

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

Technical Appendix

LÉON 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

A Summary of the Specification of the Causally Consistent Distributed Database

Given client-provided *Addrlist* and *Keys*, we provide types, functions, predicates, and lemmas as in Section 3 (Mathematical Model), and given those, the specification of the causally-consistent distributed database is as follows. The specification existentially quantifies over abstract predicates; then follows all the laws governing the abstract predicates from Section 4 (Specification) and Section 7 (HOCAP Style Specification for Write) (here we only show a few) and finally the specification for the initialization operation.

$$\begin{aligned}
 & \exists \text{GlobalInv} : iProp. \\
 & \exists (\cdot) \rightarrow_s (\cdot) : Keys \rightarrow \mathcal{P}_{\text{fin}}(\text{WriteEvent}) \rightarrow iProp. \\
 & \exists (\cdot) \rightarrow_u (\cdot) : Keys \rightarrow \mathcal{P}_{\text{fin}}(\text{WriteEvent}) \rightarrow iProp. \\
 & \exists \text{Seen} : \mathbb{N} \rightarrow \text{LocalHistory} \rightarrow iProp. \\
 & \exists \text{Snap} : Keys \rightarrow \mathcal{P}_{\text{fin}}(\text{WriteEvent}) \rightarrow iProp. \\
 & \exists \text{initToken} : \mathbb{N} \rightarrow iProp. \\
 & \exists \Phi_{\text{DB}} : \text{Message} \rightarrow iProp. \\
 & \quad \text{Snap}(x, h) * \text{Snap}(x, h') \vdash \text{Snap}(x, h \cup h') \\
 & \wedge x \rightarrow_u h \vdash x \rightarrow_u h * \text{Snap}(x, h) \\
 & \wedge \dots \\
 & \wedge \boxed{\text{GlobalInv}}^{\text{N}_{\text{Gl}}} * \text{Seen}(i, s) * x \rightarrow_u h \vdash \models_{\mathcal{E}} \forall a \in s, w \in h. w.t < a.t \Rightarrow \exists a' \in s_{|x}. [a'] = w \\
 & \wedge \text{True} \vdash \models_{\mathcal{E}} \boxed{\text{GlobalInv}}^{\text{N}_{\text{Gl}}} * \left(\bigstar_{0 \leq i < \text{length}(\text{Addrlist})} \text{initToken}(i) \right) * \left(\bigstar_{k \in \text{Keys}} k \rightarrow_u \emptyset \right) * \text{initSpec}(\text{init})
 \end{aligned}$$

where

$$\begin{aligned}
 & \text{initSpec}(\text{init}) \triangleq \\
 & \left\{ \begin{array}{l} \text{initToken}(i) * \text{Fixed}(A) * \text{IsNode}(ip_i) * \text{Addrlist}[i] = (ip_i, p) * \\ (ip_i, p) \in A * \text{isList}(\text{Addrlist}, v) * \text{FreePorts}(ip_i, \{p\}) * \bigstar_{z \in \text{Addrlist}} z \models \Phi_{\text{DB}} \end{array} \right\} \\
 & \langle ip_i; \text{init}(i, v) \rangle \\
 & \{ (\text{rd}, \text{wr}). \text{Seen}(i, \emptyset) * \text{readSpec}(\text{rd}, i) * \text{writeSpec}(\text{wr}, i) \} \\
 & \text{readSpec}(\text{read}, i) \triangleq
 \end{aligned}$$

$$\begin{aligned}
& \{ \text{Seen}(i, s) \} \\
& \langle ip_i; \text{read}(x) \rangle \\
& \left. \begin{array}{l} v. \exists s' \supseteq s. \text{Seen}(i, s') * \\ \left((v = \text{None} \wedge s'_{|x} = \emptyset) \vee \right. \\ \left. (\exists a \in s'_{|x}. v = \text{Some}(a.v) * \text{Snap}(x, \{[a]\}) * a \in \text{Maximals}(s'_{|x}) * \text{Observe}(s'_{|x}) = a) \right) \end{array} \right\} \\
& \text{writeSpec}(\text{write}, i) \triangleq \\
& \forall \mathcal{E}, k, v, s, P, Q. \mathcal{N}_{\text{GI}} \subseteq \mathcal{E} \Rightarrow \\
& \square (\forall s', a. (s \subseteq s' * a \notin s' * a.k = k * a.v = v * P) \\
& \quad \top \stackrel{\mathcal{E}}{\Rightarrow} \forall h. \left([a] \notin h * [a] \in \text{Maximals}(h \uplus \{[a]\}) * k \rightarrow_s h * \right. \\
& \quad \left. \text{Seen}(i, s' \uplus \{a\}) * \text{Maximum}(s' \uplus \{a\}) = a \right) \\
& \quad \Rightarrow_{\mathcal{E} \setminus \mathcal{N}_{\text{GI}}} k \rightarrow_s h \uplus \{[a]\} * \mathcal{E} \stackrel{\top}{\Rightarrow} Q \ a \ h \ s' \rangle \\
& \{ P * \text{Seen}(i, s) \} \langle ip_i; \text{write}(k, v) \rangle \{ v.v = () * \exists h, s', a. s \subseteq s' * Q \ a \ h \ s' \}
\end{aligned}$$

B Client Reasoning about Causality: Indirect Causal Dependency

The proof of the indirect causal dependency example [Lloyd et al. 2011] makes use of predicates

$$\begin{aligned}
& \text{pending} : iProp \\
& \text{shot} : \text{WriteEvent} \rightarrow iProp \\
& \text{hist} : \mathcal{P}_{\text{fin}}(\text{WriteEvent}) \xrightarrow{\text{fin}} iProp
\end{aligned}$$

satisfying the laws below. The predicates pending and shot are defined using the *oneshot* resource algebra [Jung et al. 2016] whereas the hist predicate is defined using a fractional agreement resource algebra. We refer to the Coq formalization for all the details and the full proof.

Laws for ghost resources

$$\begin{aligned}
& \text{pending} \Rightarrow * \text{shot}(w) \\
& \text{shot}(w) * \text{shot}w' \vdash w = w' \\
& \text{pending} * \text{shot}(w) \vdash \text{False} \\
& \overline{[\diamond]} * \overline{[\diamond]} \vdash \text{False} \\
& \text{hist}(h_1) * \text{hist}(h_2) \vdash h_1 = h_2 \\
& \text{hist}(h_1) * \text{hist}(h_1) \vdash \text{hist}(h_2) * \text{hist}(h_2)
\end{aligned}$$

Invariants

$$\begin{aligned}
& \text{Inv}_x \triangleq \exists h. x \rightarrow_u h * \text{hist}(h) * \left((\text{pending} * \forall w \in h. w.v \neq 1) \vee \right. \\
& \quad \left. (\overline{[\diamond]} * \exists w. \text{shot}(w) * w.v = 37 * \text{Maximum}(h) = \text{Some}(w) * \right. \\
& \quad \left. \forall w' \in h. w'.v = 37 \Rightarrow w = w') \right) \\
& \text{Inv}_y \triangleq \exists h. y \rightarrow_u h * \forall w \in h. w.v = 1 \Rightarrow \exists w'. w'.t < w.t * \text{shot}(w')
\end{aligned}$$

Proof sketch*Node i, proof outline*

$$\begin{array}{l}
\{ \text{Seen}(i, s) * \text{hist}(h) * \boxed{\diamond} * h \subseteq \lfloor s \rfloor * \boxed{\text{Inv}_x} \} \\
\left. \begin{array}{l}
\{ \text{hist}(h) * \text{hist}(h) * x \rightarrow_u h * \text{pending} * \boxed{\diamond} \} \\
\text{write}(x, 0) \\
\left\{ \begin{array}{l}
\exists a, s' \supseteq s. \text{Seen}(i, s' \uplus \{a\}) * \text{hist}(h \uplus \{\lfloor a \rfloor\}) \\
* \text{hist}(h \uplus \{\lfloor a \rfloor\}) * x \rightarrow_u h \uplus \{\lfloor a \rfloor\} * \text{pending} * \boxed{\diamond}
\end{array} \right\}
\end{array} \right\} \\
\{ \exists h', s'. \text{Seen}(i, s') * \text{hist}(h) * \boxed{\diamond} * h' \subseteq s' * \boxed{\text{Inv}_x} \} \\
\{ \text{Seen}(i, s') * \text{hist}(h') * \boxed{\diamond} * h' \subseteq \lfloor s' \rfloor * \boxed{\text{Inv}_x} \} \\
\left. \begin{array}{l}
\{ \text{hist}(h) * \text{hist}(h) * x \rightarrow_u h * \text{pending} * \boxed{\diamond} \} \\
\text{write}(x, 37) \\
\left\{ \begin{array}{l}
\exists s'' \supseteq s', a'. \text{Seen}(i, s'' \uplus \{a'\}) * \text{hist}(h \uplus \{\lfloor a' \rfloor\}) * \text{hist}(h \uplus \{\lfloor a' \rfloor\}) * x \rightarrow_u h \uplus \{\lfloor a' \rfloor\} \\
* \boxed{\diamond} * \text{Maximum}(h \uplus \{\lfloor a' \rfloor\}) = \text{Some}(\lfloor a' \rfloor) * \text{shot}(\lfloor a' \rfloor)
\end{array} \right\}
\end{array} \right\} \\
\{ \exists h', s''. \text{Seen}(i, s'') * \text{hist}(h') * \boxed{\text{Inv}_x} \}
\end{array}$$

Node j, proof outline

$$\begin{array}{l}
\{ \text{Seen}(j, s) * \boxed{\text{Inv}_x} * \boxed{\text{Inv}_y} \} \\
\text{wait}(x = 37) \\
\{ \exists s' \supseteq s, a_x. \text{Seen}(j, s) * \text{shot}(\lfloor a_x \rfloor) * w_x \in s' * \dots \} \\
\{ \text{Seen}(j, s') * \text{shot}(w_x) \} \\
\left. \begin{array}{l}
\{ \exists h_y, s'. x \rightarrow_u h_y * \dots \} \\
\text{write}(y, 1) \\
\left\{ \begin{array}{l}
\exists a_y, s'' \supseteq s'. \text{Seen}(j, s'' \uplus \{a_y\}) * x \rightarrow_u h_y \uplus \{\lfloor a_y \rfloor\} * \text{shot}(w_x) * \\
a_y.v = 1 * a_y \in s'' * \lfloor a_x \rfloor.t < \lfloor a_y \rfloor.t
\end{array} \right\}
\end{array} \right\} \\
\{ \text{Seen}(j, s'') * \boxed{\text{Inv}_y} \}
\end{array}$$

Node k, proof outline

$$\begin{array}{l}
\{ \text{Seen}(k, s) * \boxed{\text{Inv}_x} * \boxed{\text{Inv}_y} * \} \\
\text{wait}(y = 1) \\
\{ \exists s' \subseteq s, a_y \in s', w_x. \text{Seen}(k, s') * \text{shot}(w_x) * w_x.t < \lfloor a_y \rfloor.t * a_y.v = 1 \} \\
\{ \text{Seen}(k, s') * \text{shot}(w_x) \} \\
\left. \begin{array}{l}
\{ \exists h_x. x \rightarrow_u h_x * \text{shot}(w_x) * \text{shot}(w_x) * \text{Maximum}(h_x) = \text{Some}(w_x) * w_x.v = 37 \} \\
\text{read}(x) \\
\{ v. \exists s''. \text{Seen}(k, s'') * v = \text{Some}(37) \}
\end{array} \right\} \\
\{ v. \exists s''. \text{Seen}(k, s'') * v = \text{Some}(37) \}
\end{array}$$

C Case Study: Towards Session Guarantees for Client-Centric Consistency

C.1 Session Manager Library Implementation

```

(* Client stub *)
type db_key
type db_val

type sm_req =
  InitReq
| ReadReq of db_key
| WriteReq of db_key * db_val

type sm_res =
  InitRes
| ReadRes of db_val
| WriteRes

let rec listen_wait_for_seqid skt seq_id =
  let res_raw = listen_wait skt in
  let res = deser_res (fst res_raw) in
  let tag = fst res in
  let vl = snd res in
  if (tag = !seq_id) then
    (seq_id := !seq_id + 1;
     vl)
  else
    listen_wait_for_seqid skt seq_id

let session_exec skt seq_id lock server_addr
  req =
  acquire lock;
  let msg = ser_req req in
  sendTo skt msg server_addr;
  let res = listen_wait_for_seqid skt seq_id
    in
  release lock;
  res

(* Request handler *)
let rec request_handler skt rd_fn wr_fn =
  let req_raw = listen_wait skt in
  let sender = snd req_raw in
  let req = deser_req (fst req_raw) in
  let seq_id = fst req in
  let res =
    match (snd req) with
    | Some InitReq → InitRes
    | Some ReadReq k → ReadRes (rd_fn k)
    | Some WriteReq (k, v) → wr_fn k v;
    WriteRes
    | None → assert false
  in
  sendTo skt (ser_res (seq_id, res)) sender;
  request_handler skt rd_fn wr_fn

let server dbs db_id req_addr =
  let fns = init dbs db_id in
  let rd_fn = fst fns in
  let wr_fn = snd fns in
  let skt = socket () in
  socketbind skt req_addr;
  request_handler skt rd_fn wr_fn

let sm_setup client_addr =
  let skt = socket () in
  socketbind skt client_addr;
  let seq_id = ref 0 in
  let lock = newlock () in
  let connect_fn server_addr =
    session_exec skt seq_id lock server_addr
      InitReq;
    ()
  in
  let read_fn server_addr key =
    session_exec skt seq_id lock server_addr
      (ReadReq key)
  in
  let write_fn server_addr key vl =
    session_exec skt seq_id lock server_addr
      (WriteReq (key, vl));
    ()
  in
  (connect_fn, read_fn, write_fn)

```

C.2 Session Manager Specifications

SM-INIT

$$\{\top\} \langle ip_{client}; sconnect(ip_i) \rangle \left\{ \exists s. \text{Seen}(i, s) * \bigstar_{k \in \text{Keys}} \exists h_k. \text{Snap}(k, h_k) * \boxed{\text{GlobalInv}}^{\mathcal{N}_{\text{GI}}} \right\}$$

SM-READ

$$\{ip_i \Rightarrow \Phi_i * \text{Seen}(i, s) * \text{Snap}(k, h)\}$$

$$\langle ip_{client}; sread(ip_i, k) \rangle$$

$$\left\{ \begin{array}{l} v. \exists s' \supseteq s, h' \supseteq h. * \text{Seen}(i, s') * \text{Snap}(k, h') * \\ (v = \text{None} * s'|_k = \emptyset) \\ \vee (\exists a, w. v = \text{Some}(w) * a.v = w * a.k = k * a \in \text{Maximals}(s'|_k) * [a] \in h') \end{array} \right\}$$

SM-WRITE

$$\{ip_i \Rightarrow \Phi_i * \text{Seen}(i, s) * \text{Snap}(k, h)\}$$

$$\langle ip_{client}; swrite(ip_i, k, v) \rangle$$

$$\left\{ \begin{array}{l} \exists a, s' \subseteq s, h' \subseteq h. a.k = k * a.v = v * \text{Seen}(i, s') * \text{Snap}(k, h') \\ * a \notin s * a' \in s' * [a] \notin h * [a'] \in h' * [a] \in \text{Maximals}(h') * \text{Maximum}(s') = \text{Some}(a) \end{array} \right\}$$

C.3 Session Guarantees Specifications

SM-READ-YOUR-WRITES

$$\{ip_i \Rightarrow \Phi_i\}$$

$$\langle ip_{client}; sconnect(ip_i); swrite(ip_i, k, v_w); sread(ip_i, k) \rangle$$

$$\left\{ \begin{array}{l} \exists s, a_w, a_r, v_r. v_o = \text{Some}(v_r) * a_w.k = k * a_w.v = v_w * a_r.k = k * a_r.v = v_r \\ v_o. * \text{Seen}(a, s) * a_w, a_r \in s * \neg(a_r.t < a_w.t) \end{array} \right\}$$

SM-MONOTONIC-READS

$$\{ip_i \Rightarrow \Phi_i\}$$

$$\langle ip_{client}; sconnect(ip_i); \text{let } v_1 = \text{sread}(ip_i, k) \text{ in let } v_2 = \text{sread}(ip_i, k) \text{ in } (v_1, v_2) \rangle$$

$$\left\{ \begin{array}{l} \exists v_{o1}, v_{o2}. v_o = (v_{o1}, v_{o2}) \\ \left(\begin{array}{l} (\exists s. v_{o1} = \text{None} * v_{o2} = \text{None} * \text{Seen}(a_1, s) * s|_k = \emptyset) \\ \vee \\ (\exists s, v_2, a_2. v_{o1} = \text{None} * v_{o2} = \text{Some}(v_2) * \text{Seen}(a_1, s) \\ * a_2.k = k * a_2.v = v_2 * a_2 \in \text{Maximals}(s|_k)) \\ \vee \\ (\exists s, v_1, v_2, a_1, a_2. v_{o1} = \text{Some}(v_1) * v_{o2} = \text{Some}(v_2) * \text{Seen}(a_1, s) \\ * a_1.k = k * a_1.v = v_1 * a_2.k = k * a_2.v = v_2 * a_2 \in \text{Maximals}(s|_k) \\ * \neg(a_2.t < a_1.t)) \end{array} \right) \end{array} \right\}$$

SM-MONOTONIC-WRITES

 $\{ip_i \models \Phi_i\}$ $\langle ip_{client}; sconnect(ip_i); swrite(ip_i, k_1, v_1); swrite(ip_i, k_2, v_2) \rangle$

$$\left(\begin{array}{l} \exists s_1, a_1, a_2. a_1.k = k_1 * a_1.v = v_1 * a_2.k = k_2 * a_2.v = v_2 \\ * \text{Seen}(a_1, s_1) * a_1, a_2 \in s_1 * a_1.t < a_2.t \\ * (\forall a, s_2, a_2. \text{Seen}(a_2, s_2) * a \in s_2 * a_2.t \leq e.t \\ \quad \equiv *_T \exists a'_1, a'_2. \lfloor a'_1 \rfloor = \lfloor a_1 \rfloor * \lfloor a'_2 \rfloor = \lfloor a_2 \rfloor * a'_1, a'_2 \in s_2 * a'_1.t < a'_2.t) \end{array} \right)$$

SM-WRITES-FOLLOW-READS

 $\{ip_i \models \Phi_i\}$ $\langle ip_{client}; sconnect(ip_i); \text{let } v = \text{sread}(ip_i, k_r) \text{ in } swrite(ip_i, k_w, v_w); v \rangle$

$$\left(\begin{array}{l} \exists s_1, a_w. a_w.k = k_w * a_w.v = v_w * \text{Seen}(a_1, s_1) * a_w \in s_1 \\ v_o. \left(\begin{array}{l} v_o = \text{None} \\ \vee \\ \exists a_r, v_r. v_o = \text{Some}(v_r) * a_r.k = k_r * a_r.v = v_r * a_r \in s_1 * a_r.t < a_w.t \\ * (\forall a, s_2, a_2. \text{Seen}(a_2, s_2) * a \in s_2 * a_w.t \leq e.t \\ \quad \equiv *_T \exists a'_r, a'_w. \lfloor a'_r \rfloor = \lfloor a_r \rfloor * \lfloor a'_w \rfloor = \lfloor a_w \rfloor * a'_r, a'_w \in s_2 * a'_r.t < a'_w.t) \end{array} \right) \end{array} \right)$$

References

- Ralf Jung, Robbert Krebbers, Lars Birkedal, and Derek Dreyer. 2016. Higher-order ghost state. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming, ICFP 2016, Nara, Japan, September 18-22, 2016*. 256–269. <https://doi.org/10.1145/2951913.2951943>
- Wyatt Lloyd, Michael J. Freedman, Michael Kaminsky, and David G. Andersen. 2011. Don't settle for eventual: scalable causal consistency for wide-area storage with COPS. In *Proceedings of the 23rd ACM Symposium on Operating Systems Principles 2011, SOSP 2011, Cascais, Portugal, October 23-26, 2011*. 401–416. <https://doi.org/10.1145/2043556.2043593>