


# Aneris


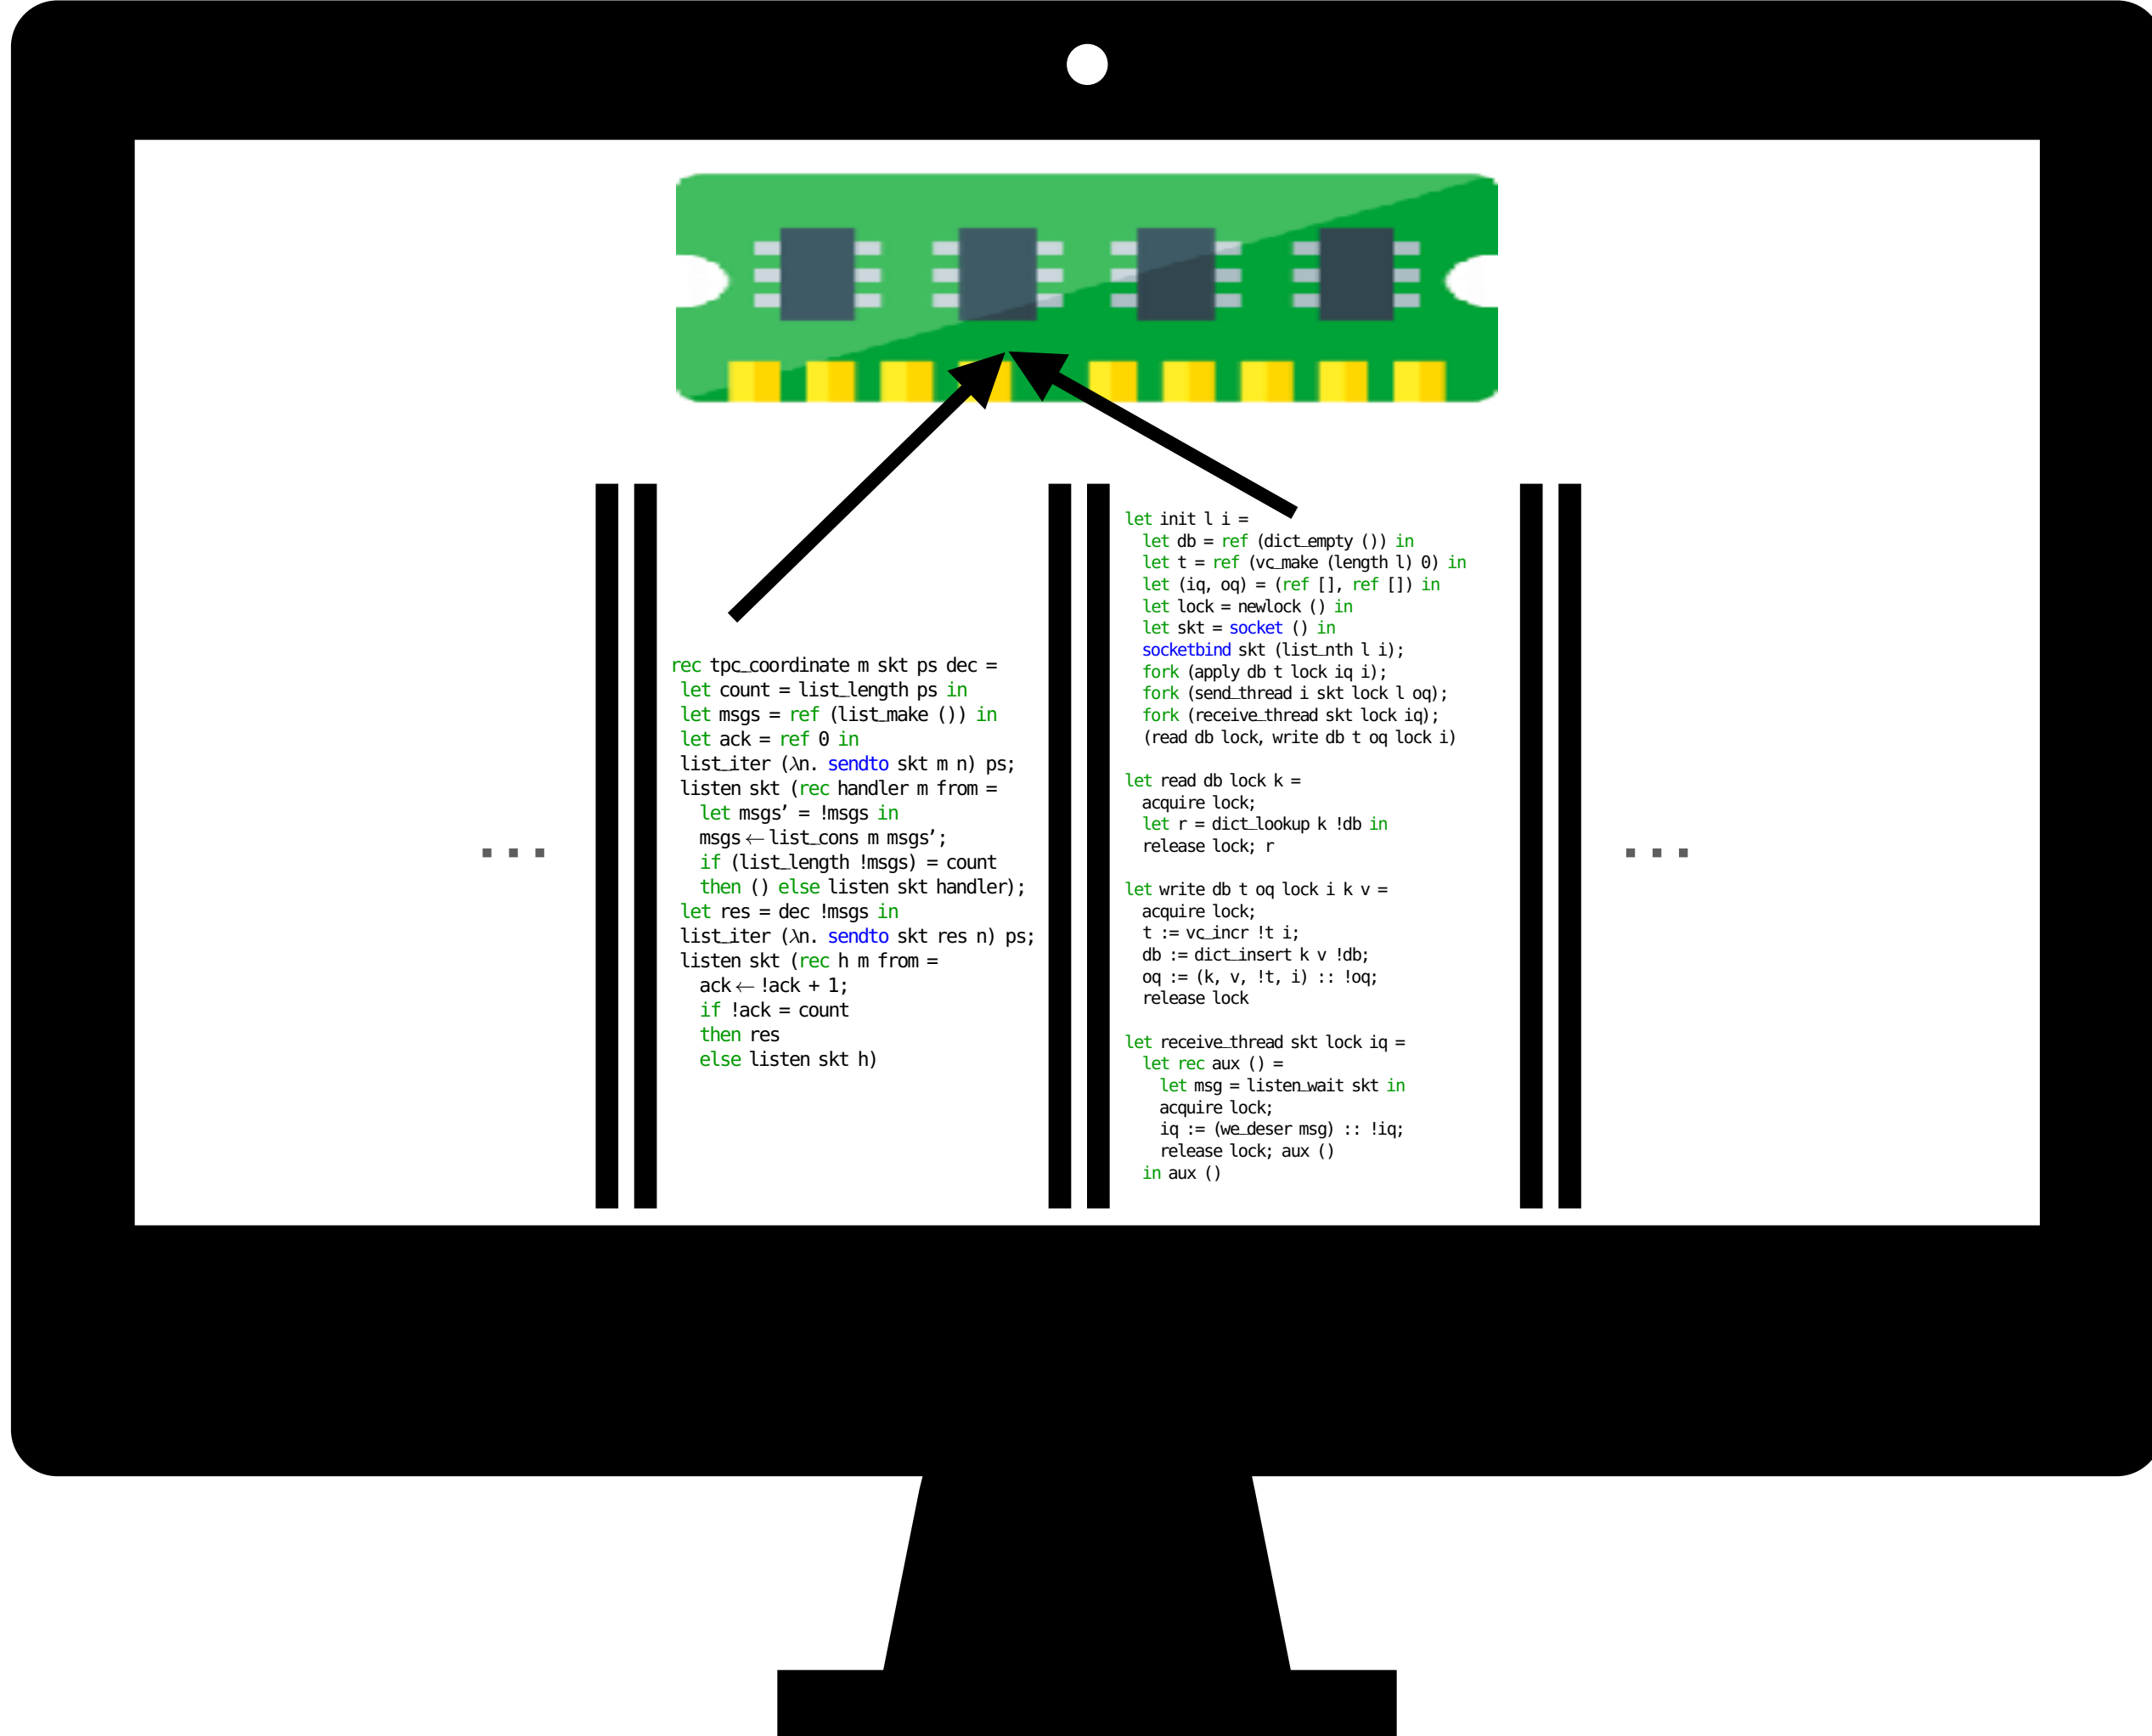
## A Mechanised Logic for Modular Reasoning about Distributed Systems

March 29, 2021 @ ESOP'20

**Simon Oddershede Gregersen**

joint work with Morten Krogh-Jespersen, Amin Timany, Marit Edna Ohlenbusch, and Lars Birkedal





```

rec tpc_coordinate m skt ps dec =
let count = list_length ps in
let msgs = ref (list_make ()) in
let ack = ref 0 in
list_iter (λn. sendto skt m n) ps;
listen skt (rec handler m from =
  let msgs' = !msgs in
  msgs ← list_cons m msgs';
  if (list_length !msgs) = count
  then () else listen skt handler);
let res = dec !msgs in
list_iter (λn. sendto skt res n) ps;
listen skt (rec h m from =
  ack ← !ack + 1;
  if !ack = count
  then res
  else listen skt h)
...

let init l i =
let db = ref (dict_empty ()) in
let t = ref (vc_make (length l) 0) in
let (iq, oq) = (ref [], ref []) in
let lock = newLock () in
let skt = socket () in
socketbind skt (list_nth l i);
fork (apply db t lock iq i);
fork (send_thread i skt lock l oq);
fork (receive_thread skt lock iq);
(read db lock, write db t oq lock i)

let read db lock k =
  acquire lock;
  let r = dict_lookup k !db in
  release lock; r

let write db t oq lock i k v =
  acquire lock;
  t := vc_incr !t i;
  db := dict_insert k v !db;
  oq := (k, v, !t, i) :: !oq;
  release lock

let receive_thread skt lock iq =
let rec aux () =
  let msg = listen_wait skt in
  acquire lock;
  iq := (we_deser msg) :: !iq;
  release lock; aux ()
in aux ()
...

```

# This work

- An ML-like language, **AnerisLang**, with higher-order store, node-local concurrency and datagram-like network sockets
- A separation logic, **Aneris**, for modular reasoning about partial correctness properties of **implementations** of distributed systems while allowing
  - **Vertical** composition
  - **Horizontal** composition
- All theory and examples are mechanised on top of the Iris framework in the Coq proof assistant

In **separation logic**, propositions denote ownership of resources, e.g.,  $\ell \mapsto v$ . Specifications only talk about the footprint of the program.

$$\{P\} e \{v.Q\}$$

These specifications can be lifted through **framing** and **binding**.

$$\frac{\{P\} e \{v.Q\}}{\{P * R\} e \{v.Q * R\}} \quad \frac{\{P\} e \{v.Q\} \quad \forall v. \{Q\} K[v] \{w.R\}}{\{P\} K[e] \{w.R\}}$$

In **concurrent separation logic**, threads are considered one at a time and we need not reason about inter-leavings of threads explicitly.

$$\frac{\{P_1\} e_1 \{v_1.Q_1\} \quad \{P_2\} e_2 \{v_2.Q_2\}}{\{P_1 * P_2\} e_1 \parallel e_2 \{v_1, v_2.Q_1 * Q_2\}}$$

Through **resource invariants**, we can express protocols on shared state.

$$\boxed{\exists v. \ell \mapsto v * \text{even}(v)}$$

In Aneris, a **distributed separation logic**, we extend these reasoning principles to allow reasoning about nodes one at a time.

$$\frac{\{P_1 * \text{FreePorts}(ip_1, A)\} e_1 \{\mathbf{True}\} \quad \{P_2 * \text{FreePorts}(ip_2, B)\} e_2 \{\mathbf{True}\}}{\{P_1 * P_2 * \text{FreeIp}(ip_1) * \text{FreeIp}(ip_2)\} e_1 \text{ } ip_1 ||| ip_2 e_2 \{\mathbf{True}\}}$$

If  $A \cap B = \emptyset$  then

$$\text{FreePorts}(ip, A) * \text{FreePorts}(ip, B) \dashv\vdash \text{FreePorts}(ip, A \cup B)$$

To express protocols on communication, we introduce **socket protocols**.


$$\Phi : \textit{Message} \rightarrow \textit{iProp}$$

With each socket address (pair of an ip and a port) we associate a protocol.

$$a \mapsto \Phi$$



# An addition service



```
rec server a =  
  let skt = socket () in  
  socketbind skt a;  
  listen skt (rec handler msg from =  
    let (n, m) = deserialize msg in  
    let res = serialize (n + m) in  
    sendto skt res from;  
  listen skt handler)
```

```
rec client x y srv a =  
  let skt = socket () in  
  socketbind skt a;  
  let m = serialize (x, y) in  
  sendto skt m srv;  
  let res = listenwait skt in  
  deserialize res
```

# Verifying an addition service

We primordially fix a socket protocol for the server.

$$\Phi_{add}(m) \triangleq \exists \Psi : \mathit{Message} \rightarrow \mathit{iProp}, (x, y : \mathbb{N}).$$

$$\mathit{body}(m) = \mathit{serialize}(x, y) *$$

$$\mathit{from}(m) \Rightarrow \Psi *$$

$$\forall m'. \mathit{body}(m') = \mathit{serialize}(x + y) \multimap \Psi(m')$$

# Verifying an addition service

$$\{(ip, p) \models \Phi_{add} * \text{FreePorts}(ip, \{p\})\}$$

server  $(ip, p)$

$$\{False\}$$

# Verifying an addition service

```
{(ip, p) ⇒ Φadd * FreePorts(ip, {p})}  
  (rec server a =  
    let skt = socket () in  
    socketbind skt a;  
    listen skt (rec handler msg from =  
      let (n, m) = deserialize msg in  
      let res = serialize (n + m) in  
      sendto skt res from;  
      listen skt handler)) (ip, p)  
  {False}
```

# Verifying an addition service

```
(rec server a =  
  let skt = socket () in  
  socketbind skt a;  
  listen skt (rec handler msg from =  
    let (n, m) = deserialize msg in  
    let res = serialize (n + m) in  
    sendto skt res from;  
    listen skt handler)) (ip, p)  
{False}
```

$$(ip, p) \models \Phi_{add}$$
$$\text{FreePorts}(ip, \{p\})$$

# Verifying an addition service

```
let skt = socket () in
socketbind skt (ip, p);
listen skt (rec handler msg from =
  let (n, m) = deserialize msg in
  let res = serialize (n + m) in
  sendto skt res from;
  listen skt handler)
```

$\{False\}$

$$(ip, p) \models \Phi_{add}$$
$$\text{FreePorts}(ip, \{p\})$$

# Verifying an addition service

$\{h \hookrightarrow \text{None}\}$

```
let skt = h in
socketbind skt (ip, p);
listen skt (rec handler msg from =
  let (n, m) = deserialize msg in
  let res = serialize (n + m) in
  sendto skt res from;
  listen skt handler)
```

$\{False\}$

$(ip, p) \models \Phi_{add}$   
 $\text{FreePorts}(ip, \{p\})$

# Verifying an addition service

```
let skt = h in
socketbind skt (ip, p);
listen skt (rec handler msg from =
  let (n, m) = deserialize msg in
  let res = serialize (n + m) in
  sendto skt res from;
  listen skt handler)
```

$\{False\}$

$$(ip, p) \models \Phi_{add}$$
$$\text{FreePorts}(ip, \{p\})$$
$$h \hookrightarrow \text{None}$$



# Verifying an addition service

```
socketbind h (ip, p);  
listen h (rec handler msg from =  
  let (n, m) = deserialize msg in  
  let res = serialize (n + m) in  
  sendto h res from;  
listen h handler)
```

$\{False\}$

$$(ip, p) \models \Phi_{add}$$
$$\text{FreePorts}(ip, \{p\})$$
$$h \hookrightarrow \text{None}$$

# Verifying an addition service

$\{h \hookrightarrow \text{Some } (ip, p)\}$

```
listen  $h$  (rec handler msg from =  
  let (n, m) = deserialize msg in  
  let res = serialize (n + m) in  
  sendto  $h$  res from;  
listen  $h$  handler)
```

$\{False\}$

$(ip, p) \models \Phi_{add}$

# Verifying an addition service

```
listen  $h$  (rec handler msg from =  
  let (n, m) = deserialize msg in  
  let res = serialize (n + m) in  
  sendto  $h$  res from;  
listen  $h$  handler)
```

$\{False\}$

$$(ip, p) \models \Phi_{add}$$
$$h \hookrightarrow \text{Some } (ip, p)$$

# Verifying an addition service

$\{\Phi_{add}(m)\}$

```
let (n, m) = deserialize body(m) in
let res = serialize (n + m) in
sendto h res from(m)
```

$(ip, p) \Rightarrow \Phi_{add}$   
 $h \hookrightarrow \mathbf{Some} (ip, p)$

# Verifying an addition service

```
let (n, m) = deserialize serialize(x, y) in  
let res = serialize (n + m) in  
sendto h res from(m)
```

$$\Psi : \text{Message} \rightarrow \text{iProp}$$
$$x, y : \mathbb{N}$$
$$(ip, p) \Vdash \Phi_{add}$$
$$h \hookrightarrow \text{Some } (ip, p)$$
$$\text{from}(m) \Vdash \Psi$$
$$\forall m'. \text{body}(m') = \text{serialize}(x + y)$$
$$\rightarrow * \Psi(m')$$

# Verifying an addition service

`sendto`  $h$   $serialize(x + y)$   $from(m)$

$$\Psi : Message \rightarrow iProp$$
$$x, y : \mathbb{N}$$
$$(ip, p) \Rightarrow \Phi_{add}$$
$$h \hookrightarrow \mathbf{Some} (ip, p)$$
$$from(m) \Rightarrow \Psi$$
$$\forall m'. \text{body}(m') = \text{serialize}(x + y)$$
$$\rightarrow * \Psi(m')$$

# Verifying a client

$$\{srv \models \Phi_{add} * \text{FreePorts}(ip, \{p\})\}$$
$$\text{client } x \ y \ srv \ (ip, p)$$
$$\{v.v = x + y\}$$

# Verifying a client

```
let skt = socket () in
socketbind skt (ip, p);
let m = serialize (x, y) in
sendto skt m srv;
let res = listenwait skt in
deserialize res
```

$$\{v.v = x + y\}$$
$$srv \models \Phi_{add}$$
$$\text{FreePorts}(ip, \{p\})$$



# Verifying a client

```
let m = serialize (x, y) in  
sendto h m srv;  
let res = listenwait h in  
deserialize res
```

$$\{v.v = x + y\}$$
$$srv \models \Phi_{add}$$
$$h \hookrightarrow \text{Some } (ip, p)$$

# Verifying a client

```
sendto h serialize(x, y) srv;  
let res = listenwait h in  
deserialize res
```

$$\{v.v = x + y\}$$
$$\Phi_{add}(m) \triangleq \exists \Psi : \text{Message} \rightarrow \text{iProp}, (x, y : \mathbb{N}).$$
$$\text{body}(m) = \text{serialize}(x, y) *$$
$$\text{from}(m) \Rightarrow \Psi *$$
$$\forall m'. \text{body}(m') = \text{serialize}(x + y) -* \Psi(m')$$
$$srv \Rightarrow \Phi_{add}$$
$$h \hookrightarrow \text{Some}(ip, p)$$

# Verifying a client

```
sendto h serialize(x, y) srv;  
let res = listenwait h in  
deserialize res
```

$$\{v.v = x + y\}$$
$$\begin{aligned}\Phi_{add}(m) &\triangleq \exists \Psi : \text{Message} \rightarrow \text{iProp}, (x, y : \mathbb{N}). \\ &\text{body}(m) = \text{serialize}(x, y) * \\ &\text{from}(m) \Rightarrow \Psi * \\ &\forall m'. \text{body}(m') = \text{serialize}(x + y) -* \Psi(m')\end{aligned}$$
$$\begin{aligned}(ip, p) &\Rightarrow \Phi_{client} \\ srv &\Rightarrow \Phi_{add} \\ h &\hookrightarrow \text{Some } (ip, p)\end{aligned}$$
$$\Phi_{client}(m) \triangleq \text{body}(m) = \text{serialize}(x + y)$$

# Verifying a client

```
let res = listenwait h in  
deserialize res
```

$$\{v.v = x + y\}$$
$$(ip, p) \models \Phi_{client}$$
$$srv \models \Phi_{add}$$
$$h \hookrightarrow \mathbf{Some} (ip, p)$$
$$\Phi_{client}(m) \triangleq \text{body}(m) = \text{serialize}(x + y)$$

# Verifying a client

$\{\Phi_{client}(m)\}$

let res = body( $m$ ) in  
deserialize res

$\{v.v = x + y\}$

$(ip, p) \models \Phi_{client}$


$srv \models \Phi_{add}$

$h \hookrightarrow \text{Some } (ip, p)$

$\Phi_{client}(m) \triangleq \text{body}(m) = \text{serialize}(x + y)$


# Horizontal composition

- The addition service showcases **horizontal compositional reasoning**
- In the paper we also introduce a load balancer, but use the existing specifications for the server and client to verify the system



# Vertical composition

- In the paper we showcase **vertical compositional reasoning** by implementing and verifying the two-phase commit protocol
- We use the two-phase commit implementation and specification as a **library** to implement and verify a replicated logging system





# Distributed Causal Memory @ POPL'21

## Modular Specification and Verification in Higher-Order Distributed Separation Logic

An implementation, modular specification, and verification of a causally-consistent distributed database and its clients in Aneris.



**Distributed Causal Memory: Modular Specification and Verification in Higher-Order Distributed Separation Logic**

LEON GONDELMAN, Aarhus University, Denmark  
SIMON ODDERSHEDE GREGERSEN, Aarhus University, Denmark  
ABEL NIETO, Aarhus University, Denmark  
AMIN TIMANY, Aarhus University, Denmark  
LARS BIRKEDAL, Aarhus University, Denmark

We present the first specification and verification of an implementation of a causally-consistent distributed database that supports modular verification of full functional correctness properties of clients and servers. We specify and reason about the causally-consistent distributed database in Aneris, a higher-order distributed separation logic for an ML-like programming language with network primitives for programming distributed systems. We demonstrate that our specifications are useful, by proving the correctness of small, but tricky, synthetic examples involving causal dependency and by verifying a session manager library implemented on top of the distributed database. We use Aneris's facilities for modular specification and verification to obtain a highly modular development, where each component is verified in isolation, relying only on the specifications (not the implementations) of other components. We have used the Coq formalization of the Aneris logic to formalize all the results presented in the paper in the Coq proof assistant.

CCS Concepts: Theory of computation → Program verification; Distributed algorithms; Separation logic

Additional Key Words and Phrases: Distributed systems; causal consistency; separation logic; higher-order logic; concurrency; formal verification

ACM Reference Format:  
Léon Gondelman, Simon Oddershede Gregersen, Abel Nieto, Amin Timany, and Lars Birkedal. 2021. Distributed Causal Memory: Modular Specification and Verification in Higher-Order Distributed Separation Logic. *Proc. ACM Program. Lang.* 5, POPL, Article 42 (January 2021), 29 pages. <https://doi.org/10.1145/3434323>

### 1 INTRODUCTION

The ubiquitous distributed systems of the present day internet often require highly available and scalable distributed data storage solutions. The CAP theorem [Gilbert and Lynch 2002] states that a distributed database cannot at the same time provide consistency, availability, and partition tolerance. Hence, many such systems choose to sacrifice aspects of data consistency for the sake of availability and fault tolerance [Bolis et al. 2013; Cheng et al. 2009; Lloyd et al. 2011; Tylenev et al. 2019]. In these systems different replicas of the database may, at the same point in time, observe different, inconsistent data. Among different notions of weaker consistency guarantees, a popular one is causal consistency. With causal consistency different replicas can observe different data, yet, it is guaranteed that data are observed in a causally related order: if a node  $u$  observes an operation

Authors' addresses: Léon Gondelman, Aarhus University, Denmark, [gondel@infos.au.dk](mailto:gondel@infos.au.dk); Simon Oddershede Gregersen, Aarhus University, Denmark, [sgregersen@cs.au.dk](mailto:sgregersen@cs.au.dk); Abel Nieto, Aarhus University, Denmark, [abelnieto@cs.au.dk](mailto:abelnieto@cs.au.dk); Amin Timany, Aarhus University, Denmark, [timany@cs.au.dk](mailto:timany@cs.au.dk); Lars Birkedal, Aarhus University, Denmark, [birkedal@cs.au.dk](mailto:birkedal@cs.au.dk).



This work is licensed under a Creative Commons Attribution 4.0 International License.  
© 2021 Copyright held by the owner(s).  
2475-3421/2021/4-ART42  
<https://doi.org/10.1145/3434323>

Proc. ACM Program. Lang., Vol. 5, No. POPL, Article 42. Publication date: January 2021.

42

Léon Gondelman, Simon Oddershede Gregersen, Abel Nieto, Amin Timany, and Lars Birkedal. 2021. Distributed Causal Memory: Modular Specification and Verification in Higher-Order Distributed Separation Logic. *Proc. ACM Program. Lang.* 5, POPL, Article 42 (January 2021), 29 pages. <https://doi.org/10.1145/3434323>



# Conclusion

- A programming language and a separation logic, **Aneris**, for modular reasoning about partial correctness properties of **implementations** of distributed systems
- Scalable development and verification of distributed systems through **vertical and horizontal compositional reasoning**
- Extensive case studies showcasing our reasoning principles
- All theory and examples are mechanised on top of the Iris separation logic framework in the Coq proof assistant

# Thank you for watching

**Contact**

[gregersen@cs.au.dk](mailto:gregersen@cs.au.dk)

**Paper**

<https://cs.au.dk/~gregersen/papers/2020-esop-aneris-final.pdf>

**Coq development**

<https://github.com/logsem/aneris>