

# From Convergence to Confidence

## Push-button verification for Replicated Data Types

**KC Sivaramakrishnan**

Joint work with Pranav Ramesh, Vimala Soundarapandian, Adarsh Kamath, Aseem Rastogi, Kartik Nagar

**PaPoC 2026**  
**27th April 2026**



# Following up on Pranav's talk from the morning

SAL: Multi-modal Verification of Replicated Data  
Types

Pranav R Vimala S KC Sivaramakrishnan

Indian Institute of Technology, Madras

April 23, 2026

***Why MRDTs?***

# Observed-Removed (OR) set CRDT

- **OR set**

- A concurrent set where add-wins in a concurrent **add(e)** and **rem(e)**

Concrete state  $:= (\Sigma_a, \Sigma_r) \xrightarrow{\text{add}(a)} (\Sigma_a \cup \{(a, id)\}, \Sigma_r)$  where id is fresh

$$(\Sigma_a, \Sigma_r) \xrightarrow{\text{rem}(a)} (\Sigma_a, \Sigma_r \cup \{(a, id) \mid (a, id) \in \Sigma_a\})$$

$$(\Sigma_a, \Sigma_r) \xrightarrow{\text{read}()} \{a \mid (a, id) \in \Sigma_a \setminus \Sigma_r\}$$



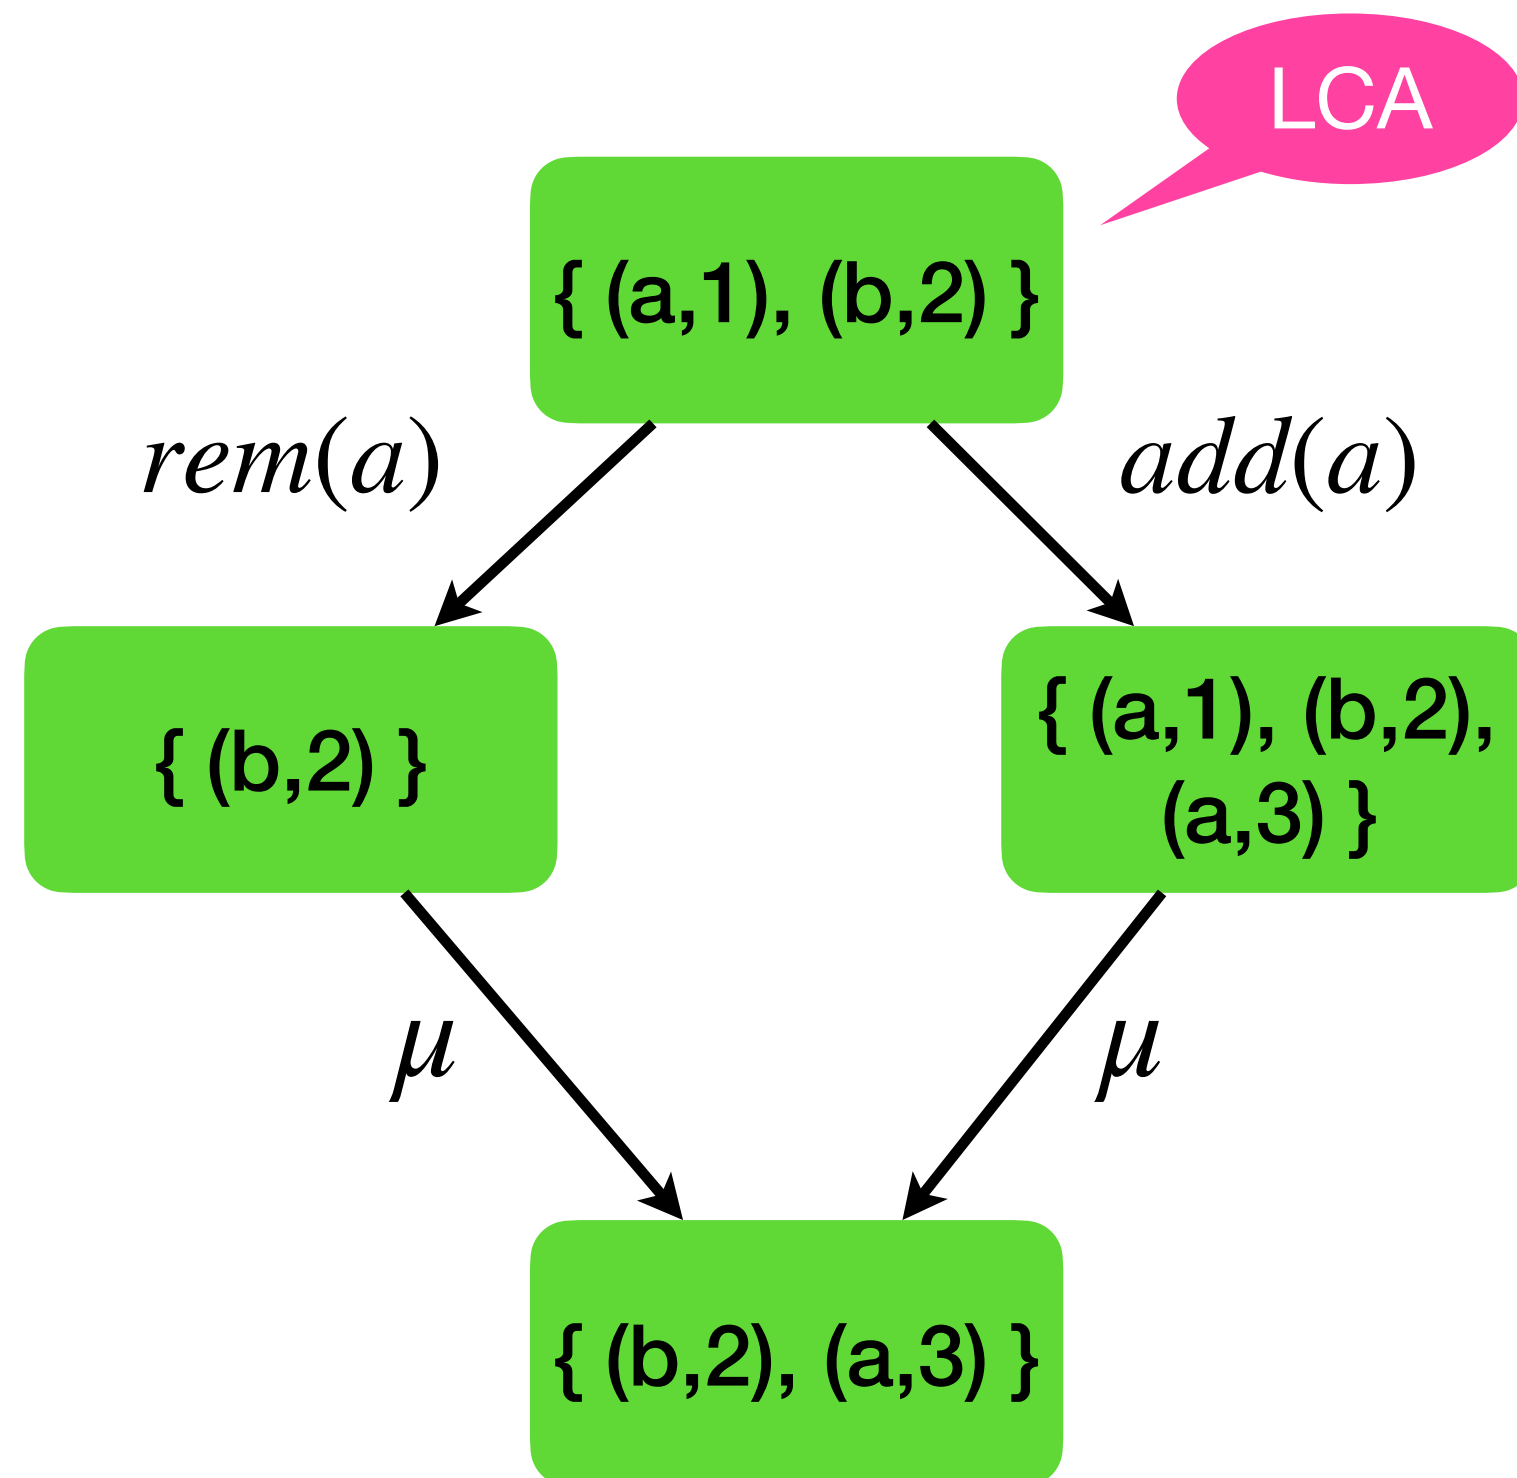
**Tombstones**

- States are asynchronously broadcast to other replicas

$$(\Sigma_a, \Sigma_r) \xrightarrow{\text{merge}(\Sigma'_a, \Sigma'_r)} (\Sigma_a \cup \Sigma'_a, \Sigma_r \cup \Sigma'_r)$$

- Space usage  $O(|\text{adds}|)$


# OR Set MRDT



- Associate *timestamp (id)* with the element
  - Add associates a *fresh* timestamp with the element
  - Remove removes all matching elements with *any id*
    - *Removed from the concrete state!*
  - Read returns the set removing ids
- ```
let merge lca v1 v2 =  
  (lca n v1 n v2) (* unchanged or removed elements *)  
  u (v1 \ lca) (* added elements on left *)  
  u (v2 \ lca) (* added elements on right *)
```
- Space usage  $O(|adds \cap live|)$ 
    - Duplicates — same element with different ids


# Runtime Support

- How did we eliminate tombstones?
  - Swept under the rug — MRDT runtime.
- MRDT runtime
  - Git-like distributed store with branching and merging
  - Support 3-way merge with lowest common ancestor (LCA)



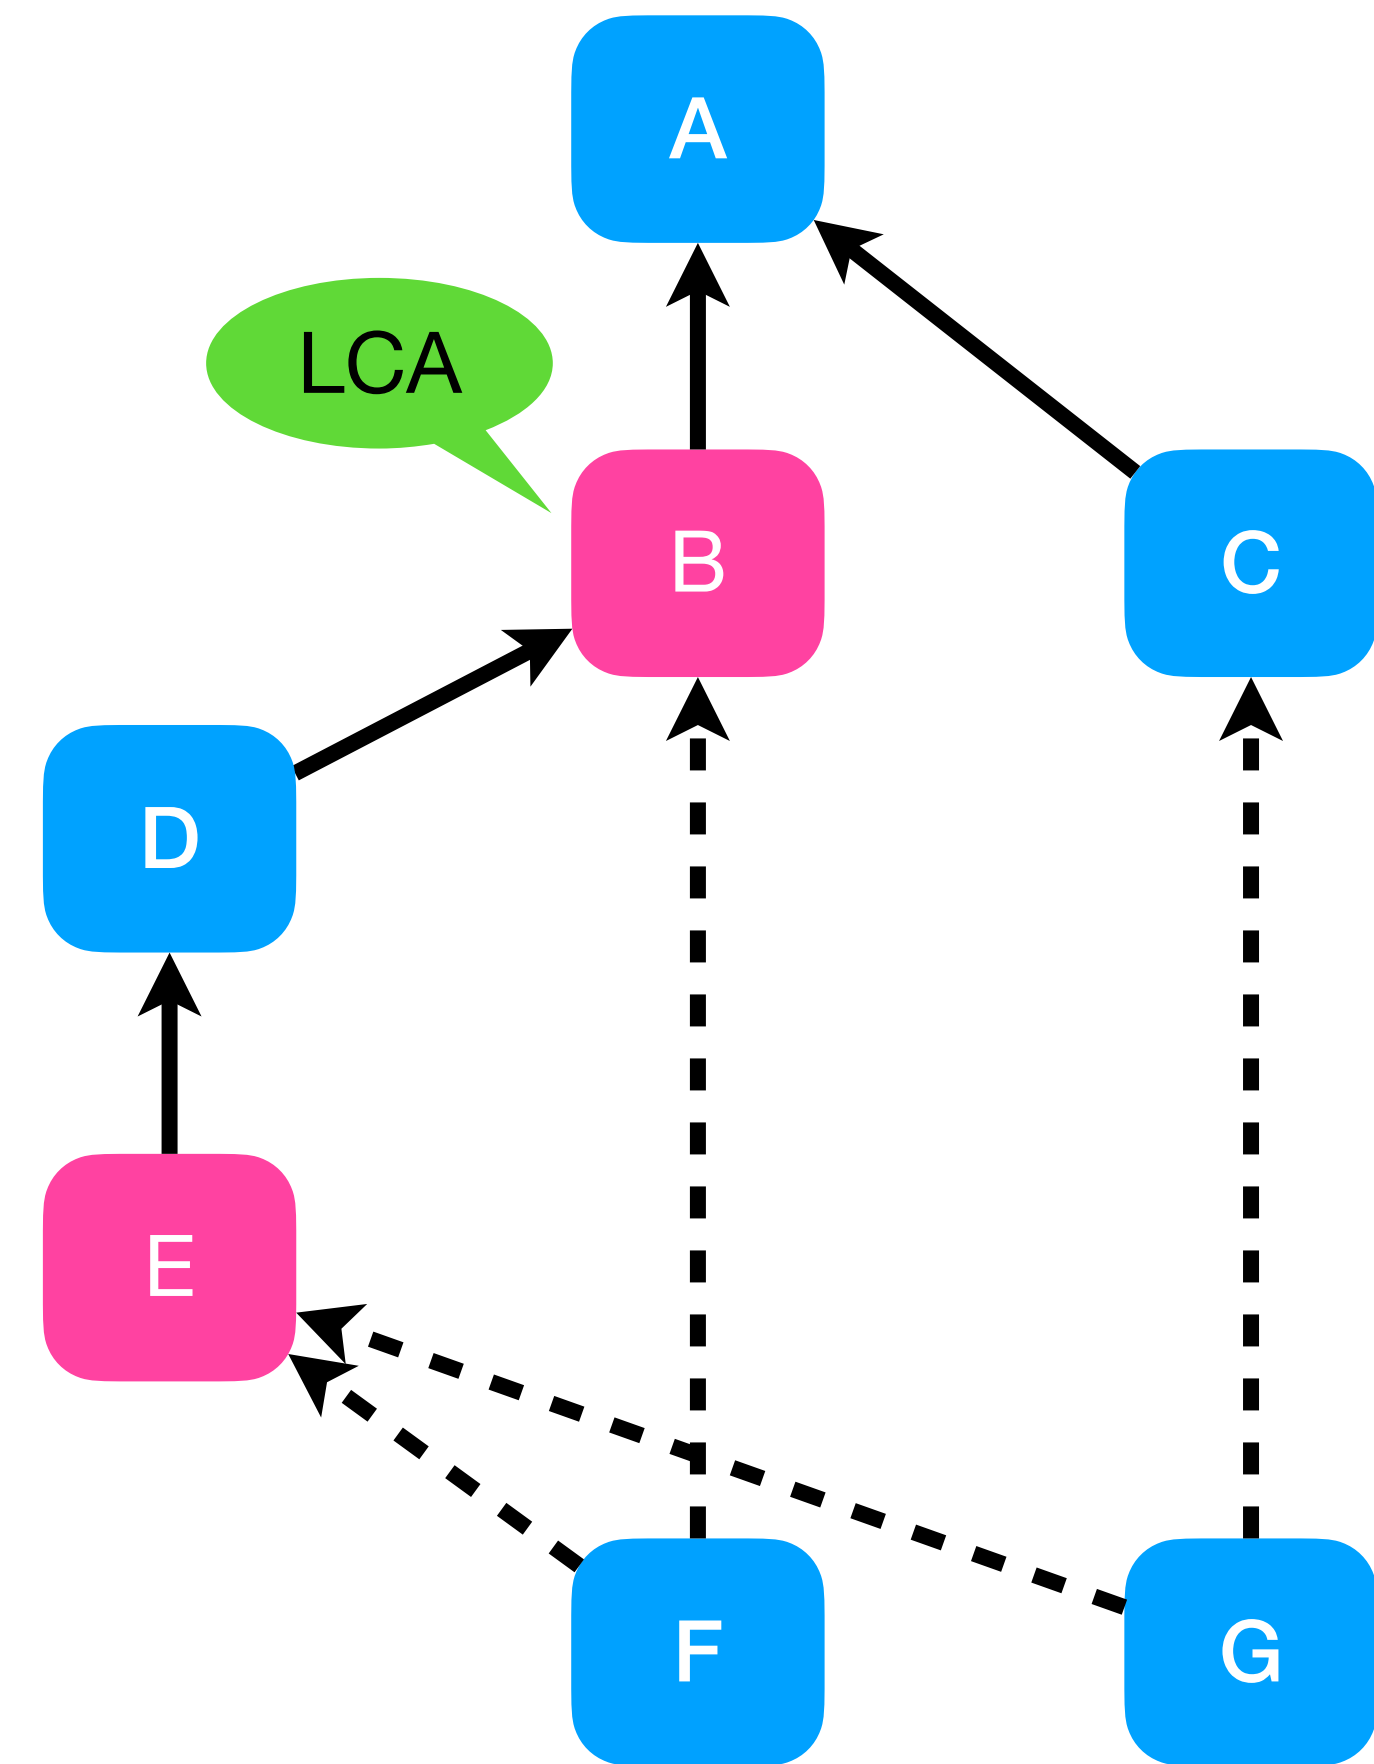
# Runtime Support

- How did we eliminate tombstones?
  - Swept under the rug — MRDT runtime.
- MRDT runtime
  - Git-like distributed store with branching and merging
  - Support 3-way merge with lowest common ancestor (LCA)
    - LCA of E & C is A



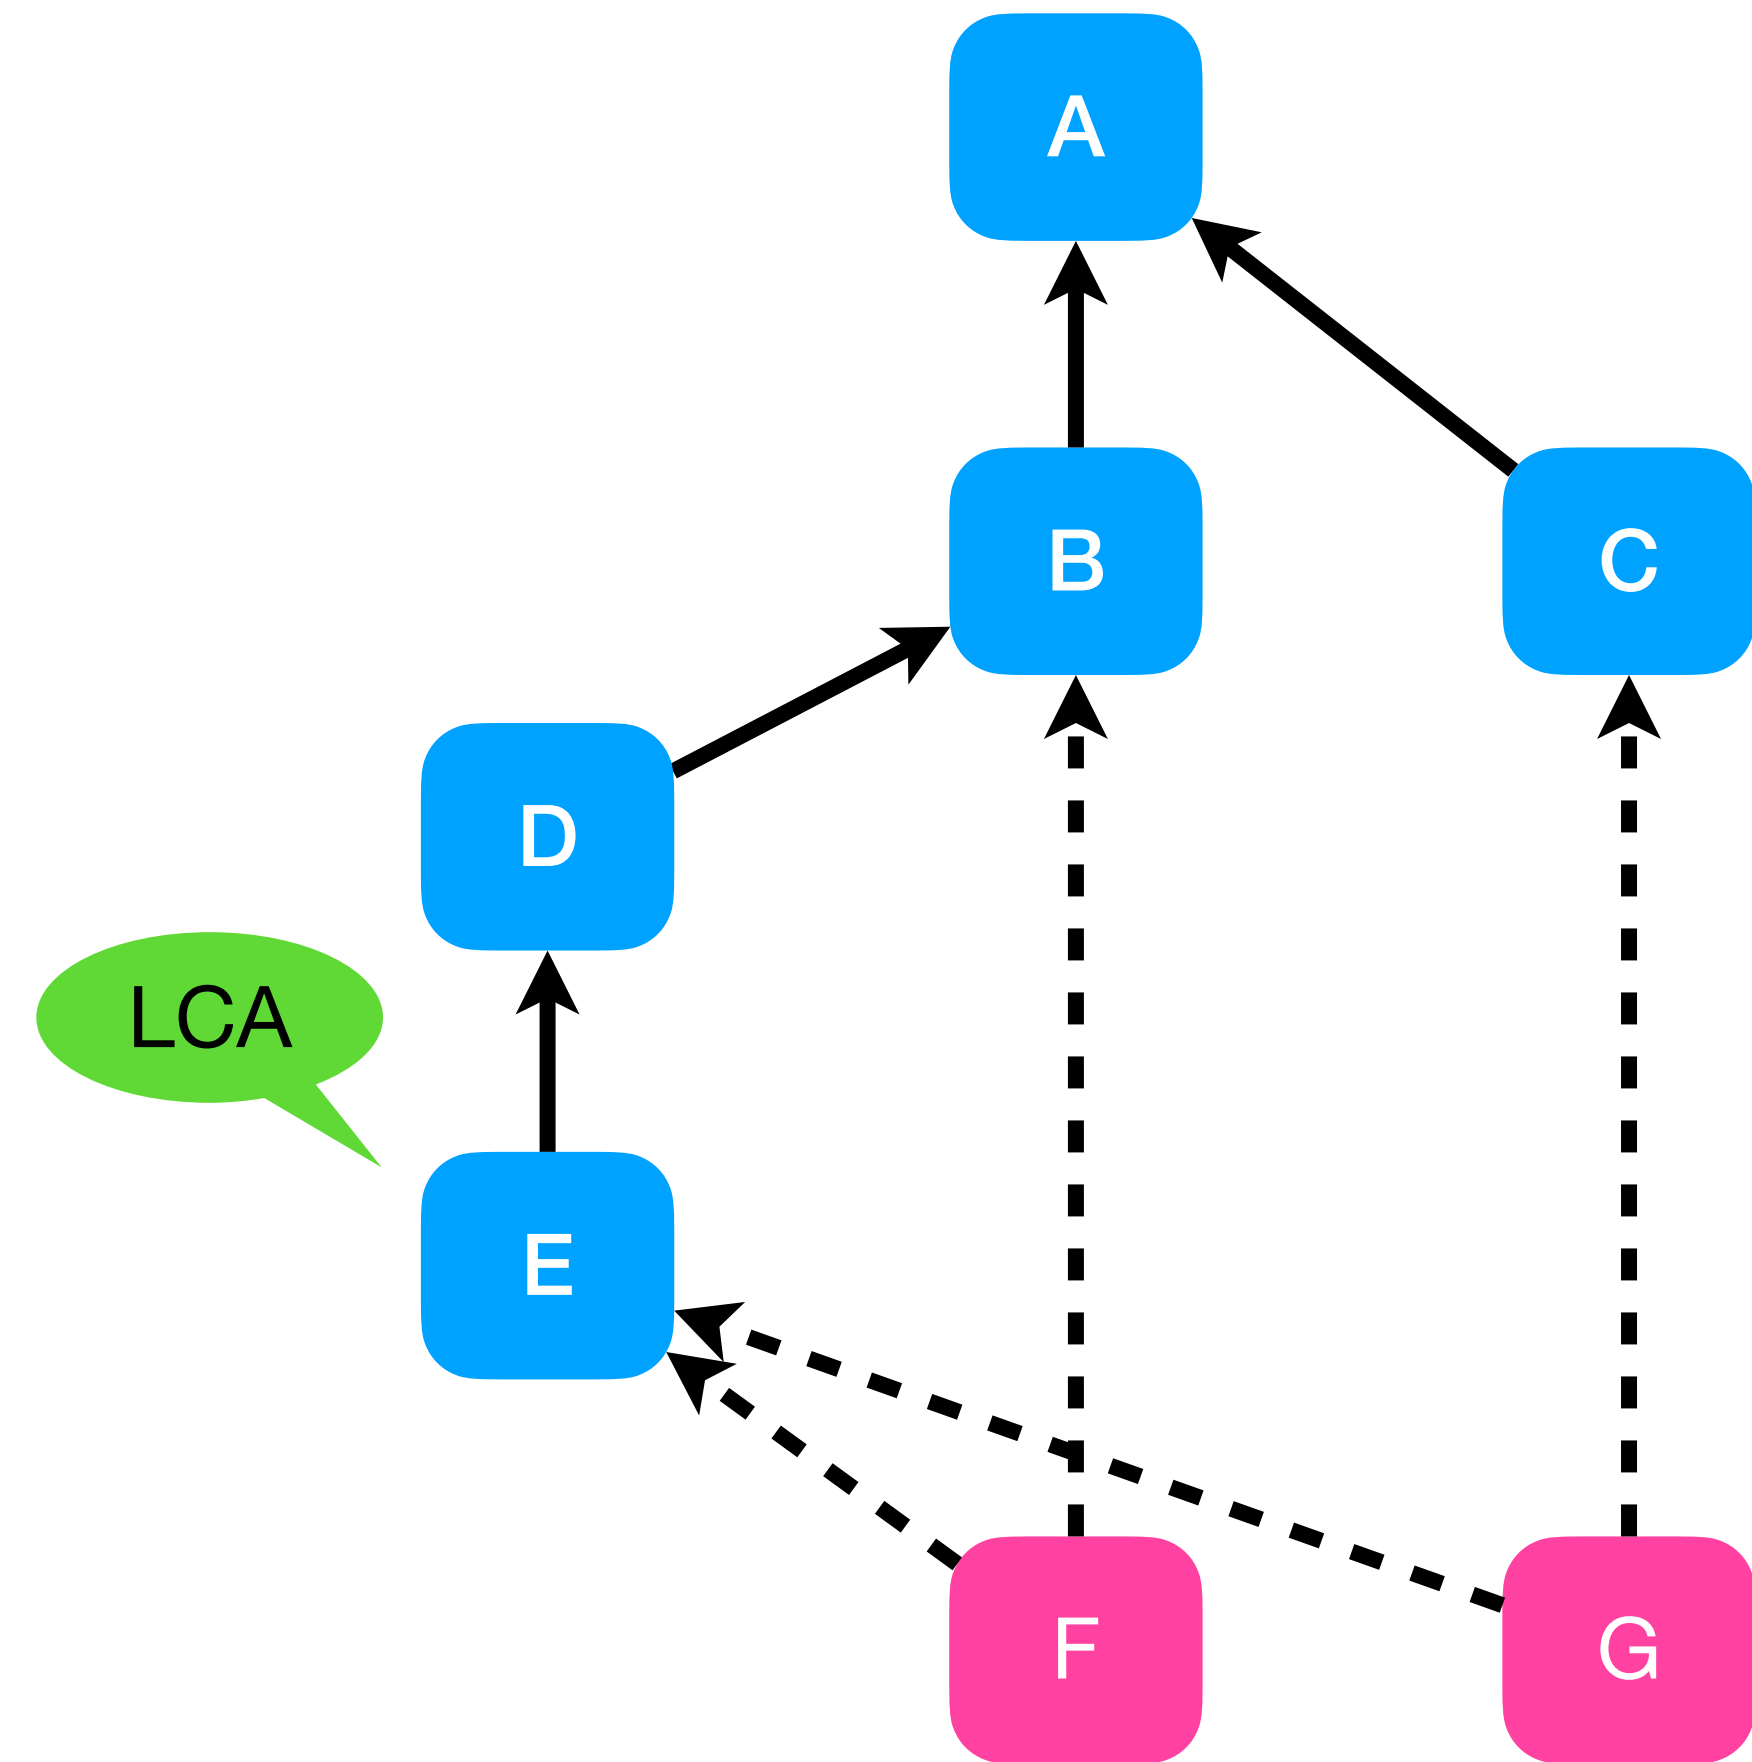
# Runtime Support

- How did we eliminate tombstones?
  - Swept under the rug — MRDT runtime.
- MRDT runtime
  - Git-like distributed store with branching and merging
  - Support 3-way merge with lowest common ancestor (LCA)
    - LCA of E & C is A
    - LCA of E & B is B



# Garbage collection of histories

- If **F** and **G** were the only two replica states
  - Only the LCA **E** is needed for any future merge



# Garbage collection of histories

- If **F** and **G** were the only two replica states
  - Only the LCA **E** is needed for any future merge
- Asynchronously garbage collect all previous commits
  - And associated states

## Banyan: Coordination-free Distributed Transactions over Mergeable Types

Shashank Shekhar Dubey<sup>1</sup>, KC Sivaramakrishnan<sup>1</sup>, Thomas Gazagnaire<sup>2</sup>, and Anil Madhavapeddy<sup>3</sup>

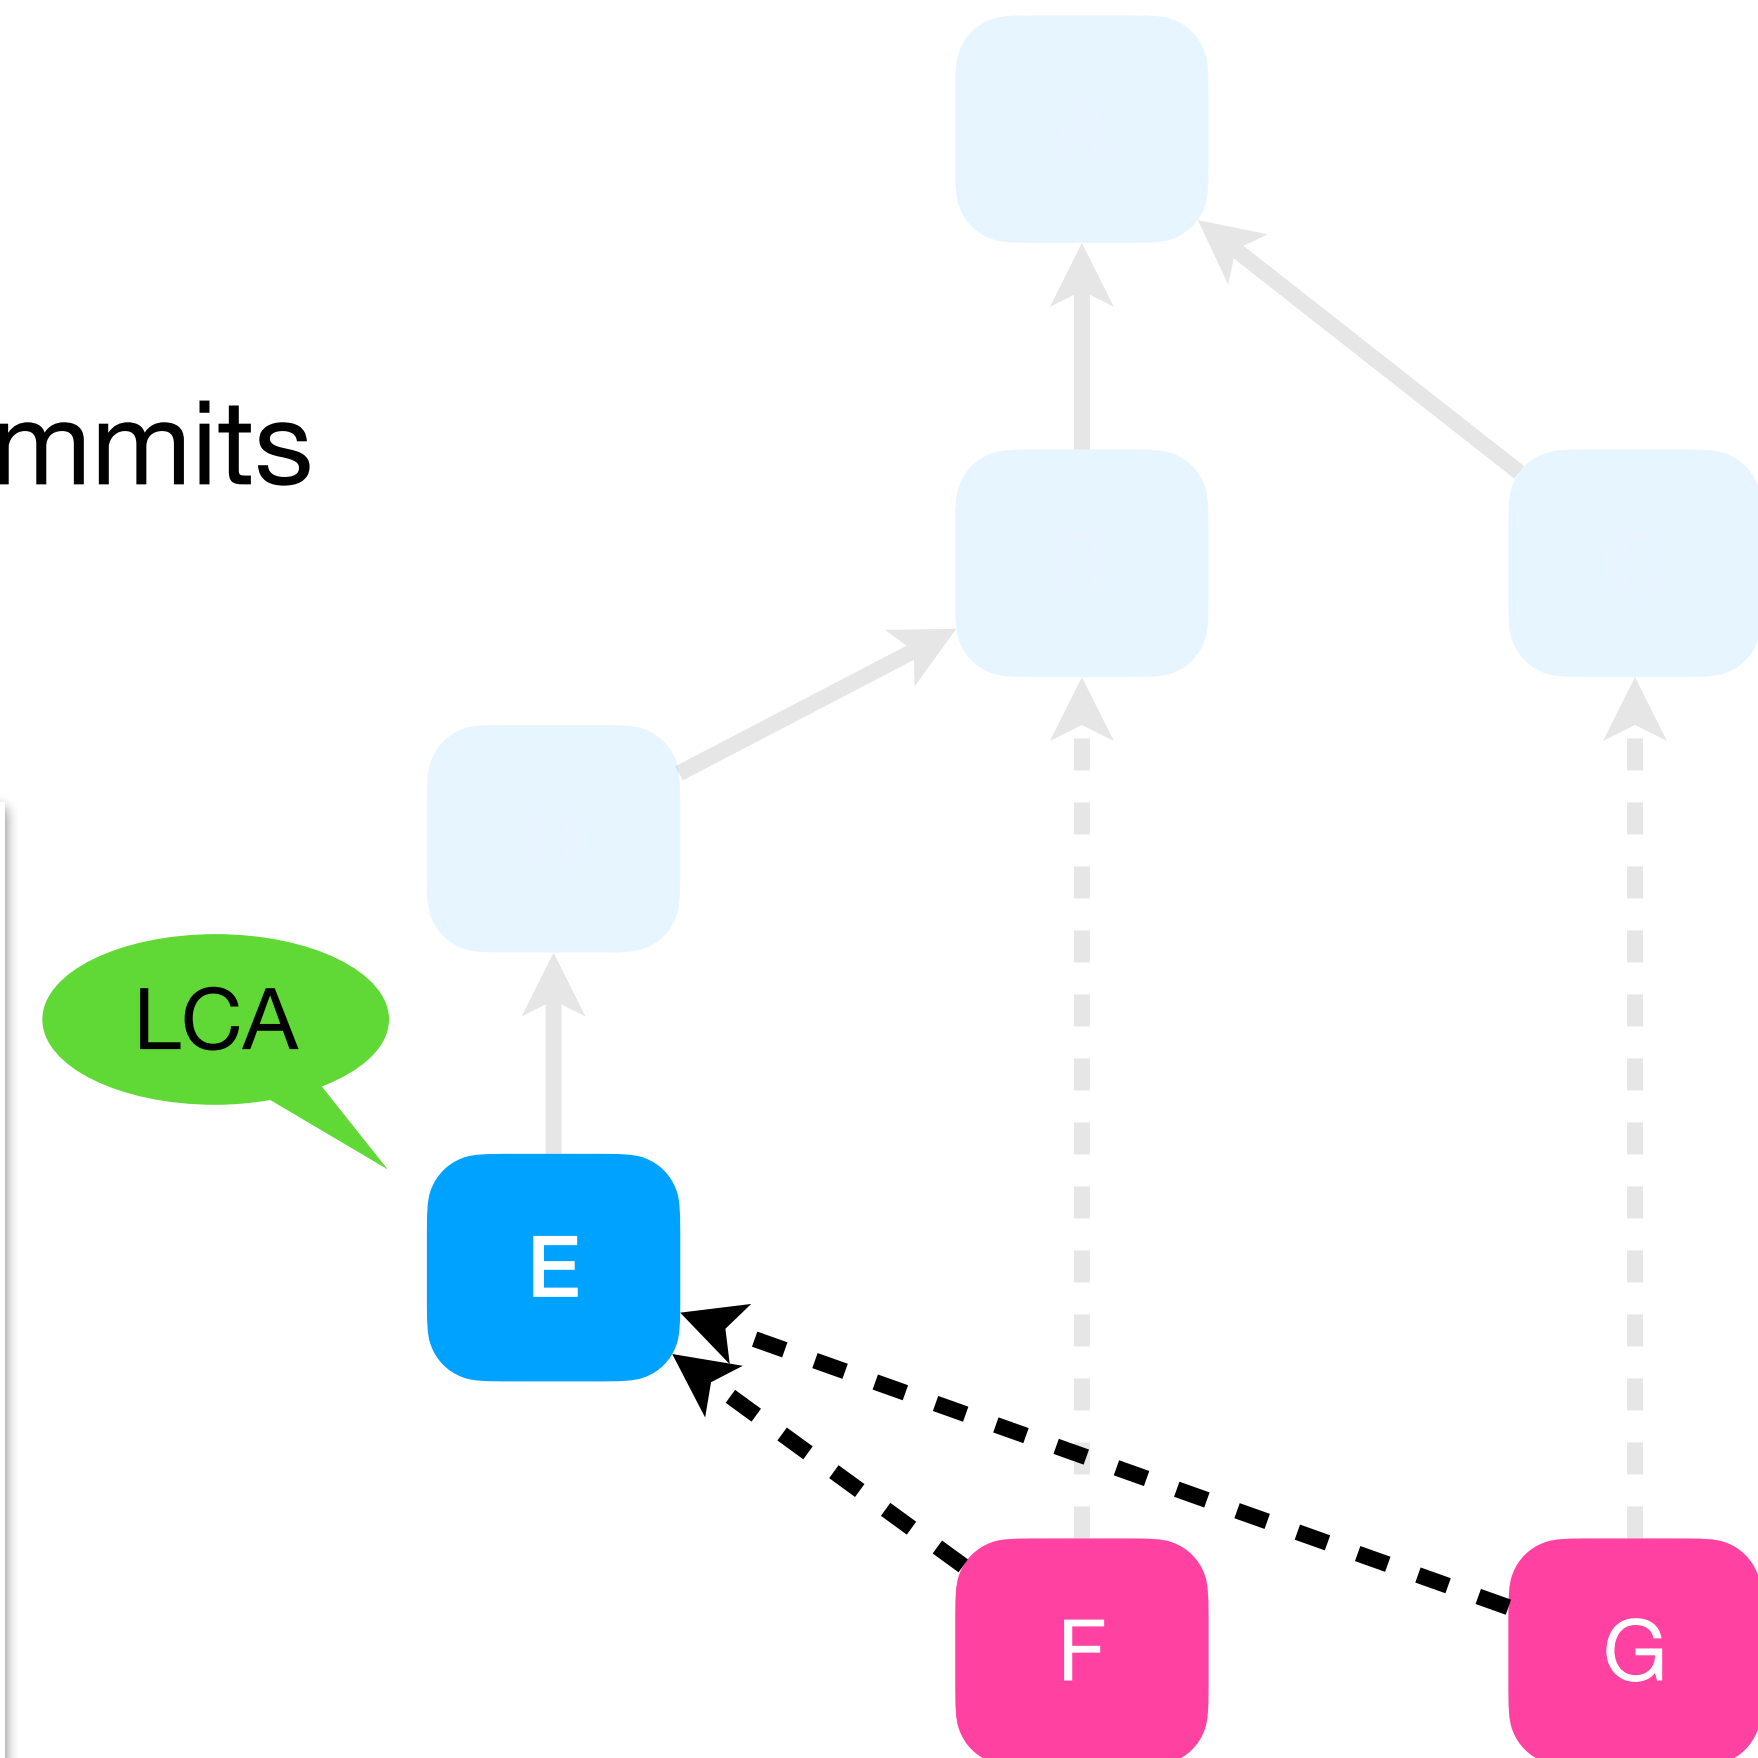
<sup>1</sup> Indian Institute of Technology, Madras, India

<sup>2</sup> Tarides, France

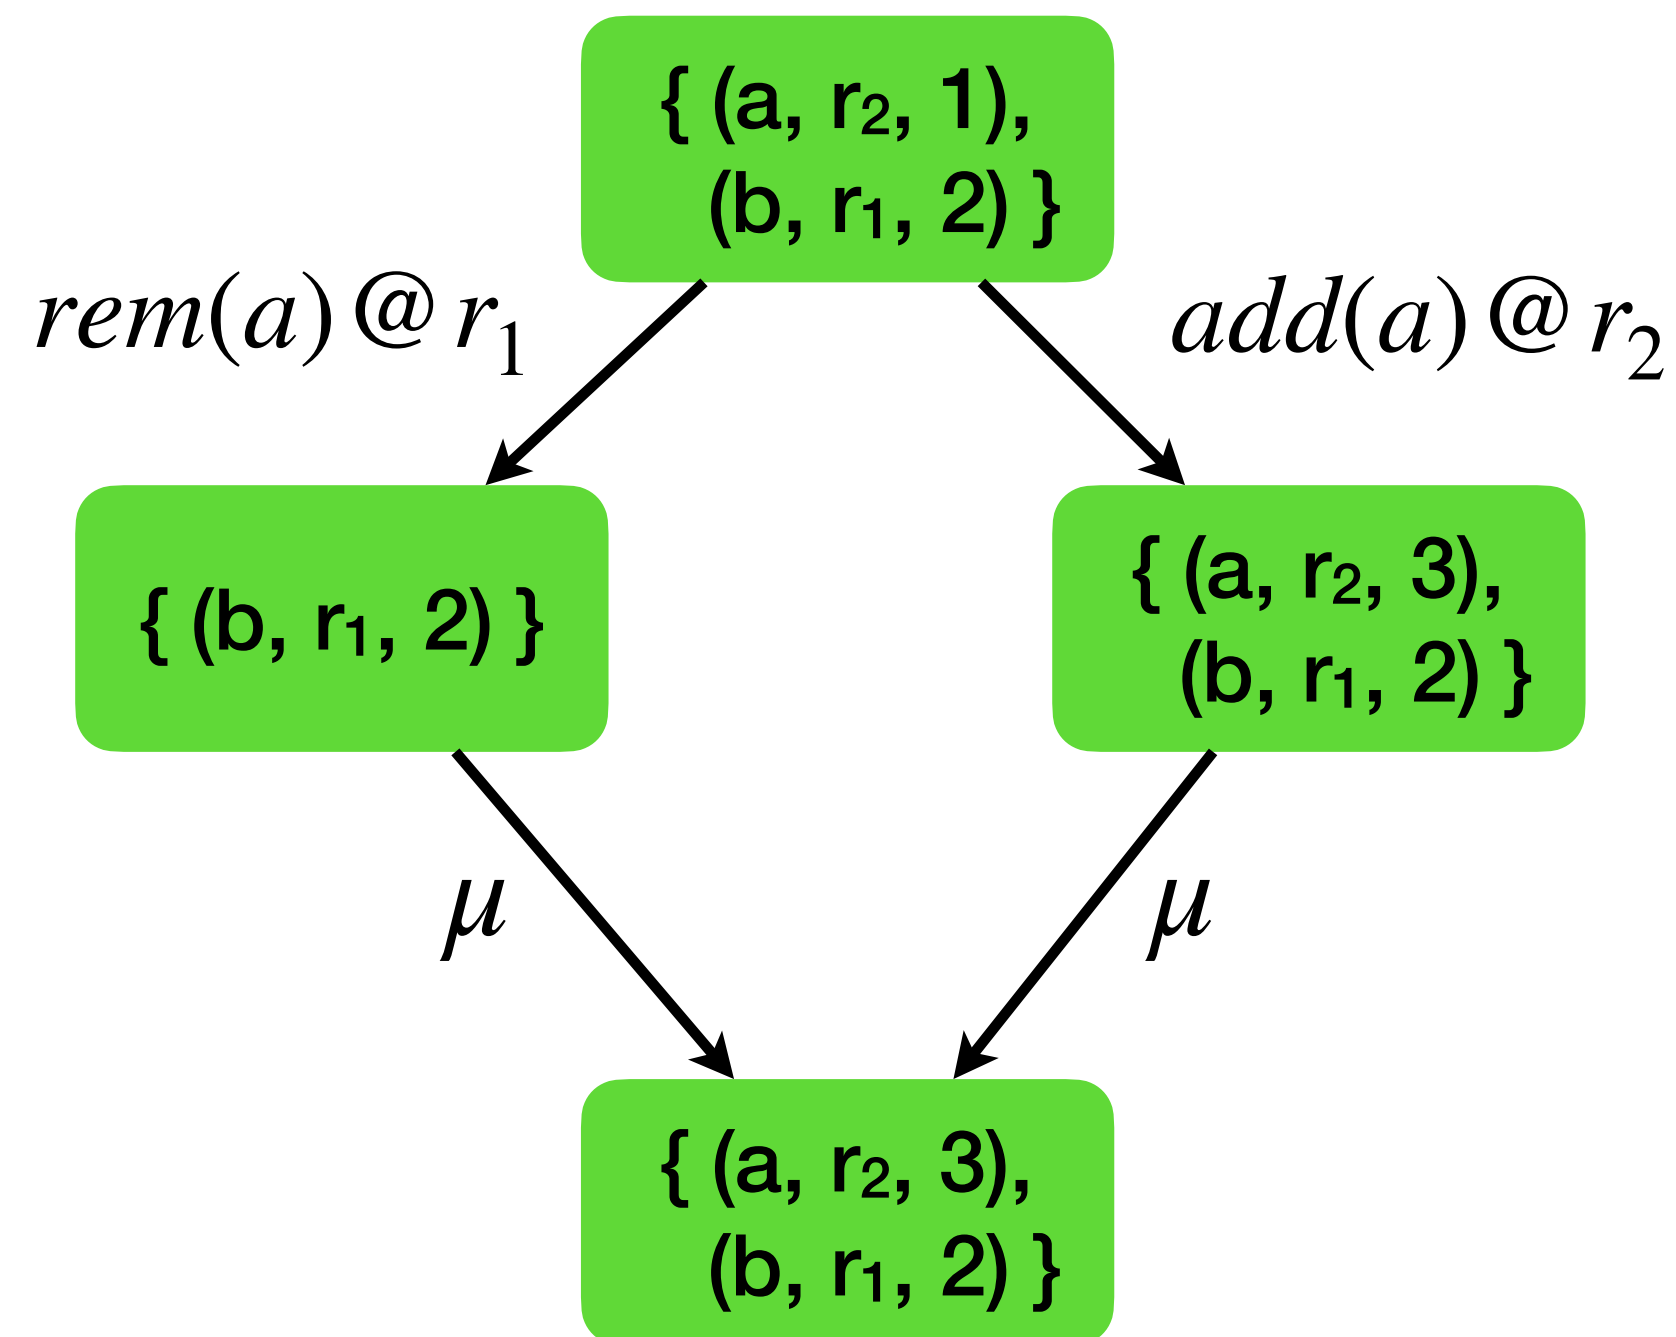
<sup>3</sup> University of Cambridge Computer Laboratory, UK

**APLAS 2020**

**Abstract.** Programming loosely connected distributed applications is a challenging endeavour. Loosely connected distributed applications such as geo-distributed stores and intermittently reachable IoT devices cannot afford to coordinate among all of the replicas in order to ensure data



# Efficient OR set MRDT



- Elements tagged with *replica id* and *timestamp*
- **add(e) @ replica r,**
  - Remove any **(e, r, ts)** for any **ts**
  - Add **(e, r, ts')** for a fresh timestamp **ts'**
- Remove, read and merge remain similar.
- Space usage  $O(|live| \cdot n)$ 
  - Duplicates — same element from different replicas

# Not the first efficient OR set

## *An Optimized Conflict-free Replicated Set*

Annette Bieniusa, INRIA & UPMC, Paris, France

Marek Zawirski, INRIA & UPMC, Paris, France

Nuno Preguiça, CITI, Universidade Nova de Lisboa, Portugal

Marc Shapiro, INRIA & LIP6, Paris, France

Carlos Baquero, HASLab, INESC TEC & Universidade do Minho, Portugal

Valter Balesgas, CITI, Universidade Nova de Lisboa, Portugal

Sérgio Duarte, CITI, Universidade Nova de Lisboa, Portugal

**N° 8083**

Octobre 2012

- Op-based CRDT
- Maintains causality info in a version vector
- Space complexity  $O(|\text{live}| \cdot n + n)$

# Not the first efficient OR set

RESEARCH-ARTICLE | FREE ACCESS

## Riak DT map: a composable, convergent replicated dictionary

Authors: [Russell Brown](#), [Sean Cribbs](#), [Christopher Meiklejohn](#), [Sam Elliott](#) | [Authors Info & Claims](#)

PaPEC '14: Proceedings of the First Workshop on Principles and Practice of Eventual Consistency • Article No.: 1, Page 1  
<https://doi.org/10.1145/2596631.2596633>

Published: 13 April 2014 [Publication History](#)  Check for updates

24 666    PDF/eReader

### Abstract

Conflict-Free Replicated Data-Types (CRDTs) [6] provide greater safety properties to eventually-consistent distributed systems without requiring synchronization. CRDTs ensure that concurrent, uncoordinated updates have deterministic outcomes via the properties of bounded join-semilattices.

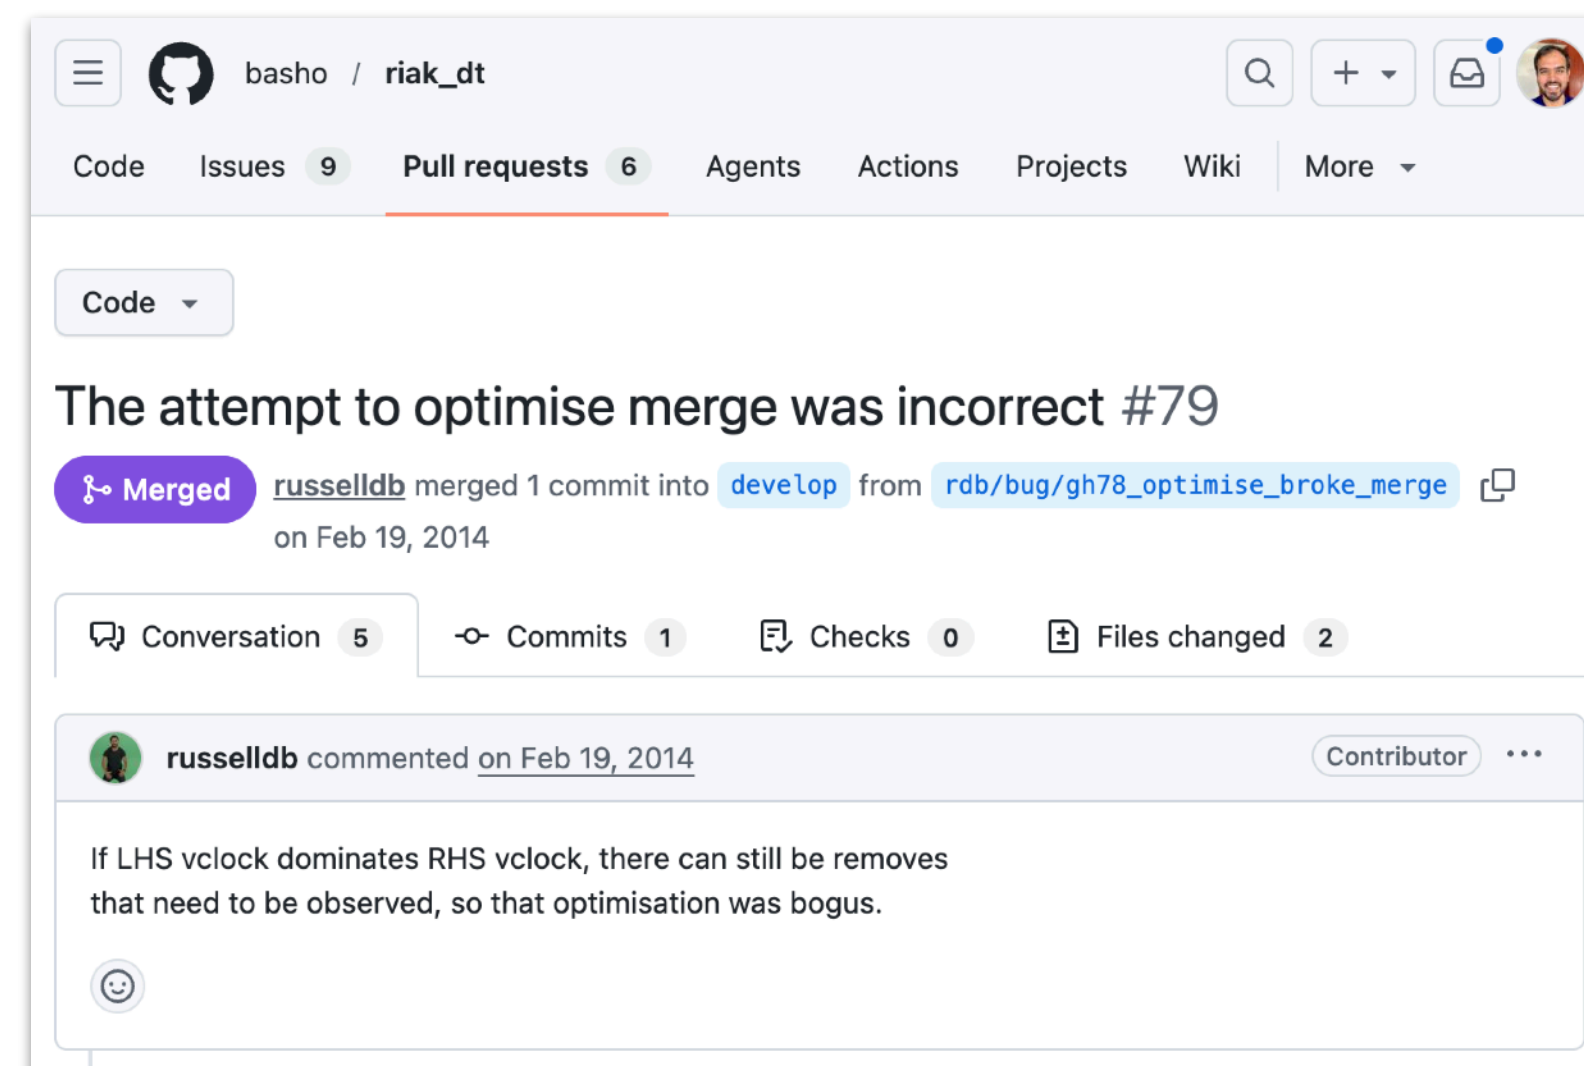
- 1st PaPoC
  - When it was PaPEC!
- State-based CRDT in Riak
- ***And buggy!***

# Buggy Riak DT Map

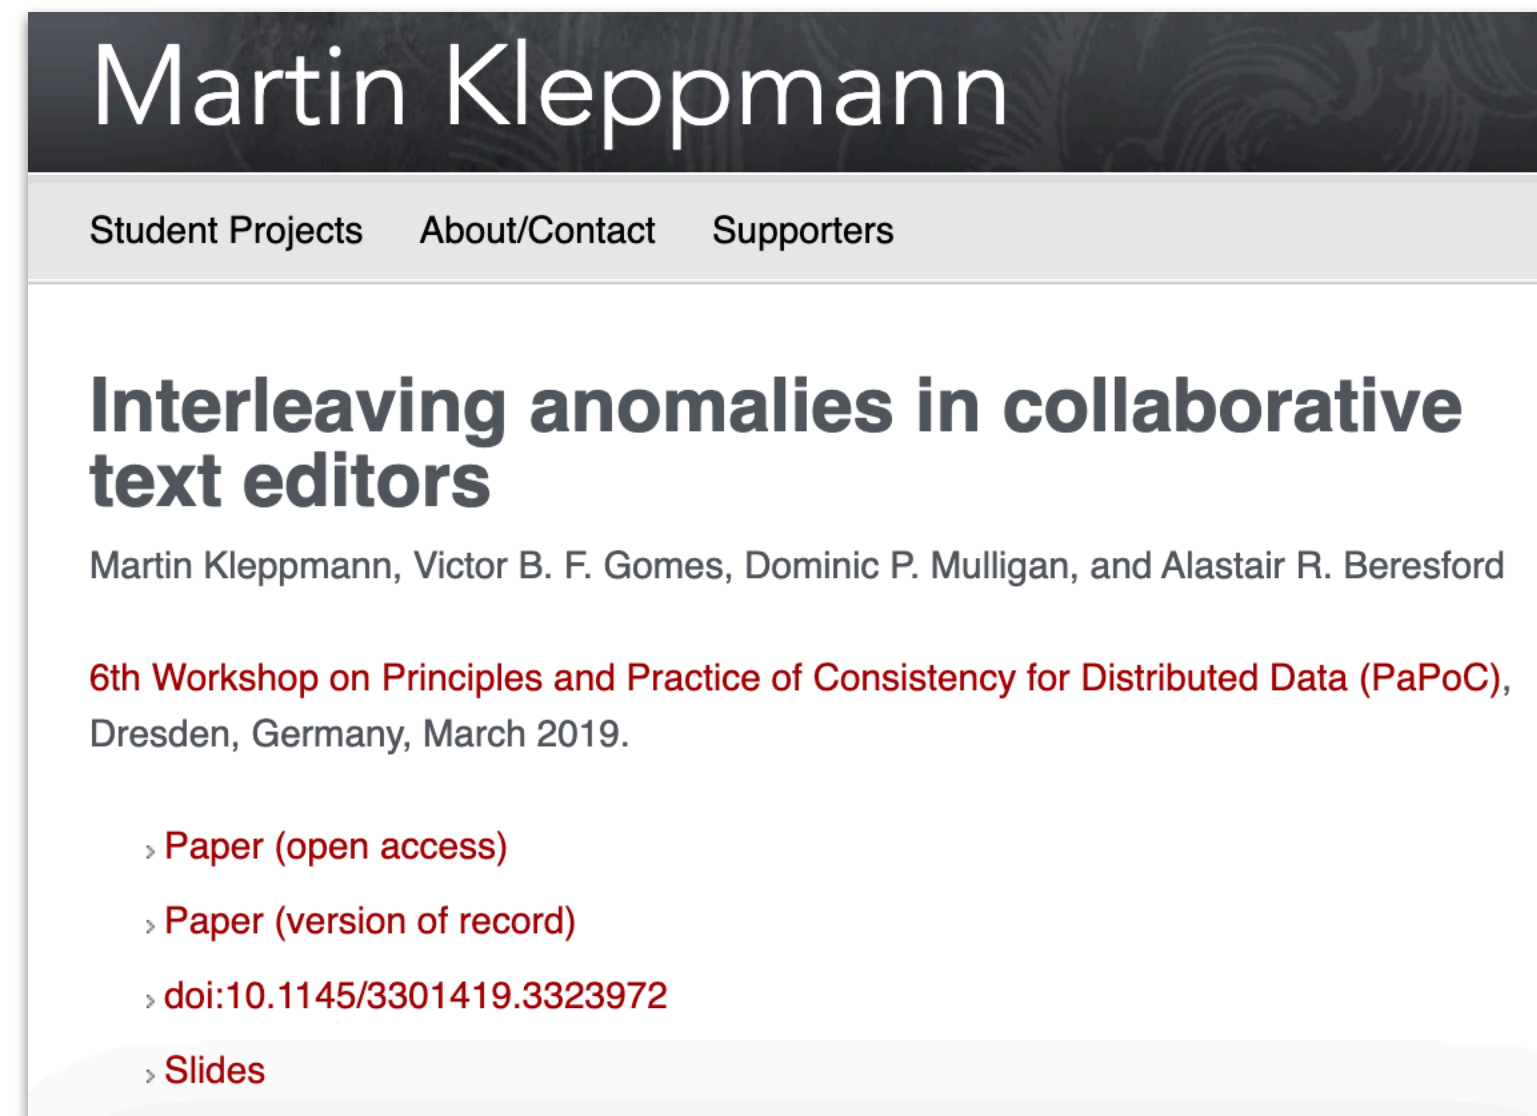
## Unsound Optimizations

Implementing CRDTs can be quite challenging. Not only does the developer who is implementing CRDTs must ensure that updates to lattices are *inflationary* – ensuring the data structure is always moving up the lattice with respect to each change – but also that the merge function must be both *deterministic* and compute the *least-upper-bound*. We have even got this wrong a few times ourselves.

<https://christophermeiklejohn.com/erlang/lasp/2019/03/08/monotonicity.html>



# Not an isolated incident



Martin Kleppmann

[Student Projects](#) [About/Contact](#) [Supporters](#)

## Interleaving anomalies in collaborative text editors

Martin Kleppmann, Victor B. F. Gomes, Dominic P. Mulligan, and Alastair R. Beresford

6th Workshop on Principles and Practice of Consistency for Distributed Data (PaPoC),  
Dresden, Germany, March 2019.

- › [Paper \(open access\)](#)
- › [Paper \(version of record\)](#)
- › [doi:10.1145/3301419.3323972](https://doi.org/10.1145/3301419.3323972)
- › [Slides](#)

6th PaPoC

## Errata

A few years after publication, we learned about some serious flaws in this paper:

- › The definition of non-interleaving in Section 2.1 cannot be guaranteed by any algorithm.
- › The algorithm presented in Section 3.1 does not guarantee convergence.





Please see the [Fugue paper](#) for more details.

# Verifying RDTs




- The goal is to *verify* the correctness of RDTs
- What's common between **designing RDTs** and **pen-and-paper proofs**?
  - I make mistakes in both!
- Machine-checked proofs
  - Use SMT-solvers, proof-oriented PLs, proof-assistants

RESEARCH-ARTICLE | OPEN ACCESS Peepul





## Certified mergeable replicated data types

Authors:  Vimala Soundarapandian,  Adharsh Kamath,  Kartik Nagar,  KC Sivaramakrishnan | [Authors Info & Claims](#)


PLDI 2022: Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation  
Pages 332 - 347 • <https://doi.org/10.1145/3519939.3523735>

RESEARCH-ARTICLE | OPEN ACCESS |   

## Automatically Verifying Replication-Aware Linearizability

Authors:  Vimala Soundarapandian,  Kartik Nagar,  Aseem Rastogi,  KC Sivaramakrishnan | [Authors Info & Claims](#)

Proceedings of the ACM on Programming Languages, Volume 9, Issue OOPSLA1  
Article No.: 111, Pages 871 - 897 • <https://doi.org/10.1145/3720452> Neem

Published: 09 April 2025 [Publication History](#) 

## SAL: Multi-modal Verification of Replicated Data Types

|                                                                                                                      |                                                                                                                               |                                                                                                                       |
|----------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| <p>Pranav Ramesh<br/>Indian Institute of Technology,<br/>Madras<br/>Chennai, India<br/>cs22b015@smail.iitm.ac.in</p> | <p>Vimala Soundarapandian<br/>Indian Institute of Technology,<br/>Madras<br/>Chennai, India<br/>cs19d750@smail.iitm.ac.in</p> | <p>KC Sivaramakrishnan<br/>Indian Institute of Technology,<br/>Madras<br/>Chennai, India<br/>kcsrk@cse.iitm.ac.in</p> |
|----------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|

**Abstract**  
Designing correct replicated data types (RDTs) is challenging because replicas evolve independently and must be merged while preserving application intent. A promising approach is correct-by-construction development in a proof-oriented programming language such as F\*, Dafny and Lean, where desired correctness guarantees are specified and checked on

and Mergeable Replicated Data Types (MRDTs) [9] are widely used to implement such replicated state. However, designing correct RDTs is subtle. Even well-known designs such as Replicated Growable Arrays (RGA) [10] have had serious anomalies discovered after publication<sup>1</sup>.  
Replication-aware (RA) linearizability [21] provides a principled correctness condition for CRDTs, relating replica exe-

# Convergence alone is insufficient

## Merge function $\mu$

$$\mu(a, b) = \mu(b, a)$$

$$\mu(a, a) = a$$

$$\mu(\mu(a, b), c) = \mu(a, \mu(b, c))$$

Commutativity


Idempotence

Associativity


Satisfies algebraic properties

```
let merge v1 v2 = max v1 v2
```





Intent is not captured



# Peepul: Capturing Intent through Axiomatic Spec

RESEARCH-ARTICLE | OPEN ACCESS X in 

## Certified mergeable replicated data types

Authors:  [Vimala Soundarapandian](#),  [Adharsh Kamath](#),  [Kartik Nagar](#),  [KC Sivaramakrishnan](#) | [Authors Info & Claims](#)

PLDI 2022: Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation  
Pages 332 - 347 • <https://doi.org/10.1145/3519939.3523735>

- Abstract state = events + visibility (partial order)
- Operations = folds over abstract state

*Specification*



MRDT operations +  
3-way merge function

*Implementation*



A proof-oriented SMT-aided programming language

# Abstract state of a data type

An abstract state for a data type  $\tau = (Op_\tau, Val_\tau)$  is a tuple

$$I = \langle E, oper, rval, time, vis \rangle,$$

where

$$E \subseteq Event$$

$$oper : E \rightarrow Op_\tau$$

$$rval : E \rightarrow Val_\tau$$

$$time : E \rightarrow Timestamp$$

$$vis \subseteq E \times E$$

(set of events),


(operation of each event),

(return value of each event),

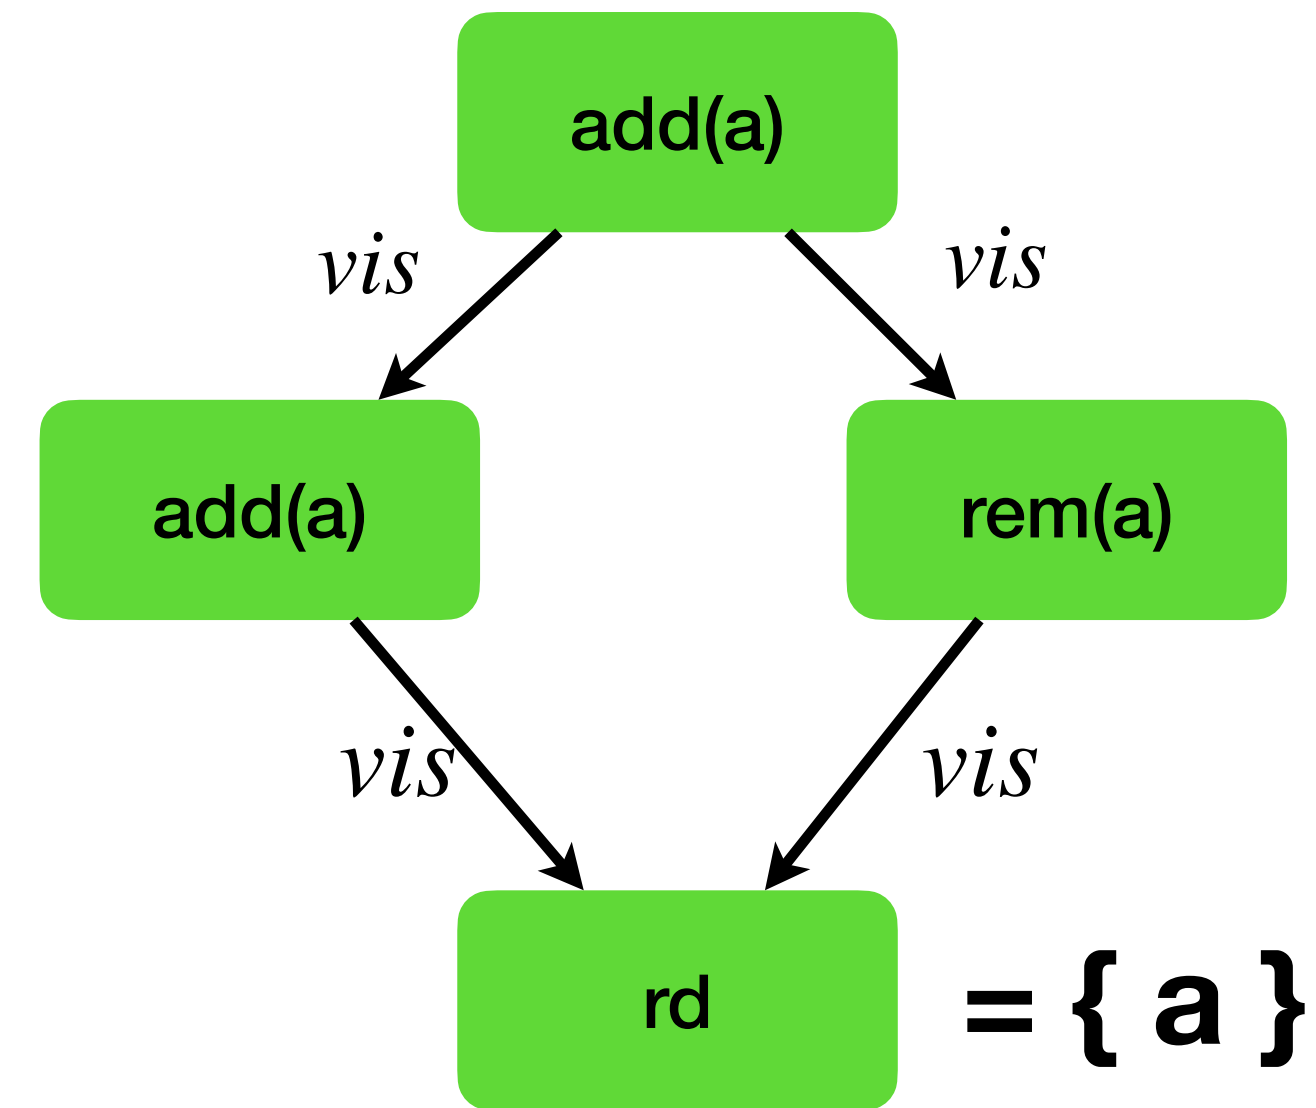
(timestamp of each event),

(visibility relation).

# Specifying OR set behaviour



**Concrete state**



**Abstract state**

$$\mathcal{F}_{orset}(rd, \langle E, oper, rval, time, vis \rangle) = \{ a \mid \exists e \in E. oper(e) = add(a) \wedge \neg(\exists f \in E. oper(f) = remove(a) \wedge e \xrightarrow{vis} f) \}$$

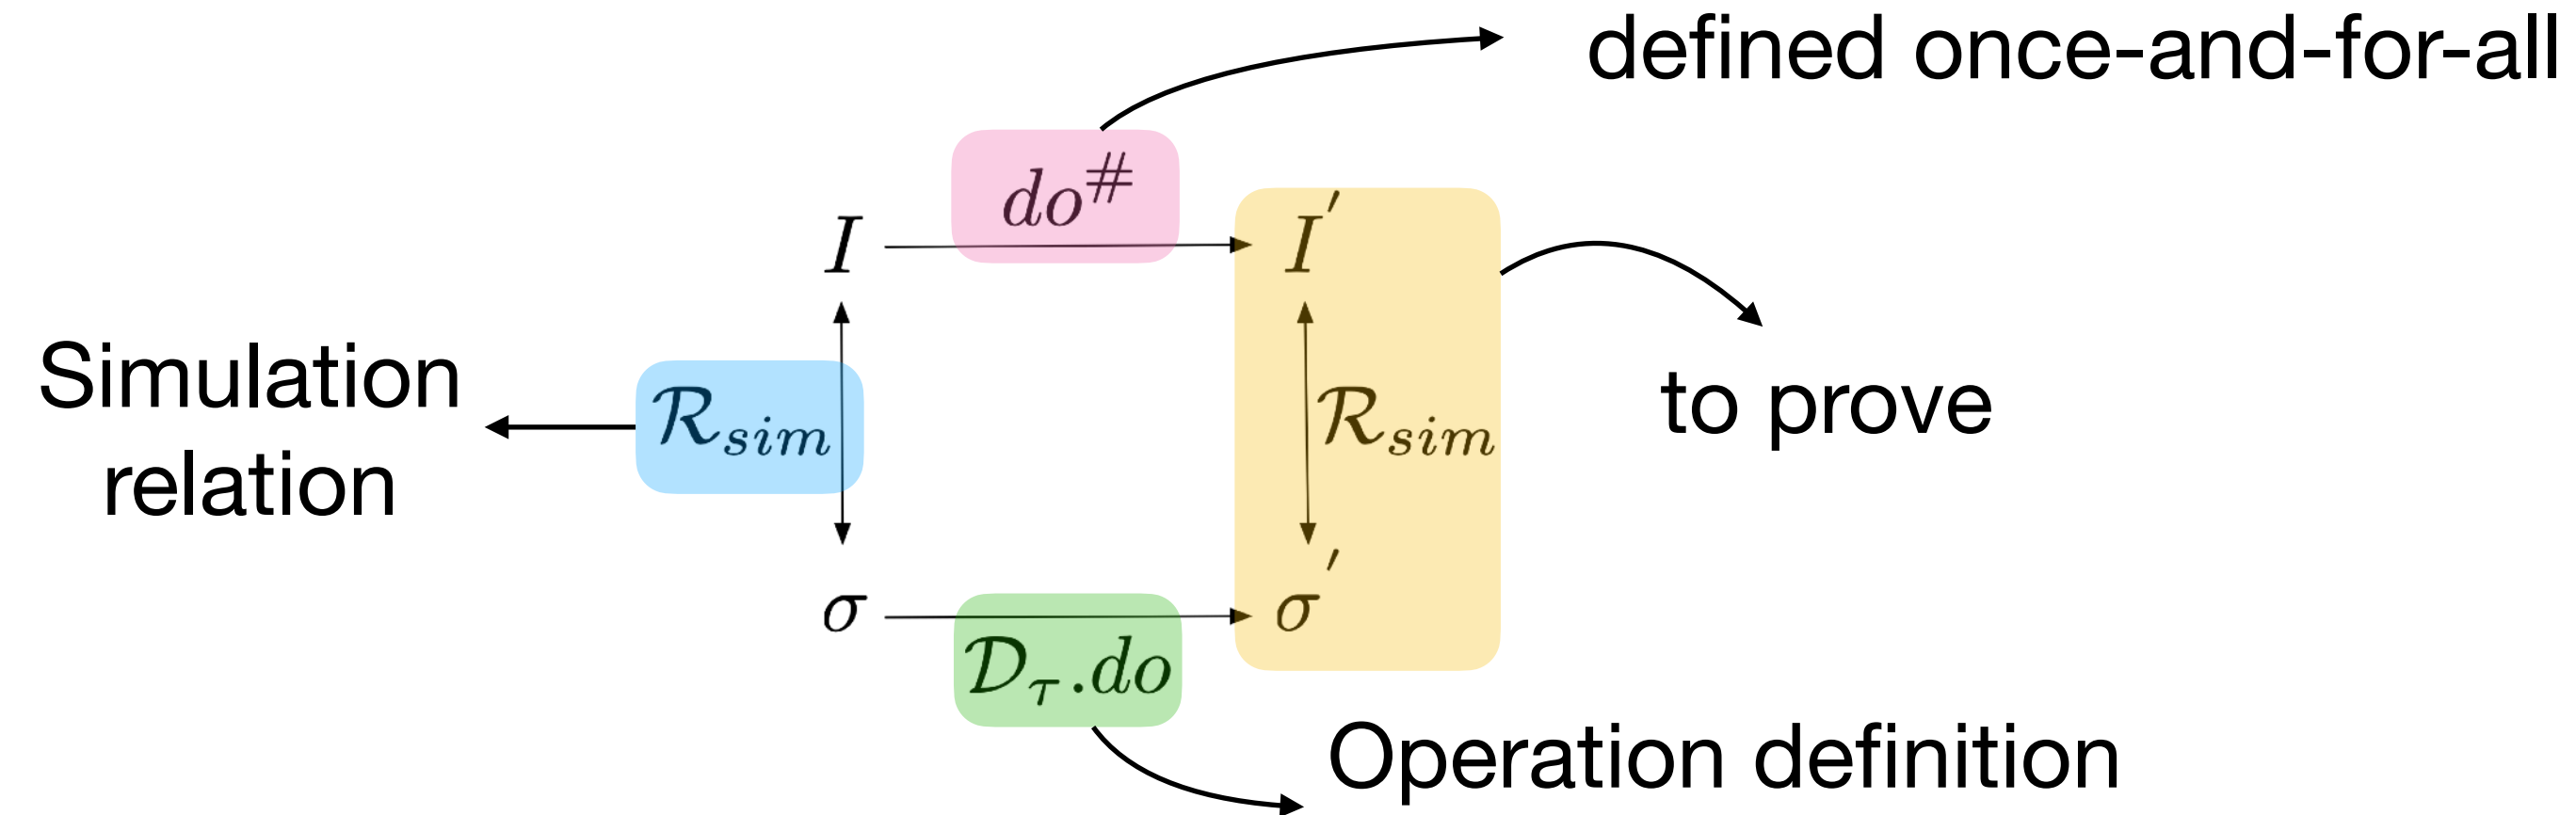
# Simulation Relation for OR set

- Connect the abstract and the concrete state

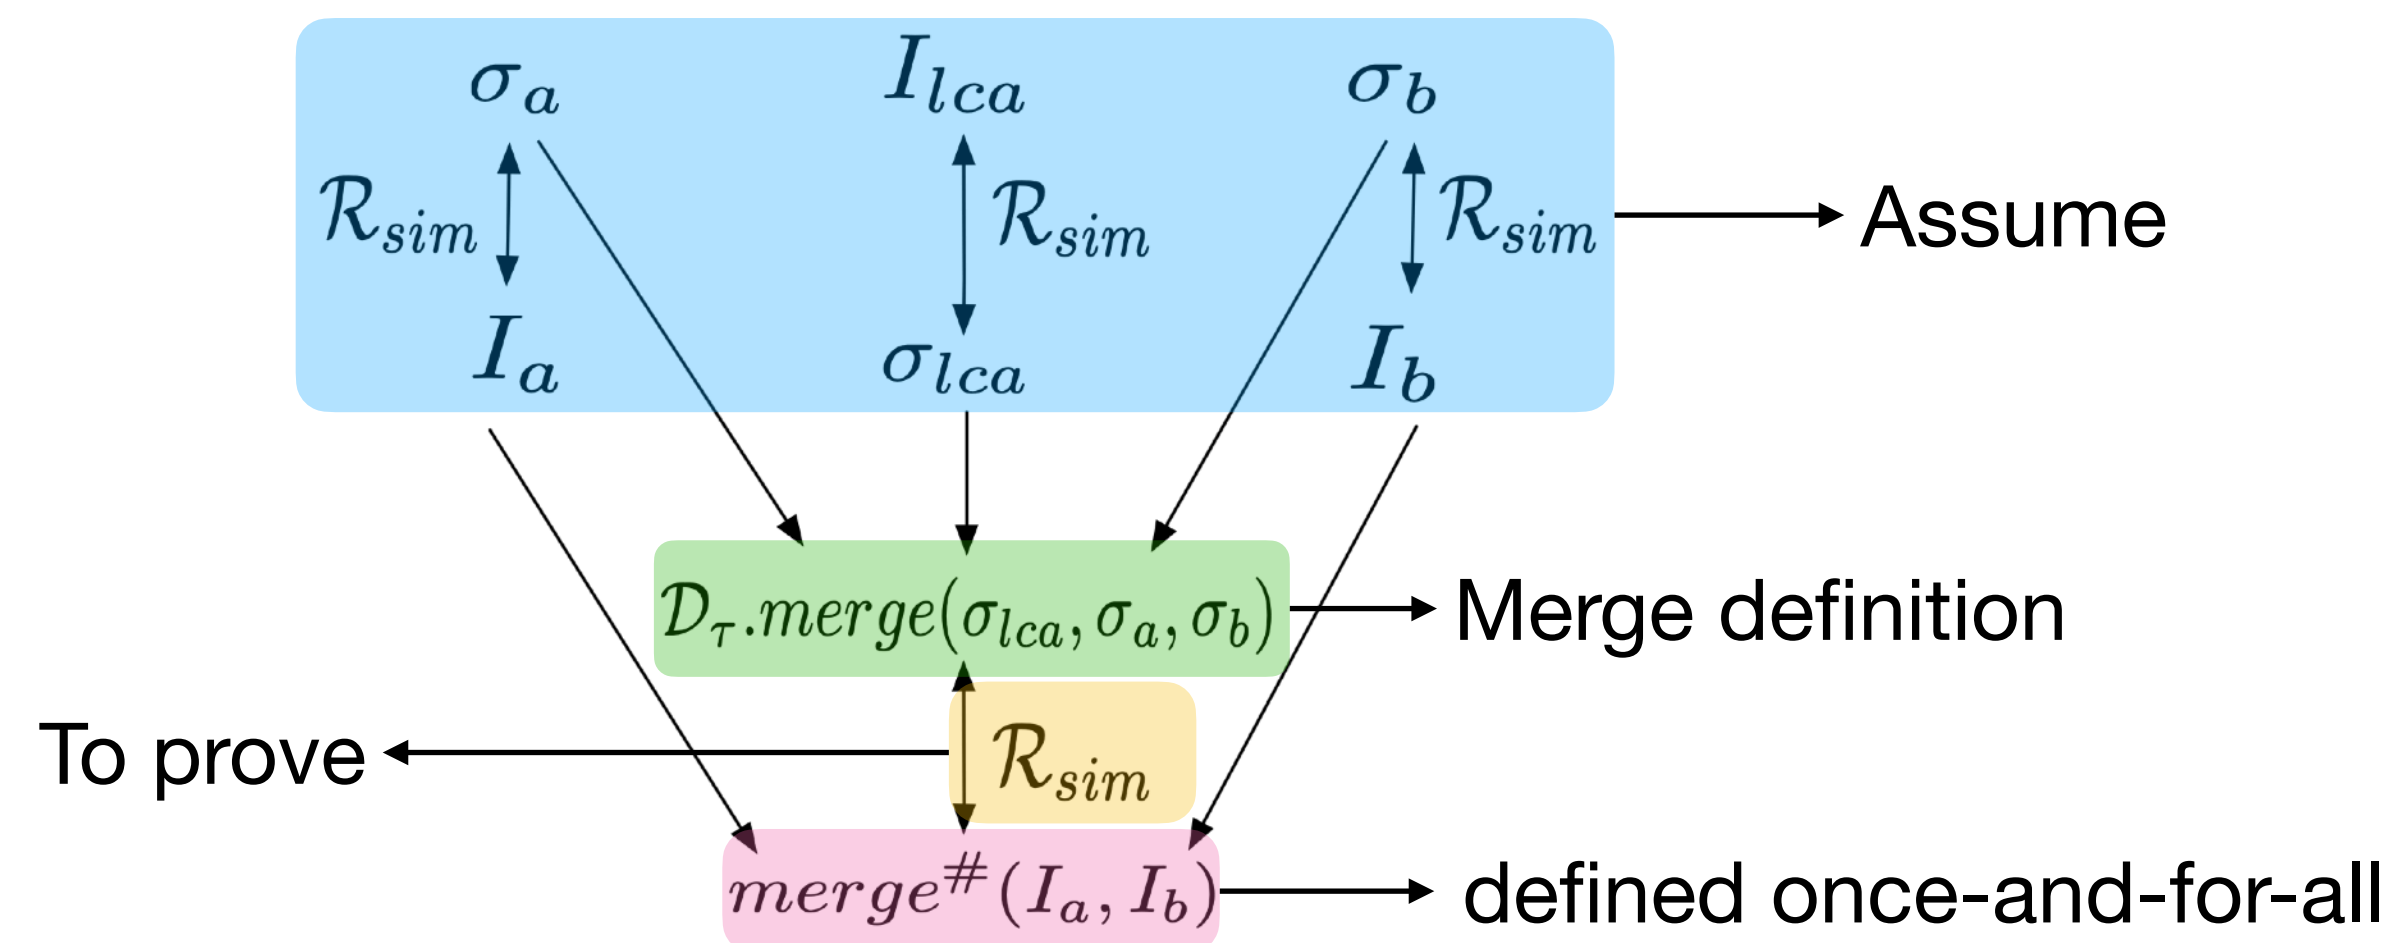
$$\begin{aligned} \mathcal{R}_{sim}(I, \sigma) \iff & (\forall (a, t) \in \sigma \iff \\ & (\exists e \in I.E \wedge I.oper(e) = add(a) \wedge I.time(e) = t \wedge \\ & \neg(\exists f \in I.E \wedge I.oper(f) = remove(a) \wedge e \xrightarrow{vis} f))) \end{aligned}$$

# Verification

- Prove that the simulation holds for operations



- Prove that the simulation holds for merge




# Verification (not done yet)

- Prove that the specification and the implementation agree on the return values of operations

$$\forall I, \sigma, e, op, a, t. \mathcal{R}_{sim}(I, \sigma) \wedge \mathcal{D}_\tau.do(op, \sigma, t) = (\sigma', a) \\ \wedge \Psi_{ts}(I) \implies a = \mathcal{F}_\tau(op, I)$$

- Prove convergence

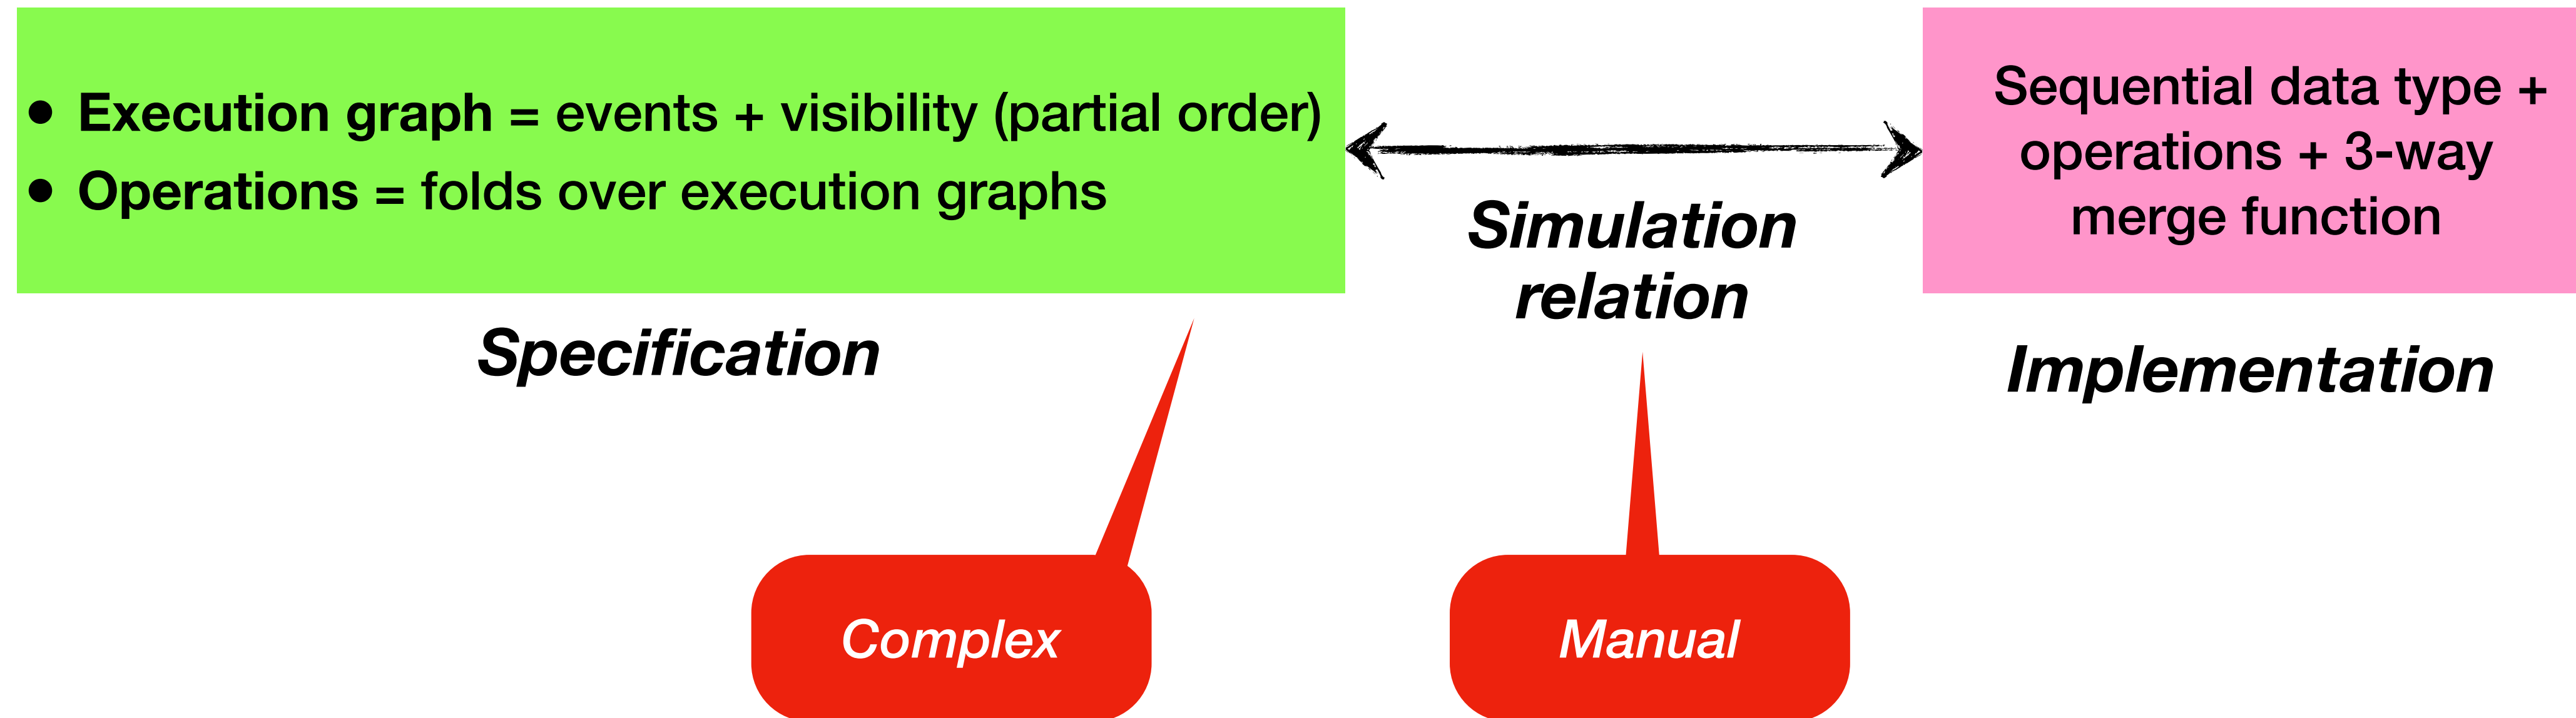
$$\forall I, \sigma_a, \sigma_b. \mathcal{R}_{sim}(I, \sigma_a) \wedge \mathcal{R}_{sim}(I, \sigma_b) \implies \sigma_a \sim \sigma_b$$



# Verified MRDTs using Peepul


| MRDTs verified         | #Lines code | #Lines proof | #Lemmas | Verif. time (s) |
|------------------------|-------------|--------------|---------|-----------------|
| Increment-only counter | 6           | 43           | 2       | 3.494           |
| PN counter             | 8           | 43           | 2       | 23.211          |
| Enable-wins flag       | 20          | 58           | 3       | 1074            |
|                        |             | 81           | 6       | 171             |
|                        |             | 89           | 7       | 104             |
| LWW register           | 5           | 44           | 1       | 4.21            |
| G-set                  | 10          | 23           | 0       | 4.71            |
|                        |             | 28           | 1       | 2.462           |
|                        |             | 33           | 2       | 1.993           |
| G-map                  | 48          | 26           | 0       | 26.089          |
| Mergeable log          | 39          | 95           | 2       | 36.562          |
| OR-set (§2.1.1)        | 30          | 36           | 0       | 43.85           |
|                        |             | 41           | 1       | 21.656          |
|                        |             | 46           | 2       | 8.829           |
| OR-set-space (§2.1.2)  | 59          | 108          | 7       | 1716            |
| OR-set-spacetime       | 97          | 266          | 7       | 1854            |
| Queue                  | 32          | 1123         | 75      | 4753            |

# Challenges with Peepul



- Queue specification more complicated than the implementation
  - Recovering a queue structure from a partially-ordered soup of events is HARD!
- Proofs are manual
  - F\* isn't great for manual proofs

# *Is there a more natural spec?*



$$\sigma_5 = \sigma_6 = \text{linearization}(\{op_1, op_2, op_3, op_4\}) \sigma_0$$

# Replication-aware Linearizability

RESEARCH-ARTICLE

## Replication-aware linearizability

**Authors:**  [Chao Wang](#),  [Constantin Enea](#),  [Suha Orhun Mutluergil](#),  [Gustavo Petri](#) | [Authors Info & Clai](#)

[PLDI 2019: Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation](#)  
<https://doi.org/10.1145/3314221.3314617>

- Replica states should be a *linearisation* of observed *update* operations
  - Linearisation total order *lo* compatible with partially-ordered visibility relation *vis*
  - No real-time ordering requirement unlike traditional linearizability
- Payoff
  - If a replicated object is RA-linearizable, reason about it using *sequential semantics*

# Using RA-linearizability for verification

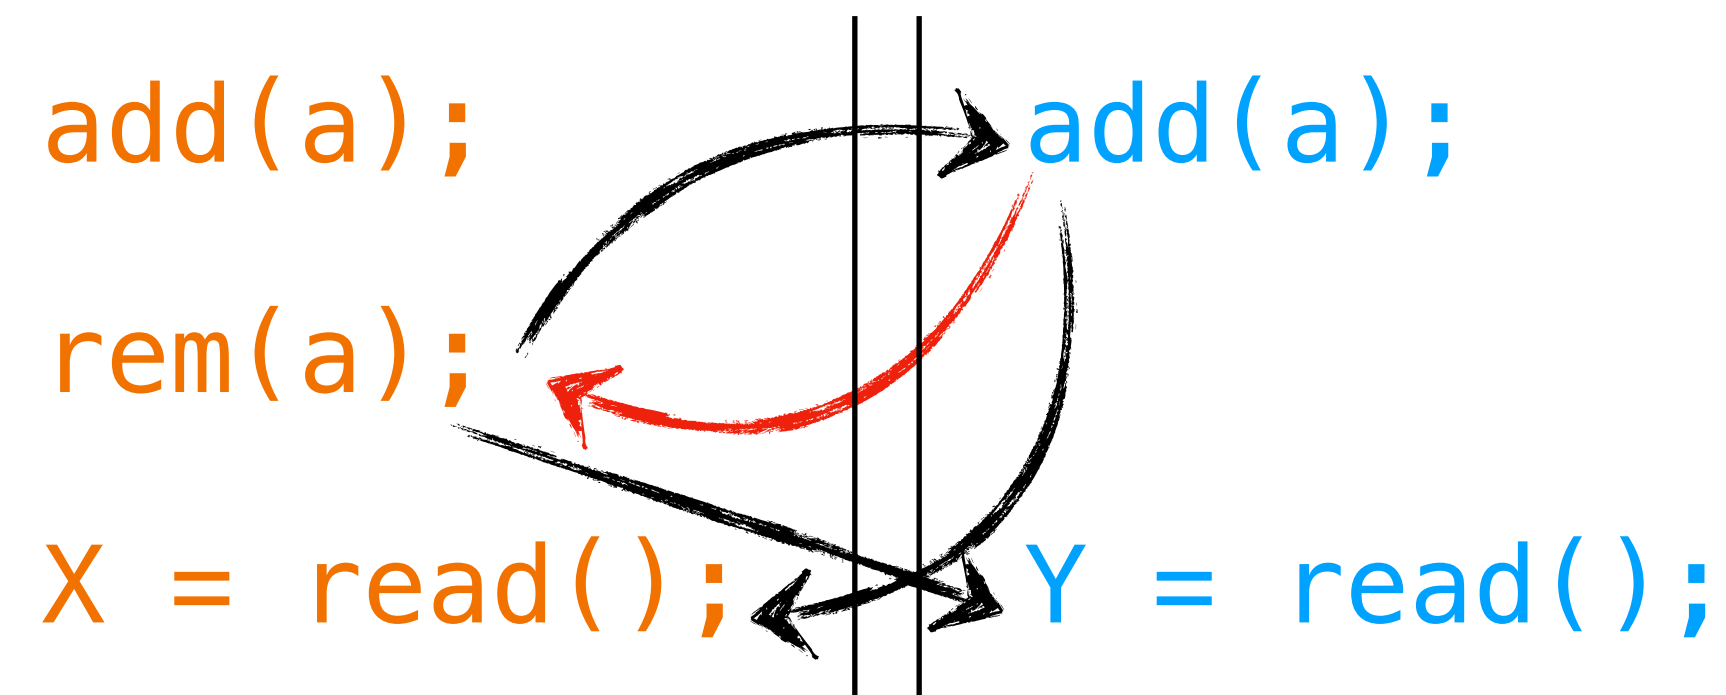
|             |  |             |
|-------------|--|-------------|
| add(a);     |  | add(a);     |
| rem(a);     |  |             |
| X = read(); |  | Y = read(); |

$$a \in X \implies a \in Y$$

- Since Add-wins set is RA-linearizable, you can use **totally ordered trace** and the **sequential reasoning**

|         |         |         |             |             |
|---------|---------|---------|-------------|-------------|
| add(a); | rem(a); | add(a); | X = read(); | Y = read(); |
| {a}     | {}      | {a}     | X = {a}     | Y = {a}     |

# Using RA-linearizability for verification



$$a \in X \implies a \in Y \quad \checkmark$$

- Let's try to make the statement false
  - Make  $a \in X$  true and  $a \in Y$  false

# Replication-aware Linearizability

RESEARCH-ARTICLE

## Replication-aware linearizability


**Authors:**  [Chao Wang](#),  [Constantin Enea](#),  [Suha Orhun Mutluergil](#),  [Gustavo Petri](#) | [Authors Info & Clai](#)

[PLDI 2019: Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation](#)  
<https://doi.org/10.1145/3314221.3314617>

- Presented a proof methodology to show that a CRDT is linearisable
- Not automated or mechanised





# Neem: Automatic verification of RDTs

- Definition of RA-linearizability for MRDTs
- A novel induction scheme for MRDTs and state-based CRDTs to **automatically** verify RA-linearizability
- Implemented in F\*


RESEARCH-ARTICLE | OPEN ACCESS | 

[X](#) [in](#) [reddit](#) [f](#) [email](#)



## Automatically Verifying Replication-Aware Linearizability

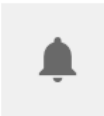




Authors:  Vimala Soundarapandian,  Kartik Nagar,  Aseem Rastogi,  KC Sivaramakrishnan | [Authors Info & Claims](#)

Proceedings of the ACM on Programming Languages, Volume 9, Issue OOPSLA1 • Article No.: 111, Pages 871 - 897  
<https://doi.org/10.1145/3720452>

Published: 09 April 2025 [Publication History](#) 

Related Artifact: [Automatically Verifying Replication-aware Linearizability - artifact](#) • April 2025 • software • <https://doi.org/10.5281/zenodo.14591614>

 0  77

    PDF  eReader

github.com/prismlab/neem

README MIT license

## Neem

Neem is a framework for automated verification of mergeable replicated data types (MRDTs) and state-based convergent replicated data types (CRDTs). See <https://dl.acm.org/doi/10.1145/3720452>.

### Development Environment




Easiest way to get started is to use the devcontainer.

```
$ git clone https://github.com/prismlab/neem
$ cd neem
$ code . # Start VSCode
```

VSCoDe will notify that there is a devcontainer associated with this repo and whether to open this repo in a devcontainer.

**Packages**  
No packages published  
[Publish your first package](#)

**Contributors** 3

-  vimcy7 Vimala S
-  kayceesrk KC Sivaramakrish...
-  aseemr Aseem Rastogi

**Languages**

- F\* 97.2%
- Shell 2.4%
- Dockerfile 0.4%


**Suggested workflows**

# Specifying RDTs in Neem

- Not all operations commute
  - **Add-wins set** —  $\text{add}(a)$  and  $\text{rem}(a)$  do not commute
  - Specify ordering using the **Conflict Resolution** relation  $rc = \{(rem_a, add_a) \mid a \in \mathbb{E}\}$
  - Linearization order  $lo$  must be compatible with  $rc$  for concurrent events
- Neem developers provide
  - MRDT = Sequential Data Type + 3-way merge
  - Conflict Resolution  $rc$  relation

# Add-wins Set


|                         |                                                                                                                                                                                         |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>State</b>            | 1: $\Sigma = \mathcal{P}(\mathbb{E} \times \mathcal{T})$                                                                                                                                |
| <b>Updates</b>          | 2: $O = \{\text{add}_a, \text{rem}_a \mid a \in \mathbb{E}\}$                                                                                                                           |
| <b>Queries</b>          | 3: $Q = \{\text{rd}\}$                                                                                                                                                                  |
| <b>Init State</b>       | 4: $\sigma_0 = \{\}$                                                                                                                                                                    |
| <b>Update behaviour</b> | 5: $\text{do}(\sigma, t, \_, \text{add}_a) = \sigma \cup \{(a, t)\}$<br>6: $\text{do}(\sigma, \_, \_, \text{rem}_a) = \sigma \setminus \{(a, i) \mid (a, i) \in \sigma\}$               |
| <b>Merge</b>            | 7: $\text{merge}(\sigma_{\top}, \sigma_1, \sigma_2) =$<br>$(\sigma_{\top} \cap \sigma_1 \cap \sigma_2) \cup (\sigma_1 \setminus \sigma_{\top}) \cup (\sigma_2 \setminus \sigma_{\top})$ |
| <b>Query behaviour</b>  | 8: $\text{query}(\sigma, \text{rd}) = \{a \mid (a, \_) \in \sigma\}$                                                                                                                    |
| <b>Resolve conflict</b> | 9: $\text{rc} = \{(\text{rem}_a, \text{add}_a) \mid a \in \mathbb{E}\}$                                                                                                                 |



$$\{(a,2)\} = \text{add}_a(\text{rem}_a\{(a,1)\})$$

# Bottom up linearisation

$$rc = \{(rem_a, add_a) \mid a \in \mathbb{E}\}$$



To show

$$\mu(\sigma, e_1(\sigma), e_2(\sigma)) = e_2(e_1(\sigma))$$

[BOTTOMUP-2-OP]

$$\frac{e_1 \neq e_2 \quad e_1 \xrightarrow{rc} e_2 \vee e_1 \rightleftharpoons e_2}{\mu(l, e_1(a), e_2(b)) = e_2(\mu(l, e_1(a), b))}$$

[BOTTOMUP-1-OP]

$$\frac{(e_{\top} \neq \epsilon \wedge e_1 \neq e_{\top}) \vee (e_{\top} = \epsilon \wedge l = b)}{\mu(e_{\top}(l), e_1(a), e_{\top}(b)) = e_1(\mu(e_{\top}(l), a, e_{\top}(b)))}$$

[BOTTOMUP-0-OP]

$$\mu(e_{\top}(l), e_{\top}(a), e_{\top}(b)) = e_{\top}(\mu(l, a, b))$$

[MERGEIDEMPOTENCE]


$$\mu(a, a, a) = a$$

[MERGECOMMUTATIVITY]

$$\mu(l, a, b) = \mu(l, b, a)$$

# Bottom up linearisation

$$rc = \{(rem_a, add_a) \mid a \in \mathbb{E}\}$$



To show

$$\mu(\sigma, e_1(\sigma), e_2(\sigma)) = e_2(e_1(\sigma))$$

[BOTTOMUP-2-OP]

$$\frac{e_1 \neq e_2 \quad e_1 \xrightarrow{rc} e_2 \quad \vee \quad e_1 \rightleftharpoons e_2}{\mu(l, e_1(a), e_2(b)) = e_2(\mu(l, e_1(a), b))}$$

[BOTTOMUP-1-OP]

$$\frac{(e_{\top} \neq \epsilon \wedge e_1 \neq e_{\top}) \vee (e_{\top} = \epsilon \wedge l = b)}{\mu(e_{\top}(l), e_1(a), e_{\top}(b)) = e_1(\mu(e_{\top}(l), a, e_{\top}(b)))}$$

[BOTTOMUP-0-OP]

$$\mu(e_{\top}(l), e_{\top}(a), e_{\top}(b)) = e_{\top}(\mu(l, a, b))$$

[MERGEIDEMPOTENCE]


$$\mu(a, a, a) = a$$

[MERGECOMMUTATIVITY]

$$\mu(l, a, b) = \mu(l, b, a)$$

# Bottom up linearisation

$$rc = \{(rem_a, add_a) \mid a \in \mathbb{E}\}$$



To show

$$e_2(\mu(\sigma, e_1(\sigma), \sigma)) = e_2(e_1(\sigma))$$

[BOTTOMUP-2-OP]

$$\frac{e_1 \neq e_2 \quad e_1 \xrightarrow{rc} e_2 \quad \vee \quad e_1 \rightleftharpoons e_2}{\mu(l, e_1(a), e_2(b)) = e_2(\mu(l, e_1(a), b))}$$

[BOTTOMUP-1-OP]

$$\frac{(e_{\top} \neq \epsilon \wedge e_1 \neq e_{\top}) \vee (e_{\top} = \epsilon \wedge l = b)}{\mu(e_{\top}(l), e_1(a), e_{\top}(b)) = e_1(\mu(e_{\top}(l), a, e_{\top}(b)))}$$

[BOTTOMUP-0-OP]

$$\mu(e_{\top}(l), e_{\top}(a), e_{\top}(b)) = e_{\top}(\mu(l, a, b))$$

[MERGEIDEMPOTENCE]


$$\mu(a, a, a) = a$$

[MERGECOMMUTATIVITY]

$$\mu(l, a, b) = \mu(l, b, a)$$

# Bottom up linearisation

$$rc = \{(rem_a, add_a) \mid a \in \mathbb{E}\}$$



To show

$$e_2(\mu(\sigma, e_1(\sigma), \sigma)) = e_2(e_1(\sigma))$$

[BOTTOMUP-2-OP]

$$\frac{e_1 \neq e_2 \quad e_1 \xrightarrow{rc} e_2 \vee e_1 \rightleftharpoons e_2}{\mu(l, e_1(a), e_2(b)) = e_2(\mu(l, e_1(a), b))}$$

[BOTTOMUP-1-OP]

$$\frac{(e_{\top} \neq \epsilon \wedge e_1 \neq e_{\top}) \vee (e_{\top} = \epsilon \wedge l = b)}{\mu(e_{\top}(l), e_1(a), e_{\top}(b)) = e_1(\mu(e_{\top}(l), a, e_{\top}(b)))}$$

[BOTTOMUP-0-OP]

$$\mu(e_{\top}(l), e_{\top}(a), e_{\top}(b)) = e_{\top}(\mu(l, a, b))$$

[MERGEIDEMPOTENCE]


$$\mu(a, a, a) = a$$

[MERGECOMMUTATIVITY]

$$\mu(l, a, b) = \mu(l, b, a)$$

# Bottom up linearisation

$$rc = \{(rem_a, add_a) \mid a \in \mathbb{E}\}$$



To show

$$e_2(e_1(\mu(\sigma, \sigma, \sigma))) = e_2(e_1(\sigma))$$

[BOTTOMUP-2-OP]

$$\frac{e_1 \neq e_2 \quad e_1 \xrightarrow{rc} e_2 \vee e_1 \rightleftharpoons e_2}{\mu(l, e_1(a), e_2(b)) = e_2(\mu(l, e_1(a), b))}$$

[BOTTOMUP-1-OP]

$$\frac{(e_{\top} \neq \epsilon \wedge e_1 \neq e_{\top}) \vee (e_{\top} = \epsilon \wedge l = b)}{\mu(e_{\top}(l), e_1(a), e_{\top}(b)) = e_1(\mu(e_{\top}(l), a, e_{\top}(b)))}$$

[BOTTOMUP-0-OP]

$$\mu(e_{\top}(l), e_{\top}(a), e_{\top}(b)) = e_{\top}(\mu(l, a, b))$$

[MERGEIDEMPOTENCE]


$$\mu(a, a, a) = a$$

[MERGECOMMUTATIVITY]

$$\mu(l, a, b) = \mu(l, b, a)$$

# Bottom up linearisation

$$rc = \{(rem_a, add_a) \mid a \in \mathbb{E}\}$$



To show

$$e_2(e_1(\mu(\sigma, \sigma, \sigma))) = e_2(e_1(\sigma))$$

[BOTTOMUP-2-OP]

$$\frac{e_1 \neq e_2 \quad e_1 \xrightarrow{rc} e_2 \vee e_1 \rightleftharpoons e_2}{\mu(l, e_1(a), e_2(b)) = e_2(\mu(l, e_1(a), b))}$$

[BOTTOMUP-1-OP]

$$\frac{(e_{\top} \neq \epsilon \wedge e_1 \neq e_{\top}) \vee (e_{\top} = \epsilon \wedge l = b)}{\mu(e_{\top}(l), e_1(a), e_{\top}(b)) = e_1(\mu(e_{\top}(l), a, e_{\top}(b)))}$$

[BOTTOMUP-0-OP]

$$\mu(e_{\top}(l), e_{\top}(a), e_{\top}(b)) = e_{\top}(\mu(l, a, b))$$

[MERGEIDEMPOTENCE]


$$\mu(a, a, a) = a$$

[MERGECOMMUTATIVITY]

$$\mu(l, a, b) = \mu(l, b, a)$$

# Bottom up linearisation

$$rc = \{(rem_a, add_a) \mid a \in \mathbb{E}\}$$



To show

$$e_2(e_1(\sigma)) = e_2(e_1(\sigma))$$

[BOTTOMUP-2-OP]

$$\frac{e_1 \neq e_2 \quad e_1 \xrightarrow{rc} e_2 \vee e_1 \rightleftharpoons e_2}{\mu(l, e_1(a), e_2(b)) = e_2(\mu(l, e_1(a), b))}$$

[BOTTOMUP-1-OP]

$$\frac{(e_{\top} \neq \epsilon \wedge e_1 \neq e_{\top}) \vee (e_{\top} = \epsilon \wedge l = b)}{\mu(e_{\top}(l), e_1(a), e_{\top}(b)) = e_1(\mu(e_{\top}(l), a, e_{\top}(b)))}$$

[BOTTOMUP-0-OP]

$$\mu(e_{\top}(l), e_{\top}(a), e_{\top}(b)) = e_{\top}(\mu(l, a, b))$$

[MERGEIDEMPOTENCE]

$$\mu(a, a, a) = a$$

[MERGECOMMUTATIVITY]

$$\mu(l, a, b) = \mu(l, b, a)$$

# Verified MRDTs

| MRDT                          | rc Policy                                            | #LOC | Verification Time (s) |
|-------------------------------|------------------------------------------------------|------|-----------------------|
| Increment-only counter [12]   | none                                                 | 6    | 0.72                  |
| PN counter [23]               | none                                                 | 10   | 1.64                  |
| Enable-wins flag*             | disable $\xrightarrow{rc}$ enable                    | 30   | 29.80                 |
| Disable-wins flag*            | enable $\xrightarrow{rc}$ disable                    | 30   | 37.91                 |
| Grows-only set [12]           | none                                                 | 6    | 0.45                  |
| Grows-only map [23]           | none                                                 | 11   | 4.65                  |
| OR-set [23]                   | rem <sub>a</sub> $\xrightarrow{rc}$ add <sub>a</sub> | 20   | 4.53                  |
| OR-set (efficient)*           | rem <sub>a</sub> $\xrightarrow{rc}$ add <sub>a</sub> | 34   | 660.00                |
| Remove-wins set*              | add <sub>a</sub> $\xrightarrow{rc}$ rem <sub>a</sub> | 22   | 9.60                  |
| Set-wins map*                 | del <sub>k</sub> $\xrightarrow{rc}$ set <sub>k</sub> | 20   | 5.06                  |
| Replicated Growable Array [1] | none                                                 | 13   | 1.51                  |
| Optional register*            | unset $\xrightarrow{rc}$ set                         | 35   | 200.00                |
| Multi-valued Register*        | none                                                 | 7    | 0.65                  |
| JSON-style MRDT*              | Fig. 13                                              | 26   | 148.84                |



*Neem also supports verification of RA-linearizability of state-based CRDTs*

<https://github.com/prismlab/neem>

# Challenges with Neem

- Automated verification in F\* returns yes / no /  $\neg$  (ツ) /
- Not pleasant for engineering
- No counterexamples!
- TCB includes the SMT solver
- F\* SMT-aided proofs are brittle
  - Some lemmas need specific heuristics for fuel, unfolding, etc.
  - z3 upgrade breaks proofs!

# Sal: Multi-modal verification of RDTs

- Move Neem from F\* to Lean
- Advantages
  - Smaller TCB
  - Property-based testing
    - Counter-example generation and visualisation
- *Agentic proof-oriented programming*
  - Fundamentally changing how we construct software
  - Coding Agents are great at Lean

# Agentic Proof Oriented Programming in Sal

- Claude Opus 4.7, with Aristotle for hard proofs
- In the last couple of weeks, compared to the Sal paper
  - **13 new CRDTs and 2 new MRDTs** have been verified
- Including Peritext, Priority Queue, Bounded Counters

## Peritext: A CRDT for Collaborative Rich Text Editing

GEOFFREY LITT, MIT CSAIL, USA  
SARAH LIM, UC Berkeley, USA  
MARTIN KLEPPMANN, University of Cambridge, United Kingdom  
PETER VAN HARDENBERG, Ink & Switch, USA

Conflict-Free Replicated Data Types (CRDTs) support decentralized collaborative editing of shared data, enabling peer-to-peer sharing and flexible branching and merging workflows. While there is extensive work on CRDTs for plain text, much less is known about CRDTs for rich text with formatting. No algorithms have been published, and existing open-source implementations do not always preserve user intent.

In this paper, we describe a model of intent preservation in rich text editing, developed through a series of concurrent editing scenarios. We then describe Peritext, a CRDT algorithm for rich text that satisfies the criteria of our model. The key idea is to store formatting spans alongside the plaintext character sequence, linked to a stable identifier for the first and last character of each span, and then to derive the final formatted text from these spans in a deterministic way that ensures concurrent operations commute.

We have prototyped our algorithm in TypeScript, validated it using randomized property-based testing, and integrated it with an editor UI. We also prove that our algorithm ensures convergence, and demonstrate its causality preservation and intention preservation properties.

531

## Conflict-free Replicated Priority Queue: Design, Verification and Evaluation

Yuqi Zhang, Lingzhi Ouyang, Yu Huang\*, Xiaoxing Ma  
State Key Laboratory for Novel Software Technology, Nanjing University  
Nanjing, Jiangsu, China  
cs.yqzhang@gmail.com, lingzhi.ouyang@smail.nju.edu.cn, {yuhuang, xxm}@nju.edu.cn

### ABSTRACT

Internet-scale distributed systems often rely on replication to achieve fault-tolerance and load distribution. To provide low latency and high availability, the systems are often required to accept updates on one replica immediately and then propagate the updates among replicas asynchronously. Conflict-free Replicated Data Type (CRDT) is a principled approach to addressing the challenge for these systems to resolve conflicts among concurrent updates. Although many CRDTs have been studied, little research has been done on Conflict-free Replicated Priority Queue (CRPQ), which is a collection of elements that focuses on maintaining element orderings based on their priority values, and can be used in many applications scenarios such as task scheduling and network routing. In this work, we discuss the design rationales of CRPQs and introduce two CRPQ designs: Add-Win CRPQ and Remove-Win CRPQ. The correctness

service availability are widely regarded as the most critical factors. To address the requirements of such applications, many distributed systems are designed for low latency and high availability in the first place [20, 24]. The server of the system handles the update requests immediately without communicating with remote replicas. The updates are propagated to remote replicas asynchronously.

These distributed systems achieve low latency and high availability by accepting weak consistency due to the CAP theorem [14, 19]. The *strong eventual consistency* is a widely accepted weak consistency that provides certain guarantees to developers of upper-layer applications [33]. It ensures that two replicas will reach the same state if they have received the same set of updates. The conflicts of concurrent updates are resolved during such process. Conflict-free Replicated Data Type (CRDT) is a framework for designing replicated data types that provide strong eventual consistency [29, 32]. It

## Extending Eventually Consistent Cloud Databases for Enforcing Numeric Invariants

Valter Balegas, Diogo Serra, Sérgio Duarte  
Carla Ferreira, Rodrigo Rodrigues, Nuno Preguiça  
NOVA LINES/FCT/Universidade Nova de Lisboa

Marc Shapiro, Mahsa Najafzadeh  
INRIA / LIP6

**Abstract**—Geo-replicated databases often operate under the principle of eventual consistency to offer high-availability with low latency on a simple key/value store abstraction. Recently, some have adopted commutative data types to provide seamless reconciliation for special purpose data types, such as counters. Despite this, the inability to enforce numeric invariants across all replicas still remains a key shortcoming of relying on the limited guarantees of eventual consistency storage.

correct operation [17]. Real world examples where enforcing invariants is essential are advertisement services, virtual wallets or to maintain stocks. However, enforcing this condition using counters implemented on eventually consistent cloud database is impossible. This is because counter updates can occur concurrently, making it impossible to detect if the limit is exceeded before the operation concludes.

Pen-and-paper proof + PBT

Verified in TLA+

Verified in TLA+

# Weakness of Sal/Neem RA-linearizability

- Original RA-linearizability assumes that a separate sequential spec

A history  $h = (E, vis)$ ,  $E \subseteq \text{Queries} \uplus \text{Updates}$ , is RA-linearizable w.r.t. a sequential specification

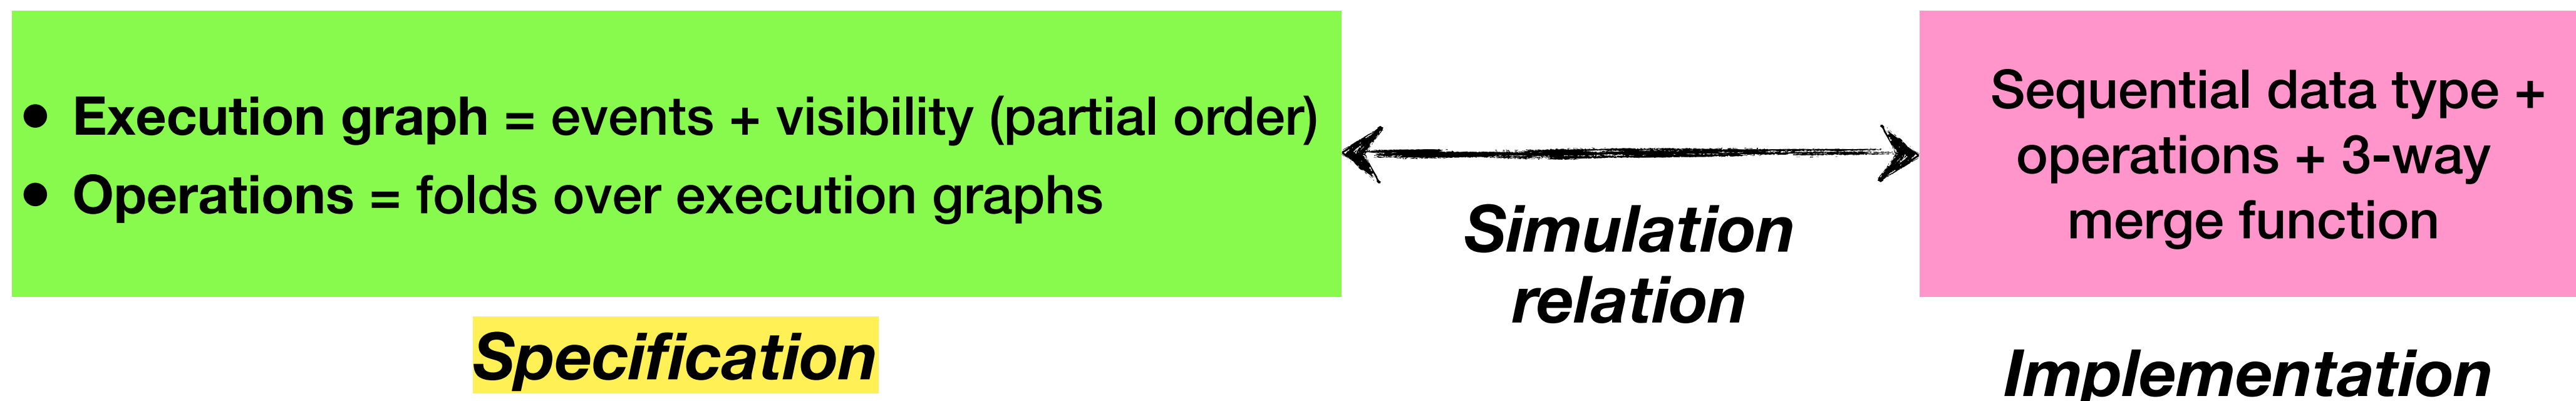
**Spec** if there exists a total order  $seq$  on  $E$  (same events) such that:

(i)  $vis \cup seq$  is acyclic;

(ii)  $seq \downarrow_{\text{Updates}} \in \text{Spec}$ ;

(iii)  $\forall \ell_{qr} \in E, (seq \downarrow_{vis^{-1}(\ell_{qr}) \cap \text{Updates}}) \cdot \ell_{qr} \in \text{Spec}$ .

- Peepul also expects a separate spec



# Weakness of Sal/Neem RA-linearizability

- Sal/Neem RA-linearizability — *operation definition itself as the sequential spec*
  - $\sigma_5 = \sigma_6 = \text{linearization}(\{op_1, op_2, op_3, op_4\}) \sigma_0$
- Works well for RDTs where
  - Concrete state  $\sim$  sequential state & **do** operations are “correct”
  - For example, efficient OR set, counters, etc.
- However
  - Broken **do** operations, broken RDT
  - RDTs that use inflationary sets + read-side heavy lifting are trivially correct
    - RGA, Peritext, priority queue...

# ReadSide Proofs for Intent Preservation

- Additional lean proofs on top of RA-linearizability
- In a multi-valued register
  - Sequential writes overwrite; concurrent writes are preserved
  - Neem implementation didn't respect the former; ***Wasn't caught!***
  - Caught by read-side proofs

```
/-- **Sequential write overwrites prior writes.** If a write at `ts1`  
is currently visible, and we apply `Write v2 O2` where `O2` includes  
`ts1` (which is what a well-formed sequential write would do - its  
prepare-time snapshot includes all currently-visible ts), then  
`v1` is no longer visible at the witness `ts1`. -/  
theorem sequential_write_supersedes  
  (s : concrete_st) (v1 v2 ts1 ts2 rid : N) (O2 : set N)  
  (h_t1_in_O2 : mem ts1 O2 = true) :  
  mem ts1 (Prod.snd (do_ s (ts2, rid, app_op_t.Write v2 O2))) = true := by  
  simp only [mem] at h_t1_in_O2  
  simp [do_, h_t1_in_O2]
```

```
/-- **Concurrent writes both survive (headline).** Two replicas  
diverge from common state `s`. Each applies its own `Write` with a  
snapshot that does NOT include the other's fresh `ts` (true by  
`distinct_ops` for any reachable execution: each replica's snapshot  
was taken before the other's write existed). After merge, both  
`v1` and `v2` are visible.  
  
Premise list:  
  * The two new ts values are fresh (not in `s.writes`, not in `s.removed`).  
  * Neither snapshot includes its own ts nor the other's ts.  
  
These all hold structurally for any well-formed replica execution. -/  
theorem concurrent_writes_both_visible  
  (s : concrete_st) (v1 v2 ts1 ts2 rid1 rid2 : N) (O1 O2 : set N)  
  (h_fresh_t1_O1 : mem ts1 O1 = false)  
  (h_fresh_t2_O1 : mem ts2 O1 = false)  
  (h_fresh_t1_O2 : mem ts1 O2 = false)
```

# ReadSide Proofs — Peritext

## Peritext: A CRDT for Collaborative Rich Text Editing

GEOFFREY LITT, MIT CSAIL, USA

SARAH LIM, UC Berkeley, USA

MARTIN KLEPPMANN, University of Cambridge, United Kingdom

PETER VAN HARDENBERG, Ink & Switch, USA

Conflict-Free Replicated Data Types (CRDTs) support decentralized collaborative editing of shared data, enabling peer-to-peer sharing and flexible branching and merging workflows. While there is extensive work on CRDTs for plain text, much less is known about CRDTs for rich text with formatting. **No algorithms have been published, and existing open-source implementations do not always preserve user intent.**

In this paper, we describe a model of intent preservation in rich text editing, developed through a series of concurrent editing scenarios. We then describe Peritext, a CRDT algorithm for rich text that satisfies the criteria of our model. The key idea is to store formatting spans alongside the plaintext character sequence, linked to a stable identifier for the first and last character of each span, and then to derive the final formatted text from these spans in a deterministic way that ensures concurrent operations commute.

We have prototyped our algorithm in TypeScript, validated it using randomized property-based testing, and integrated it with an editor UI. **We also prove that our algorithm ensures convergence, and demonstrate its causality preservation and intention preservation properties.**

# PeriText — Intent preservation

Alice: **The fox jumped.**  
Bob: The **fox jumped.**

**The fox jumped.**  
bold

## Concurrent overlapping marks

The **fox jumped.**  
bold

The quick **fox jumped over the dog.**  
bold

## Text insertion at span boundaries

# ReadSide Proofs — Peritext

| Paper claim                                                         | Sal artifact                                                                                                                                                                                                                                                                 | Status / notes                                                                                       |
|---------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------|
| §A.1 Convergence                                                    | 24 RA-linearizability VCs in <code>Peritext_CRDT.lean</code> and <code>Peritext_MRDT.lean</code> ( <code>merge_comm</code> , <code>merge_idem</code> , <code>base_2op</code> , <code>ind_lca_2op</code> , <code>inter_*</code> , <code>ind_*</code> , <code>lem_0op</code> ) | ✓ All closed via <code>by sal</code> , kernel-checked. Universally quantified over op interleavings. |
| §3.1 Causality preservation                                         | Framework assumption (causal delivery); <code>after</code> s map structurally encodes insert causality; <code>causal_order_visible_lt</code> (RGA readside) lifts it to visible-order                                                                                        | ✓ Same convention as paper §3.1.                                                                     |
| §A.2 Intent preservation — convergence at the read                  | <code>readRichText_visible_convergent</code>                                                                                                                                                                                                                                 | ✓ Read-side analogue of merge convergence.                                                           |
| §A.2 Ex 1 — insertion within a span                                 | <code>insert_within_span_in_span_visible</code> + propagation lemmas                                                                                                                                                                                                         | ✓                                                                                                    |
| §A.2 Ex 2 — overlapping same-type Adds                              | <code>partial_overlap_all_adds_formatted_visible</code>                                                                                                                                                                                                                      | ✓                                                                                                    |
| §A.2 Ex 3 — different mark types coexist                            | <code>different_type_adds_coexist_visible</code>                                                                                                                                                                                                                             | ✓                                                                                                    |
| §A.2 Ex 4 — same-type, different values (colors)                    | not formalised                                                                                                                                                                                                                                                               | ⚠ Requires <code>markValue</code> field on <code>MarkOp</code> — state-shape change.                 |
| §A.2 Ex 5 — bold vs non-bold conflict                               | <code>add_wins_over_concurrent_remove_visible</code> + <code>no_add_cover_implies_unformatted_visible</code>                                                                                                                                                                 | ✓ LWW by <code>opId</code> (paper §4.4).                                                             |
| §A.2 Ex 6 — overlapping comments via distinct <code>markType</code> | follows from <code>different_type_adds_coexist_visible</code>                                                                                                                                                                                                                | ✓                                                                                                    |
| §A.2 Ex 7 — bold-boundary insertion expands                         | <code>ex7_bold_older_sibling_in_span</code> + <code>bold_expand_in_span_visible</code> (via <code>bold_expand_reach</code> )                                                                                                                                                 | ✓                                                                                                    |
| §A.2 Ex 8 — link-boundary insertion does not expand                 | <code>ex8_link_descendant_visible_lt_endId</code> + <code>_not_in_span_visible_of_wf</code>                                                                                                                                                                                  | ✓ Uses <code>wf_after</code> s acyclicity.                                                           |
| (Beyond paper) Anchors survive tombstones                           | <code>anchors_survive_tombstones_visible</code>                                                                                                                                                                                                                              | ✓ Implicit in §4.4 paper discussion.                                                                 |

# Sal Playground

## Sal CRDT playgrounds

Interactive simulators for the CRDTs and MRDTs verified in [Sal](#). CRDTs do two-way merge (pick a source and target, target absorbs source); MRDTs do three-way merge over a git-style commit DAG with LCA computed from the history. Toggle the concrete state to see the lattice layer that makes convergence work.

### CRDTs (two-way merge)

#### Counters

##### Increment-Only Counter

The simplest CRDT: every replica owns its own counter slot; merge is pointwise max.

##### PN-Counter

Two G-counters glued together: one tracks increments, one tracks decrements. The abstract value is their difference.

##### Bounded Counter

## Peritext (text-only)

*Rich-text CRDT: RGA for the character sequence plus a flat set of anchor-attached mark ops (bold, italic). A char is formatted when the highest-opld mark covering it is an Add; concurrent Add beats concurrent Remove on the same range.*

Show concrete (lattice) state

Merge  →

#### R0

Value:  
acab

insert  after

remove

mark  start  end

#### R1

Value:  
acab

insert  after

remove

mark  start  end

[fplaunchpad.org/sal](https://fplaunchpad.org/sal)

# Take aways

- **Automated verification for RDTs is pragmatic and necessary**
- **VCs for automated verification provide guardrails for agentic proof-oriented programming**
  - John Regehr, “Zero-Degree-of-Freedom LLM Coding using Executable Oracles”, [https://john.regehr.org/writing/zero\\_dof\\_programming.html](https://john.regehr.org/writing/zero_dof_programming.html)

# Take aways

- **Agents make proof-engineering cheap**
  - But spec engineering needs human guidance; capturing intent
  - Agentic read-side proofs worked well in cases when the theorems were written down on paper
  - Small Proof-oriented Tests (SPOT), RiSE MSR blog post: <https://risemsr.github.io/blog/2026-04-16-spotting-specs/>
    - Peritext intent preservation examples are exactly this.
- **Automated verification + Agentic proof engineering allows scaling to real-world RDTs**
  - Peritext was formalised in a couple of days
  - 1100 lines of Lean proofs

*What will you prove next?*