

# **Modular Verification of Op-Based CRDTs in Separation Logic**

**Abel Nieto, Léon Gondelman, Alban Reynaud, Amin Timany, Lars Birkedal**  
**Aarhus University**

# CAP Theorem

- Informally: no distributed data store can be all of the following: (strongly) *Consistent, Available, and Partition tolerant*.
- Often presented as “choose 2 out of 3”, but some pairs don’t make sense. Additionally, given enough time partitions are *unavoidable*.
- Better phrasing: given a network partition, your system can be (strongly) consistent or available, but not *both*.

# (Strong) Eventual Consistency

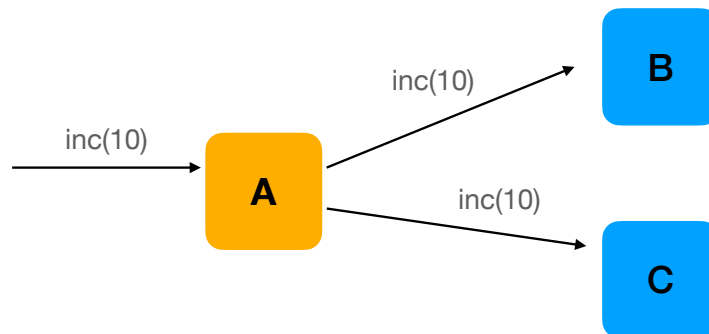
- SEC trades consistency in favour of availability.
- Replica states are allowed to diverge, but must eventually converge.
- Eventual delivery: updates eventually reach all correct replicas.
- Convergence: replicas that have delivered the same updates must be in equivalent states.
- CRDTs: a class of distributed data structures with SEC.

# State-based CRDTs

- Updates communicated by sending entire state to other replicas.
- States taken from join semi-lattice, and “merging” states is taking their LUBs.
- Cons: encoding of data type semantics into lattice state can be tricky, inefficient if state is large (but there are pros too)

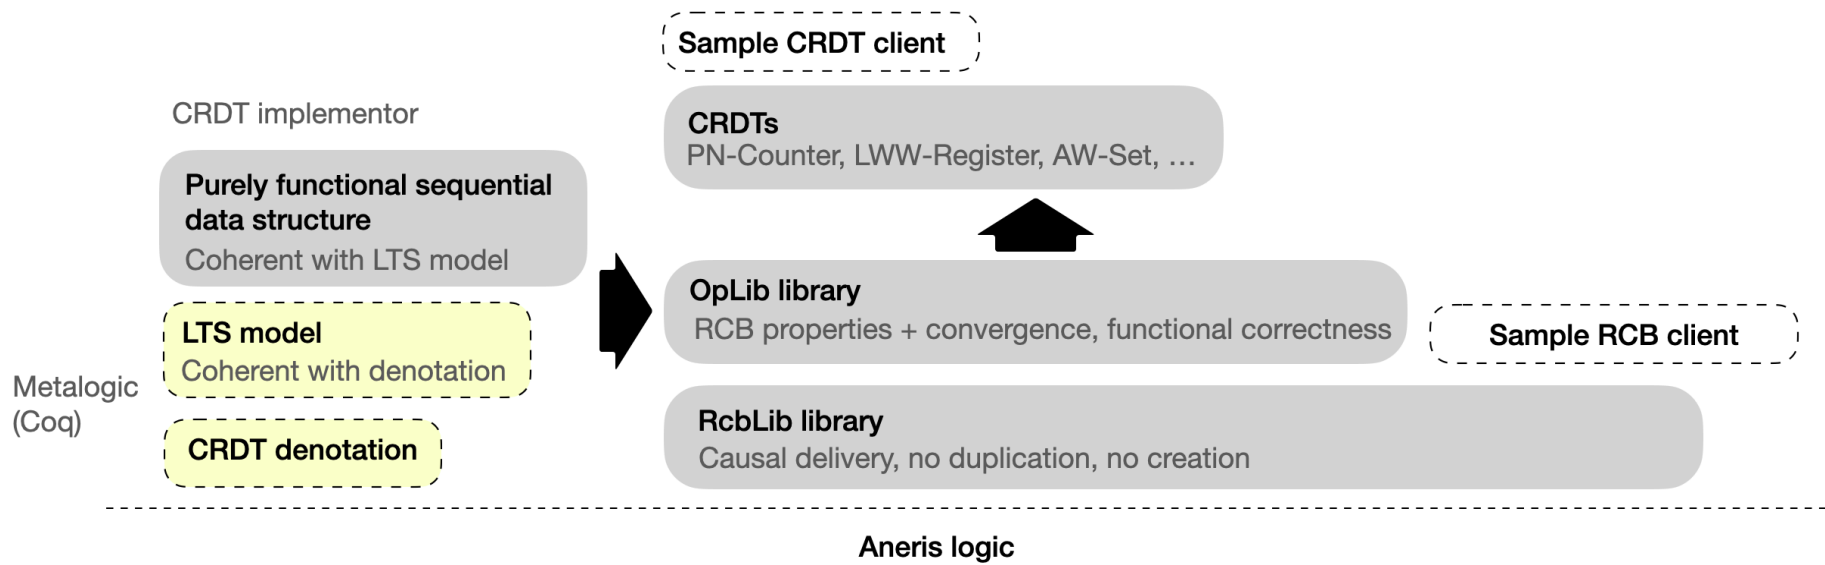
# Op-based CRDTs

- Updates communicated by sending individual operations to other replicas
- Simpler design, but requires exactly-once delivery of operations
- Multiple things can go wrong: message dropped or duplicated, or replica dies



# This Work

- We have implemented in OCaml and verified in Aneris a framework for building op-based CRDTs
- We used the framework to implement 12 example CRDTs, including higher-order combinators
- Our specifications are the first to be both *modular* and about *runnable implementations* (as opposed to protocols)
- For the first time, our formalisation of CRDTs includes a general-purpose library for Reliable Causal Broadcast (an exactly-once delivery protocol)



# Causal Broadcast

- Interface: `init(addr)`, `broadcast(msg)`, `deliver()`
- Guarantees: no duplication, no creation, and causal delivery
- Causal delivery: for any message  $m_1$  that potentially caused  $m_2$  (i.e.  $m_1 \rightarrow m_2$ ) then every node delivers  $m_1$  before delivering  $m_2$



# Resources for Causality

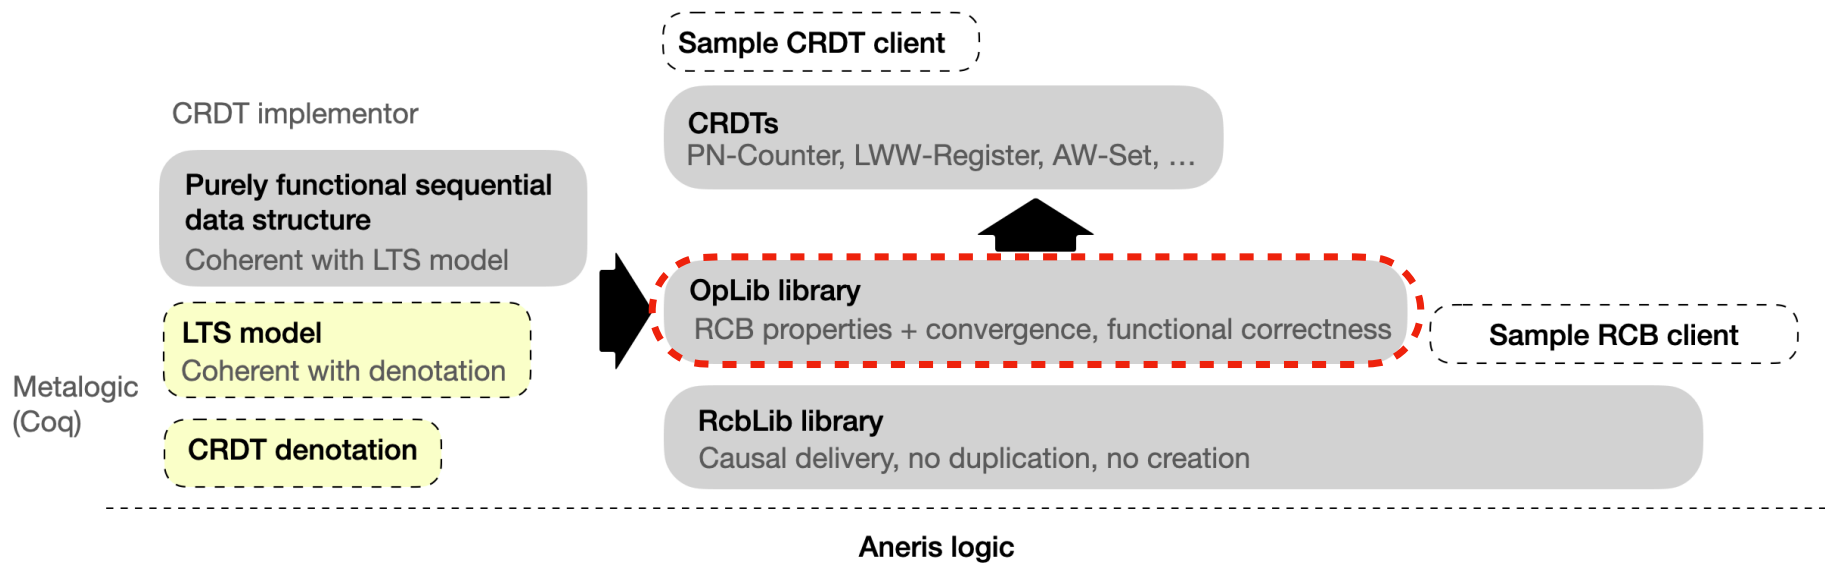
- Piggyback on Gondelman et al. (POPL'20): a causally-consistent key-value store
- Locally, track set of messages delivered at replica  $i$ :  $\text{OwnLocal}(i, s)$
- Globally, track set of all messages ever sent:  $\text{OwnGlobal}(h)$
- We can then prove resource laws: e.g. causality

$$\boxed{\text{GlobalInv}}^{\mathcal{N}_{\text{GI}}} * \text{OwnLocal}(i, s) * \text{OwnGlobalSnapshot}(h) \vdash \Rightarrow_{\mathcal{E}} \forall a \in s, w \in h. \text{vc}(w) < \text{vc}(a) \Rightarrow \exists a' \in s. [a'] = w$$

- The resources also allow us to give specifications to the broadcast and deliver functions
- Simplified broadcast spec:
 
$$\{ \text{OwnGlobal}(h) * \text{OwnLocal}(i, s) \}$$

$$\langle ip_i; \text{broadcast}(p) \rangle$$

$$\{ m. \text{payload}(m) = p * \text{OwnGlobal}(h \uplus \{m\}) * \text{OwnLocal}(i, s \uplus \{m\}) \}$$



# From Purely-Functional to CRDT

- We start with a purely-functional counter:  
an initial state (0) and  
a function to get from a state to the next ( $\text{effect}(c, n) = c + n$ )
- To turn it into a CRDT we need:
- A way to propagate operations (RCB).
- A way to (concurrently) apply remote operations.
- A way to manage mutable state (because of the above).

# OpLib

- A library for implementing operation-based CRDTs
- User (CRDT implementor) provides initial state and effect function
- They get back a fully-fledged CRDT

```
let effect msg counter =  
  let ((delta, _x), _y) = msg in  
  counter + delta  
  
let init_st () = 0  
  
let crdt = fun () -> (init_st, effect)  
  
let init addr rid =  
  let initRes = oplib_init int_ser int_deser addr rid crdt in  
  let (get_state, update) = initRes in  
  (get_state, update)
```

# Specifying OpLib

- Challenge: the CRDT's current state (e.g. the value of the counter) depends not just on local operations, but also remote ones.
- Tracking current state (Timany et al. 2021):  
$$\text{INCRSPEC}$$
$$\{ \text{gcounter}(i, k) \}$$
$$\langle ip_i; \text{incr}() \rangle$$
$$\{ () . \exists m. k < m * \text{gcounter}(i, m) \}$$
- Solution: don't track the current state. Instead, track local events *precisely* and a *lower bound* of remote events:  
$$\text{LocSt}(i, \bullet s, \circ h)$$

# Denotations

- In general, can specify CRDT with a *denotation*: partial function from set of events (including causality data) to CRDT state
- Example (multi-value register)
  - $\llbracket s \rrbracket_{\text{mv-reg}} = \{(w, vc) \mid \exists o. (\text{write}(w), vc, o) \in s \wedge vc \in \text{Maximals}(s)\}$
- Introduced in Burckhardt et al. [POPL'14] but now adapted to SL



# OpLib specs

GETSTATESPEC  
 $\langle \text{LocSt}(i, \bullet s, \circ h) \rangle$

$\langle ip_i; \text{get\_state}() \rangle$

$\left\langle \begin{array}{l} v. \exists h' w. h' \supseteq h * \text{StCoh}(w, v) * \\ \text{LocSt}(i, \bullet s, \circ h') * \llbracket s \cup h' \rrbracket = w \end{array} \right\rangle^{\mathcal{N}}$

Convergence

UPDATESPEC  
 $\langle \text{LocSt}(i, \bullet s, \circ r) * \text{GlobSt}(h) \rangle$

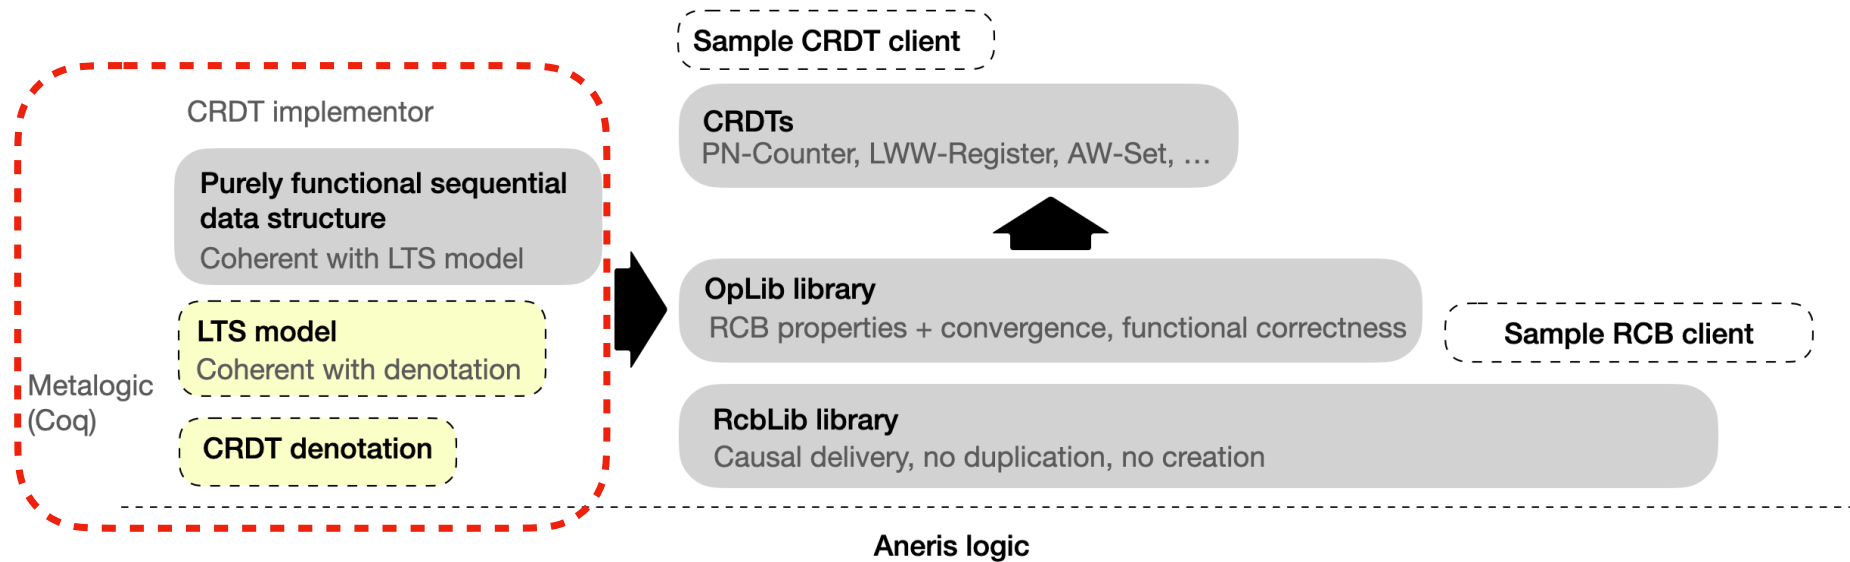
$\langle ip_i; \text{update}(v) \rangle$

$\left\langle \begin{array}{l} (). \exists ar'. r' \supseteq r * a \notin s * a \notin h * \text{payload}(a) = v * \\ \text{origin}(a) = i * a \in \text{Maximals}(h \cup \{a\}) * \\ a \in \text{Maximum}(s \cup r' \cup \{a\}) * \\ \text{LocSt}(i, \bullet s \cup \{a\}, \circ r') * \text{GlobSt}(h \cup \{a\}) \end{array} \right\rangle^{\mathcal{N}}$

# Labelled Transition Systems

- Needed: a way to connect effect function to denotation
- Done via labelled transition system  $(\text{St}, \text{Event}, \rightarrow, \sigma_0)$
- Coherence property:
  - $\llbracket \emptyset \rrbracket = \sigma_0$
  - $\forall s \ p \ e \ p'. \text{Valid}(s, e) \wedge \llbracket s \rrbracket = p \wedge p \xrightarrow{e} p' \implies \llbracket s \cup e \rrbracket = p'$
- Hoare triple for showing that `effect()` implements LTS

# OpLib recap



# Implemented CRDTs

CRDT	# lines of OCaml	# lines of Coq
Positive-Negative Counter	25	235
Grown-only Counter	26	243
Two-Part Set	25	182
Add-Wins Set	34	371
Remove-Wins Set	53	527
Grow-Only Set	22	159
Last-Writer-Wins Register	54	555
Multi-Value Register	35	334
Product Combinator	30	374
Map Combinator	34	531
Table of Positive-Negative Counters	22	74
Table of Last-Writer-Wins Registers	22	74
total	360	3585

Diagram annotations:
 

- Left side: "combined to make" with arrows pointing to the "Table of Positive-Negative Counters" and "Table of Last-Writer-Wins Registers" rows.
- Right side: Brackets grouping rows into "Simple CRDTs" (rows 1-6), "Combinators" (rows 7-8), and "Compound CRDTs" (rows 9-10).

Library	# lines of OCaml	# lines of Coq
RcbLib	196	5019
OpLib	86	3595
total	282	8614

# Conclusions

- We implemented in OCaml and verified in Aneris a framework for building op-based CRDTs, as well as many examples on top of it
- Ours is the first foundational proof of functional correctness and SEC for op-based CRDTs, as well as the first technique that is both modular and about implementations
- Future work: more complex CRDTs (collaborative text editing) and CRDTs with coordination
- Future work: state-based CRDTs

**Thank you**