangle) \in R \iff \alpha \approx \beta \wedge (x, y) \in Q$ is a (weak) simulation.

423 4.2 Emulation of state-based CRDTs by op-based CRDTs

424 Shapiro et al. provides a simple way to turn a state-based CRDT to an op-based
 425 CRDT. Given given any state-based CRDT

$$c = ((S, \sqcup), s_0, \mathbf{u}, \mathbf{q}), \quad (8)$$

426 we can construct the op-based CRDT $\mathcal{F}(c)$ by

$$c \xrightarrow{\mathcal{F}} (S, s_0, S, \mathbf{u}, \mathbf{u}, \sqcup, \mathbf{q}). \quad (9)$$

427 Indeed, the tuple $(S, s_0, S, \mathbf{u}, \mathbf{u}, \mathbf{q})$ is an instance of Definition 2. Applied to
 428 CRDTs c and $\mathcal{F}(c)$, the semantics introduced in Section 3.1 and Section 3.2 yield
 429 resp. the CRDT systems $(\rightsquigarrow_{\text{op}}, q)$ and $(\rightsquigarrow_{\text{st}}, q)$ which, according to Remarks
 430 17 and 20, induce the coalgebras

$$\begin{aligned} (\rightsquigarrow_{\text{op}}, \vec{q}) &: \mathcal{E}_{\text{op}} \times (S \times \mathcal{M}(S))^n \rightarrow \mathcal{P}(\mathcal{E}_{\text{op}} \times (S \times \mathcal{M}(S))^n) \times B^n \\ (\rightsquigarrow_{\text{st}}, \vec{q}) &: \mathcal{E}_{\text{st}} \times (S \times \mathcal{M}(S))^n \rightarrow \mathcal{P}(\mathcal{E}_{\text{st}} \times (S \times \mathcal{M}(S))^n) \times B^n. \end{aligned}$$

431 We now state the equivalence between CRDT systems c and $\mathcal{F}(c)$ in terms
 432 of two weak simulations: one of $(\rightsquigarrow_{\text{op}}, \vec{q})$ by $(\rightsquigarrow_{\text{st}}, \vec{q})$ and another one in the
 433 opposite direction.

434 First, construct the intermediate system as in Definition 22, whose transition
 435 arrows are denoted with \mapsto . The next theorem and its subsequent corollary are
 436 essentially direct, so we omit their proofs.

437 **Theorem 27.** Let $\Delta \subseteq \mathcal{E}_{\text{op}} \times (S \times \mathcal{M}(S))^n \times \mathcal{E}_{\text{op}} \times (S \times \mathcal{M}(S))^n$ be the diagonal
 438 relation, i.e.,

$$\Delta = \{ \langle \alpha, (x_i)_{i \in n} \rangle, \langle \alpha, (x_i)_{i \in n} \rangle \mid \alpha \in \mathcal{E}_{\text{op}} \wedge (x_i)_{i \in n} \in S \times \mathcal{M}(S)^n \}.$$

439 Relation Δ is a simulation of $(\rightsquigarrow_{\text{op}}, \vec{q})$ and (\mapsto, \vec{q}) .

440 **Corollary 28.** *There is a weak simulation \mathcal{Q}_1 of $(\rightsquigarrow_{\text{op}}, \vec{q})$ and $(\rightsquigarrow_{\text{st}}, \vec{q})$ arising*
 441 *from Δ which contains the initial states $\langle \top, (s^0, \emptyset)_{i \in n} \rangle \in \mathcal{E}_{\text{op}} \times (S \times \mathcal{M}(S))^n$*
 442 *and $\langle \top, (s^0, \emptyset)_{i \in n} \rangle \in \mathcal{E}_{\text{st}} \times (S \times \mathcal{M}(S))^n$.*

443 \mathcal{Q}_1 is a weak simulation of $(\rightsquigarrow_{\text{op}}, \vec{q})$ and $(\rightsquigarrow_{\text{st}}, \vec{q})$, but it is not a weak
 444 simulation of $(\rightsquigarrow_{\text{st}}, \vec{q})$ and $(\rightsquigarrow_{\text{op}}, \vec{q})$, as it is too strict a relation for this
 445 purpose. There is, however, a larger relation that is so. To define it, we use the
 446 following auxiliary definition.

447 **Definition 29.** Let $\mathcal{F}(c) = (S, s_0, S, \mathbf{u}, \mathbf{u}, \sqcup, \mathbf{q})$ be the aforementioned op-based
 448 emulator. We say a state $(s_i, \tau_i)_{i \in n}$ of $\mathcal{F}(c)$ is *synchronizable* if

$$\forall i, j \in n, C \subseteq \tau_i. \exists D \in \tau_j. \bigsqcup(\{s_i\} \cup \{s_j\} \cup C) = \bigsqcup(\{s_j\} \cup D). \quad (10)$$

449 Roughly, a state being synchronizable means that any replica can “catch up”
 450 with any other replica by delivering one or more messages from its buffer. The
 451 next proposition foreshadows how being synchronizable plays a role during a
 452 weak simulation.

453 **Proposition 30.** *Let $(x_i)_{i \in n}$ be a synchronizable state of the op-based CRDT*
 454 *system $(\rightsquigarrow_{\text{op}}, \mathbf{q})$ induced by $\mathcal{F}(c)$.*

455 *The following is true: if $\langle \alpha, (x_i)_{i \in n} \rangle \rightsquigarrow_{\text{op}}^* \langle \alpha', (x'_i)_{i \in n} \rangle$, and $\alpha' = \mathbf{bc}^j(m)$ or*
 456 *$\alpha' = \mathbf{dlvr}^{j \leftarrow k}(m)$, then $(x'_i)_{i \in n}$ is synchronizable. Otherwise $\alpha' = \mathbf{upd}^j(a)$, and*

$$\langle \mathbf{upd}^j(a), (x'_i)_{i \in n} \rangle \rightsquigarrow_{\text{op}} \langle \mathbf{bc}^j(m), z \rangle$$

457 *where $\langle \mathbf{bc}^j(m), z \rangle$ is synchronizable.*

458 We are now ready to construct our weak simulation relation. First, we define
 459 a relation $\mathcal{Q}_2 \in (S \times \mathcal{M}(S))^n \times (S \times \mathcal{M}(S))^n$ on the state spaces of $(\rightsquigarrow_{\text{st}}, \vec{q})$
 460 and $(\rightsquigarrow_{\text{op}}, \vec{q})$ as follows:

$$\begin{aligned} \mathcal{Q}_2 = \{ & ((s_i, \sigma_i)_{i \in n}, (s_i, \tau_i)_{i \in n}) \mid \\ & \forall i, j \in n. (\forall s \in \sigma_i. \exists B \subseteq \tau_i. s_i \sqcup s = \bigsqcup(\{s_i\} \cup B)) \\ & \wedge (s_i, \tau_i)_{i \in n} \text{ is synchronizable} \}. \end{aligned} \quad (11)$$

461 The intuition behind the \mathcal{Q}_2 relation is that when $(x, y) \in \mathcal{Q}_2$, the coalgebraic
 462 state y is “ahead” of x in terms of synchronization and, furthermore, it is always
 463 in the optimal synchronization state where all internal replica states converge, at
 464 the same internal state, after delivering all messages. Next, we let $\approx \subseteq \mathcal{E}_{\text{st}} \times \mathcal{E}_{\text{op}}$
 465 be a relation on events, defined as follows:

- 466 (i). $\mathbf{upd}^j(a) \approx \mathbf{bc}^j(m)$ for $a \in A, m \in M$,
- 467 (ii). $\mathbf{dlvr}^{j \leftarrow i}(s) \approx \mathbf{dlvr}^{j \leftarrow i}(s')$ for $s, s' \in S$,
- 468 (iii). $\mathbf{send}^{j \rightarrow i}(s) \approx \alpha$, for all $\alpha \in \mathcal{E}_{\text{op}}$.

469 **Theorem 31.** *Relation (\approx, \mathcal{Q}_2) is a weak simulation of $(\rightsquigarrow_{\text{st}}, \vec{q})$ and $(\rightsquigarrow_{\text{op}}, \vec{q})$.*

470 *Proof sketch.* Let $(x_i)_{i \in n} = (s_i, \sigma_i)_{i \in n}$ and $(y_i)_{i \in n} = (s_i, \tau_i)_{i \in n}$ such that for a
 471 pair of events $\beta \in \mathcal{E}_{\text{st}}, \alpha \in \mathcal{E}_{\text{op}}, ((x_i)_{i \in n}, (y_i)_{i \in n}) \in \mathcal{Q}_2$ and $\beta \approx \alpha$. By Definitions
 472 23 and 24, we have to show

$$\vec{q}((x_i)_{i \in n}) = \vec{q}((y_i)_{i \in n}) \quad (12)$$

473

$$\begin{aligned} & \langle \beta, (x_i)_{i \in n} \rangle \rightsquigarrow_{\text{st}} \langle \beta', w \rangle \\ \implies & \exists \alpha', z. \langle \alpha, (y_i)_{i \in n} \rangle \rightsquigarrow_{\text{op}}^* \langle \alpha', z \rangle \wedge \alpha' \approx \beta' \wedge (w, z) \in \mathcal{Q}_2. \end{aligned} \quad (13)$$

474 Equation (12) is immediate, as \vec{q} only depends on the internal states s_i , which
 475 are the same in both replica states x_i and y_i for all i . For (13), we proceed by
 476 case distinction on the global semantics of state-based systems (7). We identify
 477 and sketch three cases on the transition $\langle \beta, (x_i)_{i \in n} \rangle \rightsquigarrow_{\text{st}} \langle \beta', w \rangle$.

478 1. ($\beta' = \text{upd}^j(a)$). There are two sub-cases: if we have the sub-case $\alpha \notin \text{update}$,
 479 then simulate with the transitions

$$\langle \alpha, (y_i)_{i \in n} \rangle \rightsquigarrow_{\text{op}} \langle \text{upd}^j(a), (s'_i, \tau'_i)_{i \in n} \rangle \rightsquigarrow_{\text{op}} \langle \text{bc}^j(s'_j), z \rangle,$$

480 where $s'_j = \text{u}(s_j, a)$ in the previous step. Then $\text{upd}^j(a) \approx \text{bc}^j(m)$ and by
 481 Proposition 30, z is synchronizable. We thus have to show $\forall i \in n, i \neq j$,

$$\forall s \in \sigma_i. \exists B \subseteq \tau_i \cup \{s_j\}. s_i \sqcup s = \bigsqcup(\{s_i\} \cup B). \quad (14)$$

$$\forall s \in \sigma_j. \exists B \subseteq \tau_j. \text{u}(s_j, a) \sqcup s = \bigsqcup(\{\text{u}(s_j, a)\} \cup B). \quad (15)$$

482 Statement (14) immediately holds because of $((x_i)_{i \in n}, (y_i)_{i \in n}) \in \mathcal{Q}_2$. State-
 483 ment (15) follows from the fact that $\forall s \in S, s \sqcup \text{u}(s, a) = \text{u}(s, a)$ in con-
 484 junction with $((x_i)_{i \in n}, (y_i)_{i \in n}) \in \mathcal{Q}_2$. This closes the sub-case. For the other
 485 sub-case that $\alpha = \text{upd}^k(a')$, we must first complete the update-broadcast
 486 cycle (recall they are essentially atomic in the semantics 4), then perform
 487 the simulating transitions above, and apply the same argument. This close
 488 the second sub-case.

489 2. ($\beta' = \text{send}^{j \rightarrow i}(s)$). If $\alpha \in \text{update}$, then finish the update-broadcast cy-
 490 cle yielding configuration $\langle \alpha', z \rangle$ where α' is a broadcast event, and z is
 491 synchronizable by Proposition 30. Then simulate with the reflexive step
 492 $\langle \alpha', z \rangle \rightsquigarrow_{\text{op}}^* \langle \alpha', z \rangle$. The remaining details follow similarly to the previous
 493 case. Note that if $\alpha \notin \text{update}$ then the reflexive step alone suffices.

494 3. ($\beta' = \text{dlvr}^{j \leftarrow i}(s)$). We need to simulate by delivering the right sequence
 495 of messages. It thus suffices to only deal with the $\alpha \notin \text{update}$ sub-case,
 496 since finishing the update-broadcast cycle would yield buffers τ'_i which would
 497 contain the buffers τ_i as subsets had we not done this. Hence the proof for
 498 $\alpha \notin \text{update}$ implies the other subcase $\alpha \in \text{update}$.

499 The event $\text{dlvr}^{j \leftarrow i}(s)$ implies $s \in \sigma_j$, so from the hypothesis, there is a
 500 subset $B \subseteq \tau_j$ so that $s_j \sqcup s = \bigsqcup(\{s_j\} \cup B)$. Therefore, there is a (possibly
 501 empty) trace (Definition 15) $\langle \alpha_i, z^i \rangle_{i \in k}$ from $\langle \alpha, (y_i)_{i \in n} \rangle$ s.t. the $(\alpha_i)_{i \in k}$ are
 502 all deliver events on replica j , and thus correspond to the subset $B \subseteq \tau_j$. One
 503 can show that the resulting global state z^k is synchronizable, and replica j
 504 in z^k agrees with replica j in w . The simulating transitions is thus implied
 505 by the trace $\langle \alpha_i, z^i \rangle_{i \in k}$. \square

506 4.3 Emulation of op-based CRDTs by state-based CRDTs

507 The next translation is from op-based to state-based CRDTs, which is a slightly
 508 modified version of the one in Shapiro et al. [43]. The original translation takes a
 509 given op-based CRDT $o = (S, s_0, M, \mathbf{u}_{\text{op}}, \mathbf{t}, \mathbf{e}, \mathbf{q}_{\text{op}})$ and translates it into a state-
 510 based CRDT with replica state space $S \times \mathcal{P}_{\text{fin}}(M) \times \mathcal{P}_{\text{fin}}(M)$. The emulating
 511 replica state is thus a triple (s, H, D) of an internal state, a set of known messages
 512 H , and a set of delivered messages D . Merges on the emulating system work by
 513 creating a new set of known messages on the merged replica via set-theoretic
 514 union. Finally, updating is implemented in terms of some recursive function d
 515 which, apart from applying a given operation, also delivers qualified messages
 516 that are still in H .

517 In our modified translation, we observe that it is sufficient to simply use the
 518 sets $H \in \mathcal{P}_{\text{fin}}(M)$ as the emulating states, since one can recover the states s
 519 on-demand by computation (thus also eliminating the need for D). The idea
 520 is that, since M is equipped with the causal relation \prec (a partial order - see
 521 Remark 16), H essentially represents an equivalence class of strings of messages
 522 $m_1 m_2 \cdots m_{|H|}$, which all result in the same end state s when applied to the
 523 initial state s^0 .

524 *Op-based to state-based CRDT translation.* We show how the aforementioned H
 525 sets arise in our semantics.

526 Let $(S, s_0, M, \mathbf{u}_{\text{op}}, \mathbf{t}, \mathbf{e}, \mathbf{q}_{\text{op}})$ be an op-based CRDT, and $(\rightsquigarrow_{\text{op}}, \mathbf{q}_{\text{op}})$ its cor-
 527 responding system, taking $\langle \top, (s^0, \emptyset)_{i \in n} \rangle$ as the initial configuration. Following
 528 Remark 16, any trace $\langle \alpha_i, z^i \rangle_{i \in k}$ of $\langle \top, (s^0, \emptyset)_{i \in n} \rangle$ (Definition 15) yields a par-
 529 tially ordered (by causality) set E of events with partition $(E_i)_{i \in n}$.

530 Then there is an obvious map $E_i \mapsto H_i \in \mathcal{P}_{\text{fin}}(M)$ which projects the event
 531 sets down to finite message sets, inheriting a partial order \prec (we assume unique
 532 events project to unique messages), and thus also the concurrency relation \parallel .
 533 That is, $m \parallel m' \iff (m \not\prec m') \wedge (m' \not\prec m)$. We call H_i the *history* for replica
 534 i , and they can be thought of as denoting states $s_i \in S$ by the following lemma.

535 **Lemma 32.** *Let (M, \prec) be the partial order given by causality, and let $s^0 \in S$
 536 be a local state. Then, there is a well-defined map $\llbracket \cdot \rrbracket : \mathcal{P}_{\text{fin}}(M) \rightarrow S$.*

537 *Proof sketch.* Define the equivalence relation $w \equiv_{\mathbf{e}} w' \iff \mathbf{e}_w = \mathbf{e}_{w'}$ (Notation
 538 3). By Remark 4 and Remark 16, for all $m, m' \in M$ if $m \parallel m'$ then necessarily
 539 $m \cdot m' \equiv_{\mathbf{e}} m' \cdot m$. We note the following:

- 540 1. There is a map $H \mapsto \mathcal{L}(H)$, where $\mathcal{L}(H) \in \mathcal{P}_{\text{fin}}(M^*)$ is the set of *linear*
 541 *extensions* of \prec on H s.t. $\forall w, w' \in \mathcal{L}(H)$, we have $w \equiv_{\mathbf{e}} w'$.
- 542 2. Hence, there is a map $\mu : \mathcal{P}_{\text{fin}}(M) \rightarrow M^* / \equiv_{\mathbf{e}}$ which sends $H \mapsto [w]_{\equiv_{\mathbf{e}}}$, and
 543 an injection $g : M^* / \equiv_{\mathbf{e}} \rightarrow M$ which selects a representative $w' \in [w]_{\equiv_{\mathbf{e}}}$.
- 544 3. Finally, we define $\llbracket H \rrbracket = \mathbf{e}_{w'}(s^0)$ where $w' = (g \circ \mu)(H)$.

545 The map $\llbracket \cdot \rrbracket$ is well-defined in the sense that it does not depend on choice of
 546 $w' \in \mu(H)$. \square

547 **Corollary 33.** Let $\llbracket \cdot \rrbracket : \mathcal{P}_{fin}(M) \rightarrow S$ be given by Lemma 32, and let $H \in$
 548 $\mathcal{P}_{fin}(M)$ be a partially ordered set. For all $m \in M$,

$$(\forall m' \in H. (m' \prec m) \vee (m' \parallel m)) \implies \llbracket H \cup \{m\} \rrbracket = \mathbf{e}(\llbracket H \rrbracket, m).$$

549 *Proof sketch.* If w is any linear extension of (H, \prec) , and if $m' \prec m$ or $m' \parallel m$
 550 holds $\forall m' \in H$, then any linear extension of w' of $H \cup \{m\}$ belongs to the
 551 equivalence class $[w \cdot m]_{\equiv_o}$, hence $\mathbf{e}_{w'}(s^0) = \mathbf{e}_{w \cdot m}(s^0) = \mathbf{e}(\llbracket H \rrbracket, m)$. \square

552 In all that follows, we assume for simplicity fixed message set (M, \prec) and
 553 initial state s^0 . We can apply Lemma 32 to (M, \prec) and s^0 to obtain an *inter-*
 554 *pretation function* $\llbracket \cdot \rrbracket : \mathcal{P}_{fin}(M) \rightarrow S$ and construct the state-based CRDT $\mathcal{G}(o)$
 555 as follows:

$$o \xrightarrow{\mathcal{G}} ((\mathcal{P}_{fin}(M), \cup), \emptyset, \mathbf{u}_{st}, \mathbf{q}_{st}), \quad (16)$$

556 where $\mathbf{u}_{st}(H, a) = H \cup \{\mathbf{t}(\llbracket H \rrbracket, a)\}$ and $\mathbf{q}_{st}(H) = \mathbf{q}(\llbracket H \rrbracket)$. The fact that $\mathcal{G}(o)$ em-
 557 ulates o is a straightforward corollary of Lemma 32. In particular, definition (16)
 558 induces the CRDT systems $(\rightsquigarrow_{op}, \mathbf{q}_{op})$ and $(\rightsquigarrow_{st}, \mathbf{q}_{st})$, which weakly simulate
 559 one another.

560 The weak simulation argument consists of several steps. First, we construct
 561 an intermediate op-based CRDT system $(\longrightarrow, \mathbf{q})$ which behaves the same as the
 562 original system $(\rightsquigarrow_{op}, \mathbf{q}_{op})$, except now each replica state additionally carries
 563 not only state s_i , but also the partially ordered (by causality) set of d_i of *delivered*
 564 messages such that $\llbracket d_i \rrbracket = s_i$ (Lemma 32) holds as an *invariant*. Second, we also
 565 pass from $(\rightsquigarrow_{st}, \mathbf{q}_{st})$ to an intermediate system $(\dashrightarrow, \mathbf{q}')$ via Definition 22, also
 566 with histories H_i as the replica states.

567 We then show a weak simulation from $(\longrightarrow, \mathbf{q})$ to $(\dashrightarrow, \mathbf{q}')$ which maintains
 568 the equality $d_i = H_i$, hence $\llbracket H_i \rrbracket = s_i$.

569 We now construct the intermediate $(\longrightarrow, \mathbf{q})$ for the $(\rightsquigarrow_{op}, \mathbf{q}_{op})$ system. De-
 570 fine the map

$$o \mapsto (S \times \mathcal{P}_{fin}(M), (s_0, \emptyset), M, \mathbf{u}, \mathbf{t}', \mathbf{e}', \mathbf{q}) \quad (17)$$

571 where $\mathbf{t}'((s, d), a) = \mathbf{t}(s, a)$, and $\mathbf{e}'((s, d), m) = (\mathbf{e}(s, m), d \cup \{m\})$, and $\mathbf{u}(z, a) =$
 572 $\mathbf{e}'(z, \mathbf{t}'(z, a))$. Finally, set $\mathbf{q}(s, d) = \mathbf{q}_{op}(s)$. This induces an intermediate op-based
 573 system $(\longrightarrow, \mathbf{q})$ with the same set of events \mathcal{E}_{op} as $(\rightsquigarrow_{op}, \mathbf{q}_{op})$ system induced
 574 from o .

575 The next lemma expresses that $(\longrightarrow, \mathbf{q})$ has the same behavior as $(\rightsquigarrow_{op}, \mathbf{q}_{op})$.

576 **Lemma 34.** *The forgetful map*

$$f : \mathcal{E}_{op} \times (S \times \mathcal{P}_{fin}(M) \times \mathcal{M}(M))^n \rightarrow \mathcal{E}_{op} \times (S \times \mathcal{M}(M))^n$$

$$f(\alpha, (s_i, d_i, \sigma_i)_{i \in n}) = \langle \alpha, (s_i, \sigma_i)_{i \in n} \rangle$$

577 *is a coalgebra homomorphism, hence its graph a bisimulation.*

578 Lemma 32 and Corollary 33 establishes the relation between the s_i states
 579 and the delivered messages d_i , summarized as a lemma.

580 **Lemma 35.** Define $\text{Reach}(\alpha, z) = \{\langle \alpha', z' \rangle \mid \langle \alpha, z \rangle \longrightarrow^* \langle \alpha', z' \rangle\}$ as the set
 581 of reachable states from $\langle \alpha, z \rangle$. Then the predicate $P \subseteq \mathcal{E}_{\text{op}} \times (S \times \mathcal{P}_{\text{fin}}(M) \times$
 582 $\mathcal{M}(M))^n$ defined by

$$P = \{\langle \alpha, (s_i, d_i, \sigma_i)_{i \in n} \rangle \mid \forall i \in n. s_i = \llbracket d_i \rrbracket\} \cap \text{Reach}(\top, (s^0, \emptyset, \emptyset)) \quad (18)$$

583 is non-empty and an invariant of the intermediate system $(\longrightarrow, \mathbf{q})$, where $\llbracket \cdot \rrbracket : \mathcal{P}_{\text{fin}}(M) \rightarrow S$ is given by Lemma 32.

585 Lemmas 34 and 35 allow us to denote a replica state s_i by a set of delivered
 586 messages d_i under the op-based semantics. This informs us that the desired sim-
 587 ulation is that the state-based semantics can equivalently represent the delivered
 588 set d_i , though perhaps with multiple steps, and ignoring event labels.

589 To that end, apply Definition 22 and pass from $(\rightsquigarrow_{\text{st}}, \mathbf{q}_{\text{st}})$ to the interme-
 590 diate system (\mapsto, \mathbf{q}') (so that our events \mathcal{E} correspond to the broadcasting and
 591 delivering of *histories* of messages).

592 Our weak simulation argument is dependent on the relation $\approx \subseteq \mathcal{E}_{\text{op}} \times \mathcal{E}$
 593 on events, which we promptly define. Let $m \in M$ and define $H(m) \in \mathcal{P}_{\text{fin}}(M)$
 594 as the finite downward closed set of messages wrt \prec with m at the top.⁴ We
 595 inductively define $\approx \subseteq \mathcal{E}_{\text{op}} \times \mathcal{E}$ as follows:

$$\begin{aligned} \text{upd}^i(a) &\approx \text{upd}^i(a) \\ \text{bc}^i(m) &\approx \text{bc}^i(H(m)) \\ \text{dlvr}^{j \leftarrow i}(m) &\approx \text{dlvr}^{j \leftarrow i}(H(m)). \end{aligned} \quad (19)$$

596 We now give the statement of the main theorem and sketch its proof.

597 **Theorem 36.** Let $\mathcal{H} = \mathcal{P}_{\text{fin}}(M)$. For all states s_i , write $(s_i, d_i) \sim_1 H_i \iff$
 598 $\llbracket H_i \rrbracket = s_i = \llbracket d_i \rrbracket$. Also write $\sigma_i \sim_2 \tau_i$ if $m \in \sigma_i \iff H(m) \in \tau_i$. Then the pair
 599 of relations

$$(\approx, \mathcal{R}_1) \subseteq (\mathcal{E}_{\text{st}} \times \mathcal{E}) \times ((\mathcal{H} \times \mathcal{M}(\mathcal{H}))^n \times (S \times \mathcal{H} \times \mathcal{M}(M))^n)$$

600 where $\mathcal{R}_1 = \prod_{i \in n} R_i$, and

$$R_i = \{(s_i, d_i, \sigma_i), (H_i, \tau_i) \mid (s_i, d_i) \sim_1 H_i \wedge \sigma_i \sim_2 \tau_i\}, \quad (20)$$

601 is a simulation of $(\longrightarrow, \mathbf{q})$ and (\mapsto, \mathbf{q}') .

602 *Proof sketch.* On related configurations $\langle \alpha, (s_i, d_i, \sigma_i)_{i \in n} \rangle$ and $\langle \beta, (H_i, \tau_i)_{i \in n} \rangle$,
 603 proceed by case analysis on the transitions $\langle \alpha, (s_i, d_i, \sigma_i) \rangle \longrightarrow \langle \gamma, (y_i)_{i \in n} \rangle$ and
 604 find the simulating transition $\langle \beta, (H_i, \tau_i)_{i \in n} \rangle \mapsto \langle \phi, (z_i)_{i \in n} \rangle$. The cases are:

- 605 1. $\gamma = \text{upd}^j(a)$, then y_j has updated both s_j and d_j by generating locally and
 606 applying the message m . This same message m can be generated and applied
 607 locally from (H_j, τ_j) and event-transition $\phi = \text{upd}^j(a)$.

⁴ We may essentially think of $H(m)$ as a *causal history* [42] of m .

- 608 2. $\gamma = \mathbf{bc}^j(m)$, then each $y_i \neq y_j$ has input m in their buffer. From the op-
609 based rules, we must have $\alpha = \mathbf{upd}^j(a)$ as well and therefore $\beta = \mathbf{upd}^j(a)$.
610 By the previous case, and the semantics of causality, $H_j = H(m)$, so we
611 choose $\phi = \mathbf{bc}^j(H_j)$.
612 3. $\gamma = \mathbf{dlvr}^{j \leftarrow k}(m)$, then $m \in \sigma_j$ and was deliverable. Apply the hypothesis
613 and do the corresponding $\phi = \mathbf{dlvr}^{j \leftarrow k}(H(m))$ event. \square

614 The following corollary coincides with weak simulation, and is ultimately
615 what we wanted to show.

616 **Corollary 37.** *For every execution $\langle \top, (s^0, \emptyset)_{i \in n} \rangle \rightsquigarrow_{\text{op}}^* \langle \alpha, (s_i, \sigma_i)_{i \in n} \rangle$, there is
617 a corresponding execution $\langle \top, (\emptyset, \emptyset)_{i \in n} \rangle \rightsquigarrow_{\text{st}}^* \langle \beta, (H_i, \tau_i)_{i \in n} \rangle$ where $\forall i \in n. \llbracket H_i \rrbracket =$
618 s_i . Hence every $(\rightsquigarrow_{\text{op}}, \mathbf{q}_{\text{op}})$ behavior yields a $(\rightsquigarrow_{\text{st}}, \mathbf{q}_{\text{st}})$ behavior.*

619 The other weak simulation can be argued more directly. It's proof is remark-
620 ably similar to that of Theorem 31. Define the relation \approx on events $\mathcal{E}_{\text{st}} \times \mathcal{E}_{\text{op}}$ as
621 follows:

- 622 (i). $\mathbf{upd}^j(a) \approx \mathbf{bc}^j(m)$ for $a \in A, m \in M$,
623 (ii). $\mathbf{dlvr}^{j \leftarrow i}(H) \approx \mathbf{dlvr}^{j \leftarrow i}(m)$ for $H \in \mathcal{P}_{\text{fin}}(M)$ and $m \in M$,
624 (iii). $\mathbf{send}^{j \rightarrow i}(H) \approx \alpha$, for all $\alpha \in \mathcal{E}_{\text{op}}$.

625 **Theorem 38.** *Define the relation \mathcal{R}_2 on global states as*

$$\mathcal{R}_2 = \prod_{i \in n} \{(H_i, \tau_i), (s_i, d_i, \sigma_i) \mid (s_i, d_i) \sim_1 H_i \wedge \tau_i \sim_2 \sigma_i\}$$

626 where \sim_1 is as in Theorem 36 and

$$\sigma_i \sim_2 \tau_i \iff \forall H \in \tau_i. \exists \{m_1, \dots, m_k\} \subseteq \sigma_i. \llbracket H_i \cup H \rrbracket = \mathbf{e}_{m_1 \dots m_k}(s_i).$$

627 Then (\approx, \mathcal{R}_2) is a weak simulation from $(\rightsquigarrow_{\text{st}}, \mathbf{q}_{\text{st}})$ to $(\rightsquigarrow_{\text{op}}, \mathbf{q}_{\text{op}})$.

628 *Proof sketch.* Follow the proof structure of Theorem 31. The only tricky case is
629 when we must deliver a sequence of messages $m_1 \dots m_k$ to simulate the delivery
630 of a set $H \in \mathcal{P}_{\text{fin}}(M)$. The key insight is that H itself contains only a subset of
631 messages $O \subseteq H$ which produce an observable effect in the computation. That
632 is, $\llbracket H_i \cup H \rrbracket = \llbracket H \cup O \rrbracket$. The set O of messages is the set $\{m_1, \dots, m_k\}$ we are
633 looking for. \square

634 5 Related Work

635 *Specification and verification of CRDTs.* The primary focus of the most exist-
636 ing research on verified CRDTs has been the verification of strong convergence
637 and other safety properties for either state-based [45, 12, 31, 44, 33, 22] or op-
638 based [15, 30, 25, 24, 32] CRDTs.⁵ One exception is the work of Burckhardt et

⁵ Timany et al. [44] also consider verification of *liveness* properties, such as eventual delivery of messages.

639 al. [9], who give a framework for axiomatic specification and verification of both
640 op-based and state-based CRDTs, inspired by previous work on axiomatization
641 of weak memory models.

642 To our knowledge, our work is the first to take the approach of modeling
643 CRDTs with coalgebra in mind. However, the use of (bi)simulation relations
644 in CRDT verification is not new. For instance, Burckhardt et al. [9]’s frame-
645 work is based on *replication-aware simulations*, and Nair et al. [31] use a strong
646 bisimulation argument to justify the use of simpler, easier-to-implement proof
647 rules in an automated verification tool for state-based CRDTs, using the more
648 complicated semantics as a reference implementation.

649 Nieto et al. [33] observe that whether a CRDT is op-based or state-based is
650 in fact an implementation detail that should be hidden from clients, and their
651 verification approach centers around this *representation independence* property.
652 They contribute the first mechanically verified representation independence re-
653 sult for a specific CRDT, the *pn-counter*, demonstrating that a particular client
654 program cannot distinguish between (handwritten) op-based and state-based
655 implementations of this CRDT. Our goal in this work, on the other hand, is to
656 make precise the sense in which Shapiro et al. [43]’s *general* op-to-state-based and
657 state-to-op-based emulation algorithms result in the same observable behavior
658 for the original and the emulating object. However, we have not yet attempted
659 any proof mechanization. More generally, mechanized verification of state-based
660 and op-based CRDTs, both interactive [45, 15, 44, 32, 33] and automated [31,
661 30, 10, 22], is an active area of research. Our goal in this work is to complement
662 these existing verification efforts by making precise the sense in which results for
663 op-based CRDTs can be said to transfer to state-based CRDTs and vice versa.

664 *Coalgebraic reasoning about concurrent and distributed systems.* The theory of
665 universal coalgebra [20, 37, 36, 16] is a general model for state-based systems,
666 and can be seen as a generalization of classical work on process calculi [38, 17,
667 28, 29]. The typical notion of equivalence when modeling a system as a process
668 calculus or as a coalgebra is that of bisimulation [38, 28, 29], since classical
669 definitions of bisimulation on process calculi generalize nicely to coalgebras as
670 Hermida-Jacobs bisimulation [20], or Aczel-Mendler bisimulations [2, 37].

671 The transition system coalgebras we give here are presented on a mostly se-
672 mantic level, i.e., there is no corresponding process calculus. Thus our models
673 of CRDTs are closer in spirit to the the classical state-machine models of dis-
674 tributed computing theory [23, 41, 4, 26]. Moreover, the concept of simulation
675 (as opposed to bisimulation) in coalgebra and classical distributed computing
676 theory is not new [19, 4, 1], although, unlike in the present work, classical dis-
677 tributed computing theory does not make use of the aforementioned coalgebraic
678 notion.

679 **References**

- 680 [1] Martin Abadi and Leslie Lamport. “The Existence of Refinement Map-
681 pings”. In: *Proceedings of the 3rd Annual Symposium on Logic in Com-*
682 *puter Science*. LICS 1988 Test of Time Award. 1988, pp. 165–175. URL:
683 [https://www.microsoft.com/en-us/research/publication/the-](https://www.microsoft.com/en-us/research/publication/the-existence-of-refinement-mappings/)
684 [existence-of-refinement-mappings/](https://www.microsoft.com/en-us/research/publication/the-existence-of-refinement-mappings/).
- 685 [2] Peter Aczel and Nax Mendler. “A final coalgebra theorem”. In: *Category*
686 *Theory and Computer Science*. Ed. by David H. Pitt et al. Berlin, Hei-
687 delberg: Springer Berlin Heidelberg, 1989, pp. 357–365. ISBN: 978-3-540-
688 46740-3.
- 689 [3] Paulo Sérgio Almeida, Ali Shoker, and Carlos Baquero. “Efficient State-
690 Based CRDTs by Delta-Mutation”. In: *Networked Systems*. Ed. by Ahmed
691 Bouajjani and Hugues Fauconnier. Cham: Springer International Publish-
692 ing, 2015, pp. 62–76. ISBN: 978-3-319-26850-7.
- 693 [4] Hagit Attiya and Jennifer Welch. *Distributed Computing: Fundamentals,*
694 *Simulations and Advanced Topics*. Hoboken, NJ, USA: John Wiley & Sons,
695 Inc., 2004. ISBN: 0471453242.
- 696 [5] Falk Bartels, Ana Sokolova, and Erik P. de Vink. “A hierarchy of proba-
697 bilistic system types”. In: *6th International Workshop on Coalgebraic Meth-*
698 *ods in Computer Science, CMCS 2003, Satellite Event for ETAPS 2003,*
699 *Warsaw, Poland, April 5-6, 2003*. Ed. by H. Peter Gumm. Vol. 82. Elec-
700 tronic Notes in Theoretical Computer Science 1. Elsevier, 2003, pp. 57–75.
701 DOI: [10.1016/S1571-0661\(04\)80632-7](https://doi.org/10.1016/S1571-0661(04)80632-7). URL: [https://doi.org/10.](https://doi.org/10.1016/S1571-0661(04)80632-7)
702 [1016/S1571-0661\(04\)80632-7](https://doi.org/10.1016/S1571-0661(04)80632-7).
- 703 [6] K. Birman and T. Joseph. “Exploiting Virtual Synchrony in Distributed
704 Systems”. In: *SIGOPS Oper. Syst. Rev.* 21.5 (Nov. 1987), pp. 123–138.
705 ISSN: 0163-5980. DOI: [10.1145/37499.37515](https://doi.org/10.1145/37499.37515). URL: [https://doi.org/](https://doi.org/10.1145/37499.37515)
706 [10.1145/37499.37515](https://doi.org/10.1145/37499.37515).
- 707 [7] Kenneth Birman, André Schiper, and Pat Stephenson. “Lightweight Causal
708 and Atomic Group Multicast”. In: *ACM Trans. Comput. Syst.* 9.3 (Aug.
709 1991), pp. 272–314. ISSN: 0734-2071. DOI: [10.1145/128738.128742](https://doi.org/10.1145/128738.128742). URL:
710 <https://doi.org/10.1145/128738.128742>.
- 711 [8] Kenneth P. Birman and Thomas A. Joseph. “Reliable Communication in
712 the Presence of Failures”. In: *ACM Trans. Comput. Syst.* 5.1 (Jan. 1987),
713 pp. 47–76. ISSN: 0734-2071. DOI: [10.1145/7351.7478](https://doi.org/10.1145/7351.7478). URL: [https://doi.](https://doi.org/10.1145/7351.7478)
714 [org/10.1145/7351.7478](https://doi.org/10.1145/7351.7478).
- 715 [9] Sebastian Burckhardt et al. “Replicated Data Types: Specification, Verifi-
716 cation, Optimality”. In: *Proceedings of the 41st ACM SIGPLAN-SIGACT*
717 *Symposium on Principles of Programming Languages*. POPL ’14. San
718 Diego, California, USA: Association for Computing Machinery, 2014, pp. 271–
719 284. ISBN: 9781450325448. DOI: [10.1145/2535838.2535848](https://doi.org/10.1145/2535838.2535848). URL: [https:](https://doi.org/10.1145/2535838.2535848)
720 [//doi.org/10.1145/2535838.2535848](https://doi.org/10.1145/2535838.2535848).
- 721 [10] Kevin De Porre, Carla Ferreira, and Elisa Gonzalez Boix. “VeriFx: Cor-
722 rect Replicated Data Types for the Masses”. In: *37th European Conference*
723 *on Object-Oriented Programming (ECOOP 2023)*. Ed. by Karim Ali and

- 724 Guido Salvaneschi. Vol. 263. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum
725 für Informatik, 2023, 9:1–9:45. ISBN: 978-3-95977-281-5. DOI: [10.4230/
726 LIPIcs.ECOP.2023.9](https://doi.org/10.4230/LIPIcs.ECOP.2023.9). URL: [https://drops.dagstuhl.de/opus/
727 volltexte/2023/18202](https://drops.dagstuhl.de/opus/volltexte/2023/18202).
- 728
- 729 [11] C. J. Fidge. “Timestamps in message-passing systems that preserve the
730 partial ordering”. In: *Proceedings of the 11th Australian Computer Science
731 Conference* 10.1 (1988), pp. 56–66.
- 732 [12] Fabio Gadducci, Hernán Melgratti, and Christian Roldán. “On the seman-
733 tics and implementation of replicated data types”. In: *Science of Com-
734 puter Programming* 167 (2018), pp. 91–113. ISSN: 0167-6423. DOI: [https:
735 //doi.org/10.1016/j.scico.2018.06.003](https://doi.org/10.1016/j.scico.2018.06.003). URL: [https://www.
736 sciencedirect.com/science/article/pii/S0167642318302429](https://www.sciencedirect.com/science/article/pii/S0167642318302429).
- 737 [13] Seth Gilbert and Nancy Lynch. “Brewer’s Conjecture and the Feasibility of
738 Consistent, Available, Partition-Tolerant Web Services”. In: *SIGACT News*
739 33.2 (2002), pp. 51–59. ISSN: 0163-5700. DOI: [10.1145/564585.564601](https://doi.org/10.1145/564585.564601).
740 URL: <https://doi.org/10.1145/564585.564601>.
- 741 [14] Seth Gilbert and Nancy Lynch. “Perspectives on the CAP Theorem”. In:
742 *Computer* 45.2 (2012), pp. 30–36. DOI: [10.1109/MC.2011.389](https://doi.org/10.1109/MC.2011.389).
- 743 [15] Victor B. F. Gomes et al. “Verifying Strong Eventual Consistency in Dis-
744 tributed Systems”. In: *Proc. ACM Program. Lang.* 1.OOPSLA (2017). DOI:
745 [10.1145/3133933](https://doi.org/10.1145/3133933). URL: <https://doi.org/10.1145/3133933>.
- 746 [16] Sergey Goncharov et al. “Towards a Higher-Order Mathematical Opera-
747 tional Semantics”. In: *Proc. ACM Program. Lang.* 7.POPL (2023). DOI:
748 [10.1145/3571215](https://doi.org/10.1145/3571215). URL: <https://doi.org/10.1145/3571215>.
- 749 [17] Ichiro Hasuo, Bart Jacobs, and Ana Sokolova. “The Microcosm Principle
750 and Concurrency in Coalgebra”. In: *Foundations of Software Science and
751 Computational Structures*. Ed. by Roberto Amadio. Berlin, Heidelberg:
752 Springer Berlin Heidelberg, 2008, pp. 246–260. ISBN: 978-3-540-78499-9.
- 753 [18] Maurice P. Herlihy and Jeannette M. Wing. “Linearizability: A Correctness
754 Condition for Concurrent Objects”. In: *ACM Trans. Program. Lang. Syst.*
755 12.3 (1990), pp. 463–492. ISSN: 0164-0925. DOI: [10.1145/78969.78972](https://doi.org/10.1145/78969.78972).
756 URL: <https://doi.org/10.1145/78969.78972>.
- 757 [19] Jesse Hughes and Bart Jacobs. “Simulations in coalgebra”. In: *Theor. Com-
758 put. Sci.* 327.1-2 (2004), pp. 71–108. DOI: [10.1016/J.TCS.2004.07.022](https://doi.org/10.1016/J.TCS.2004.07.022).
759 URL: <https://doi.org/10.1016/j.tcs.2004.07.022>.
- 760 [20] Bart Jacobs. *Introduction to Coalgebra: Towards Mathematics of States
761 and Observation*. Cambridge Tracts in Theoretical Computer Science. Cam-
762 bridge University Press, 2016.
- 763 [21] Bart Jacobs. *Introduction to Coalgebra: Towards Mathematics of States
764 and Observation*. Vol. 59. Cambridge Tracts in Theoretical Computer Sci-
765 ence. Cambridge University Press, 2016. ISBN: 9781316823187. DOI: [10.
766 1017/CB09781316823187](https://doi.org/10.1017/CB09781316823187). URL: <https://doi.org/10.1017/CB09781316823187>.

- 767 [22] Shadaj Laddad et al. “Katara: Synthesizing CRDTs with Verified Lift-
768 ing”. In: *Proc. ACM Program. Lang.* 6.OOPSLA2 (2022). DOI: [10.1145/
769 3563336](https://doi.org/10.1145/3563336). URL: <https://doi.org/10.1145/3563336>.
- 770 [23] Leslie Lamport. “Time, Clocks, and the Ordering of Events in a Distributed
771 System”. In: *Commun. ACM* 21.7 (July 1978), pp. 558–565. ISSN: 0001-
772 0782. DOI: [10.1145/359545.359563](https://doi.org/10.1145/359545.359563). URL: [http://doi.acm.org/10.
773 1145/359545.359563](http://doi.acm.org/10.1145/359545.359563).
- 774 [24] Hongjin Liang and Xinyu Feng. “Abstraction for Conflict-Free Replicated
775 Data Types”. In: *Proceedings of the 42nd ACM SIGPLAN International
776 Conference on Programming Language Design and Implementation*. PLDI
777 2021. Virtual, Canada: Association for Computing Machinery, 2021, pp. 636–
778 650. ISBN: 9781450383912. DOI: [10.1145/3453483.3454067](https://doi.org/10.1145/3453483.3454067). URL: [https:
779 //doi.org/10.1145/3453483.3454067](https://doi.org/10.1145/3453483.3454067).
- 780 [25] Yiyun Liu et al. “Verifying Replicated Data Types with Typeclass Re-
781 finements in Liquid Haskell”. In: *Proc. ACM Program. Lang.* 4.OOPSLA
782 (2020). DOI: [10.1145/3428284](https://doi.org/10.1145/3428284). URL: [https://doi.org/10.1145/
783 3428284](https://doi.org/10.1145/3428284).
- 784 [26] Nancy A Lynch and Mark R Tuttle. *An introduction to input/output au-
785 tomata*. 1988.
- 786 [27] Friedemann Mattern. “Virtual Time and Global States of Distributed
787 Systems”. In: *Parallel and Distributed Algorithms*. North-Holland, 1989,
788 pp. 215–226.
- 789 [28] R. Milner. *A Calculus of Communicating Systems*. Berlin, Heidelberg:
790 Springer-Verlag, 1982. ISBN: 0387102353.
- 791 [29] Robin Milner, Joachim Parrow, and David Walker. “A calculus of mobile
792 processes, I”. In: *Information and Computation* 100.1 (1992), pp. 1–40.
793 ISSN: 0890-5401. DOI: [https://doi.org/10.1016/0890-5401\(92\)90008-
794 4](https://doi.org/10.1016/0890-5401(92)90008-4). URL: [https://www.sciencedirect.com/science/article/pii/
795 0890540192900084](https://www.sciencedirect.com/science/article/pii/S0890540192900084).
- 796 [30] Kartik Nagar and Suresh Jagannathan. “Automated Parameterized Veri-
797 fication of CRDTs”. In: *Computer Aided Verification*. Ed. by Isil Dillig and
798 Serdar Tasiran. Cham: Springer International Publishing, 2019, pp. 459–
799 477. ISBN: 978-3-030-25543-5.
- 800 [31] Sreeja S. Nair, Gustavo Petri, and Marc Shapiro. “Proving the Safety of
801 Highly-Available Distributed Objects”. In: *Programming Languages and
802 Systems*. Ed. by Peter Müller. Cham: Springer International Publishing,
803 2020, pp. 544–571. ISBN: 978-3-030-44914-8.
- 804 [32] Abel Nieto et al. “Modular Verification of Op-Based CRDTs in Separation
805 Logic”. In: *Proc. ACM Program. Lang.* 6.OOPSLA2 (2022). DOI: [10.1145/
806 3563351](https://doi.org/10.1145/3563351). URL: <https://doi.org/10.1145/3563351>.
- 807 [33] Abel Nieto et al. “Modular Verification of State-Based CRDTs in Sepa-
808 ration Logic”. In: *37th European Conference on Object-Oriented Program-
809 ming (ECOOP 2023)*. Ed. by Karim Ali and Guido Salvaneschi. Vol. 263.
810 Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Ger-
811 many: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023, 22:1–

- 812 22:27. ISBN: 978-3-95977-281-5. DOI: [10.4230/LIPIcs.ECOOP.2023.22](https://doi.org/10.4230/LIPIcs.ECOOP.2023.22).
813 URL: <https://drops.dagstuhl.de/opus/volltexte/2023/18215>.
- 814 [34] Nuno Preguiça, Carlos Baquero, and Marc Shapiro. “Conflict-Free Replicated
815 Data Types CRDTs”. In: *Encyclopedia of Big Data Technologies*. Ed.
816 by Sherif Sakr and Albert Zomaya. Cham: Springer International Publishing,
817 2018, pp. 1–10. ISBN: 978-3-319-63962-8. DOI: [10.1007/978-3-319-63962-8_185-1](https://doi.org/10.1007/978-3-319-63962-8_185-1). URL: https://doi.org/10.1007/978-3-319-63962-8_185-1.
- 818
819
- 820 [35] Hyun-Gul Roh et al. “Replicated abstract data types: Building blocks for
821 collaborative applications”. In: *Journal of Parallel and Distributed Computing* 71.3 (2011), pp. 354–368. ISSN: 0743-7315. DOI: <https://doi.org/10.1016/j.jpdc.2010.12.006>. URL: <https://www.sciencedirect.com/science/article/pii/S0743731510002716>.
- 822
823
824
- 825 [36] Jan J. M. M. Rutten. “Algebraic Specification and Coalgebraic Synthesis
826 of Mealy Automata”. In: *Proceedings of the International Workshop on
827 Formal Aspects of Component Software, FACS 2005, Macao, October 24-
828 25, 2005*. Ed. by Zhiming Liu and Luís Soares Barbosa. Vol. 160. Electronic
829 Notes in Theoretical Computer Science. Elsevier, 2005, pp. 305–319. DOI:
830 [10.1016/j.entcs.2006.05.030](https://doi.org/10.1016/j.entcs.2006.05.030). URL: <https://doi.org/10.1016/j.entcs.2006.05.030>.
- 831
- 832 [37] Jan J. M. M. Rutten. “Universal coalgebra: a theory of systems”. In: *Theor.
833 Comput. Sci.* 249.1 (2000), pp. 3–80. DOI: [10.1016/S0304-3975\(00\)00056-6](https://doi.org/10.1016/S0304-3975(00)00056-6). URL: [https://doi.org/10.1016/S0304-3975\(00\)00056-6](https://doi.org/10.1016/S0304-3975(00)00056-6).
- 834
- 835 [38] Davide Sangiorgi. *Introduction to Bisimulation and Coinduction*. Cambridge
836 University Press, 2011.
- 837 [39] André Schiper, Jorge Egli, and Alain Sandoz. “A New Algorithm to Implement
838 Causal Ordering”. In: *Proceedings of the 3rd International Workshop on
839 Distributed Algorithms*. Berlin, Heidelberg: Springer-Verlag, 1989,
840 pp. 219–232. ISBN: 3540516875.
- 841 [40] Frank B Schmuck. “The use of efficient broadcast protocols in asynchronous
842 distributed systems”. PhD thesis. Cornell University, 1988.
- 843 [41] Fred B. Schneider. “Implementing Fault-Tolerant Services Using the State
844 Machine Approach: A Tutorial”. In: *ACM Comput. Surv.* 22.4 (1990),
845 pp. 299–319. ISSN: 0360-0300. DOI: [10.1145/98163.98167](https://doi.org/10.1145/98163.98167). URL: <https://doi.org/10.1145/98163.98167>.
- 846
- 847 [42] Reinhard Schwarz and Friedemann Mattern. “Detecting causal relationships
848 in distributed computations: In search of the holy grail”. In: *Distributed
849 computing* 7.3 (1994), pp. 149–174.
- 850 [43] Marc Shapiro et al. “Conflict-Free Replicated Data Types”. In: *Stabilization,
851 Safety, and Security of Distributed Systems*. Ed. by Xavier Défago,
852 Franck Petit, and Vincent Villain. Berlin, Heidelberg: Springer Berlin Heidelberg,
853 2011, pp. 386–400. ISBN: 978-3-642-24550-3.
- 854 [44] Amin Timany et al. “Trillium: Higher-Order Concurrent and Distributed
855 Separation Logic for Intensional Refinement”. In: *Proc. ACM Program.*

856 *Lang.* 8.POPL (2024). DOI: [10.1145/3632851](https://doi.org/10.1145/3632851). URL: [https://doi.org/](https://doi.org/10.1145/3632851)
857 [10.1145/3632851](https://doi.org/10.1145/3632851).
858 [45] Peter Zeller, Annette Bieniusa, and Arnd Poetsch-Heffter. “Formal Speci-
859 fication and Verification of CRDTs”. In: *Formal Techniques for Distributed*
860 *Objects, Components, and Systems*. Ed. by Erika Ábrahám and Catuscia
861 Palamidessi. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 33–
862 48. ISBN: 978-3-662-43613-4.