

# LINEARIZABILITY

PAST WEEKS

- TRACES
- CORRECTNESS → CONSTRAIN ADMISSIBLE TRACES
  - SOUNDNESS
  - LIVENESS
- RPC:
  - REQUEST  $\leftrightarrow$  RESPONSE
- TODAY ◦ LINEARIZABILITY, A SPECIFIC CORRECTNESS CONDITION

LINEARIZABILITY IS REALLY ABOUT CONCURRENCY

↳ "How Should A **DATA TYPE** BEHAVE WITH CONCURRENT ACCESSSES"

◦ DOES NOT CONSIDER FAILURES, LOSSES, ETC.

◦ WHY INTERESTING?

ABSTRACT DATA TYPE

SOME OBJECT    **QUEUE**    **STACK**    ...

WITH OPERATIONS    **ENQUEUE,**    **PUSH**    ...  
                                  **DEQUEUE**    **POP**

AND A SEQUENTIAL SPECIFICATION

BEHAVIOR WHEN A SINGLE PROCESS ISSUES ONE  
REQUEST AT A TIME.

**QUEUES** ◦ FIFO

p<sub>0</sub> q.enqueue(1)  
                  OK ←

**STACK** ◦ LIFO

p<sub>0</sub> s.push(1)  
                  OK ←

P<sub>0</sub> q.enqueue(2) }  
OK ←  
P<sub>0</sub> q.dequeue() }  
OK(1) ←

P<sub>0</sub> s.push(2) }  
OK ←  
P<sub>0</sub> s.pop() }  
OK(2) ←

---

P<sub>0</sub> q.dequeue() }  
←

---

P<sub>0</sub> s.pop() }  
←

### MAIN TAKEAWAY

TELLS US BEHAVIOR WHEN A SINGLE  
PROCESS PERFORMS ONE OPERATION  
AT A TIME

### CONSISTENCY IN THIS LECTURE


GIVEN A PARTIALLY ORDERED SET OF  
CONCURRENT OPERATIONS  
SPECIFY HOW ADT BEHAVES

# ↳ EQUIVALENT SEQUENTIAL HISTORY

## LINEARIZABILITY

PARTIAL ORDER  $<_H$

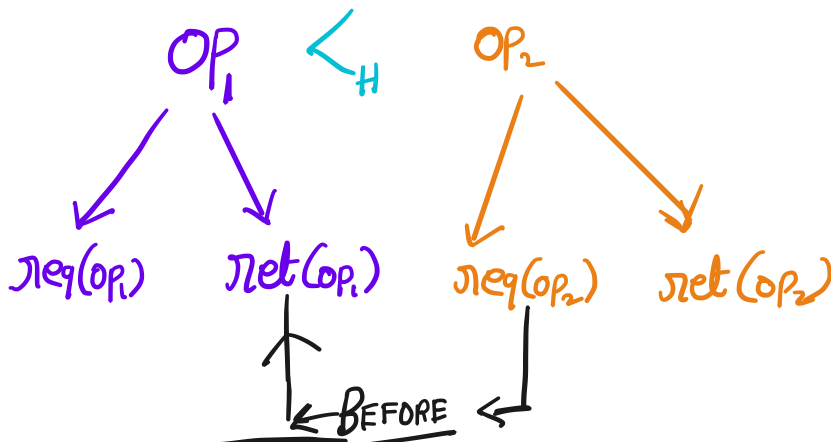
EQUIVALENT SEQ. HISTORY  $S \supseteq$




↳ CORRECT SEQ. SPECIFICATION

Here lie errors  
& dragons.  
Seda, Herlihy, Pentlak  
Proc '21


$<_H$  : Partial order on operations



↳ IN WHOSE FRAME OF REFERENCE / ACCORDING TO WHOM?



→ Might really be a group of processes rather than one.



Some object  $O$  is linearizable implementation of ADT

$\Rightarrow$

Given  $H, \leq_H$  can extend  $H$  to  $H'$  so

$\exists$  sequential history  $S(\leq_S)$  compliant with ADT  
s.t.

(1)  $\text{complete}(H')$  is equivalent to  $S$

① complete(H)

$$\textcircled{2} \quad \langle H \subseteq \langle S$$

Walk through this step by step

① complete(H): Maximal subsequence of H where every req has a response

$$H: \langle req, op_1, p_0 \rangle$$

complete(H) =

$$H: \langle req, op_1, p_0 \rangle \langle ret, op_1, p_0 \rangle$$

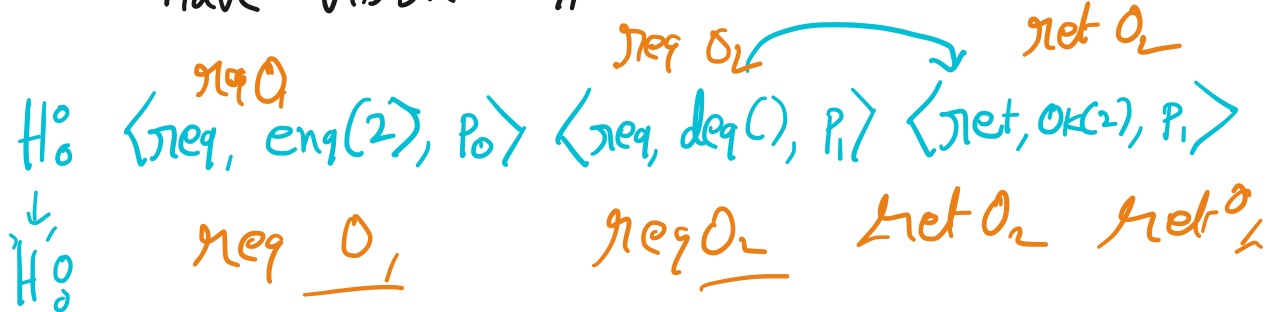
complete(H) =

② Extend H to H'

↳ Can add missing responses/rets.

→ Necessary to deal with operations that

have visible effect before response




③ Find S, < S

∴ S 3

$S$   $O_1$   $O_2$   
 $\prec_S$

Q1: Is  $S$  UNIQUE GIVEN  $H, \prec_H$ ?



$D, H, \prec_H \rightarrow S$

Q2 We produce  $S, \prec_S$  equivalent to  $complete(H')$ , but only require  $\prec_H \subseteq \prec_S$ . Is that OK?

Core problem:  $\prec_H$  only orders "operations" in  $H$

"Proof by citation :)"

Both invocation ( $\pi_{req}$ )  
 & response ( $\pi_{ret}$ )  
 in  $H$

As written some operations can be reordered

$inv\ q.enq(1) \prec_H ok(1) \prec_H inv\ q.enq(2) \prec_H ok(2)$

$$H_0 \quad \langle \underbrace{\text{req}, \text{req}(i), p_0}_{\text{req}(op_i)}, \underbrace{\text{ret}, \text{OK}(i), p_0}_{\text{ret}(op_i)}, \underbrace{\langle \text{req}, \text{req}(i), p_0 \rangle}_{\text{req}(op_2)} \rangle$$

$$H'_0 \quad \text{req}(op_i) \text{ ret}(op_i) \text{ req}(op_2) \text{ ret}(op_2)$$

$$\langle_H : \{ \}$$

$$\langle_{H'} : \{ (op_i, op_2) \}$$

Allows  
locality to  
be broken

$\left\{ \begin{array}{l} S \\ O \end{array} \right.$

$$\{ [op_2; op_i] \}$$

$$\langle_H \subset \langle_S$$

$$\langle_{H'} \not\subset \langle_S$$

Herlihy's connection [PODC 2021]

Should be

$$\langle_{\text{complete}(H')} \subseteq \langle_S$$

$$\text{ctx.set}(2) \quad \text{OK}(1); \quad \text{ctx.set}(1)$$

LOCALITY

• If a system has multiple objects  $o_1, o_2, \dots, o_n$   
the system as a whole is linearizable

$$\Leftrightarrow o_1, o_2, \dots, o_n \text{ linearizable}$$

$$H \text{ linearizable} \Leftrightarrow H|_{o_i} \text{ linearizable}$$



↑ Only operations on  $O_i$

$H$  linearizable  $\Rightarrow H|O_i$  linearizable

"Easy direction"

- Observation

$H, H', \text{complete}(H')$   
 $\sqcup \sqcup \sqcup$   
 $H|O_i, H'|O_i, \text{complete}(H'|O_i)$

$\angle_{H|O_i} \subseteq \angle_H$

$\Rightarrow \angle_{H|O_i} \subseteq \angle_S$  & easy to see  $\angle_{H|O_i} \subseteq \angle_{S|O_i}$

Also  $S$  equivalent to  $\text{complete}(H') \Rightarrow$

$S|O_i$  equivalent to  $\text{complete}(H'|O_i)$

$S$  compliant with  $O_1, \dots, O_n \Rightarrow$

$S|O_i$  compliant with  $O_i$

$H|O_i$  linearizable  $\forall O_i \Rightarrow H$  linearizable

• More complicated. Why care?

• Instructive to see what it means when this is not true.

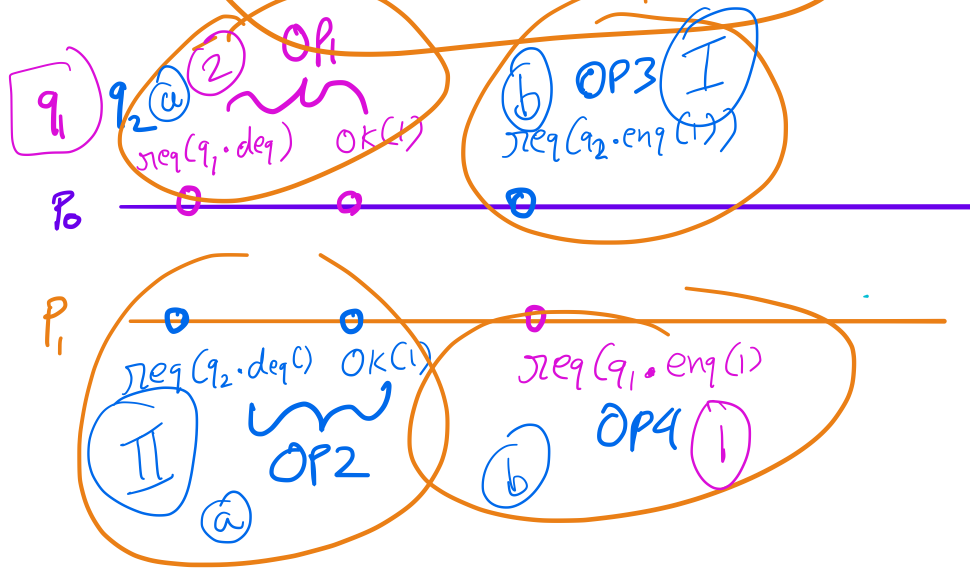
# Two examples

① Broken linearizability definition  
 ①  $\prec_H$  only defined on operations with request & response

② Require  $\prec_H \subseteq \prec$

$H|Q_i$

$H|Q_i \text{ lin} \Rightarrow$   
 $H \text{ lin}$   
 $\times$



$\prec_H : \{ \}$

[Note: Sequential history requires maintaining process order  $\Rightarrow OP_3$  req must be after  $OP_2$  req]

$H|Q_1$  can be linearized


$OP_4 ; OP_1$  [Same way as before]

$H|Q_2$  can be linearized

$OP_3 ; OP_2$

$H$  cannot be linearized

It cannot



Why is this not a problem with the fix?

② Sequential Consistency


- Different model.

For process  $P_0, P_1, P_2, \dots, P_n$  requires ordering

$$S_i <_s \text{ s.t.}$$

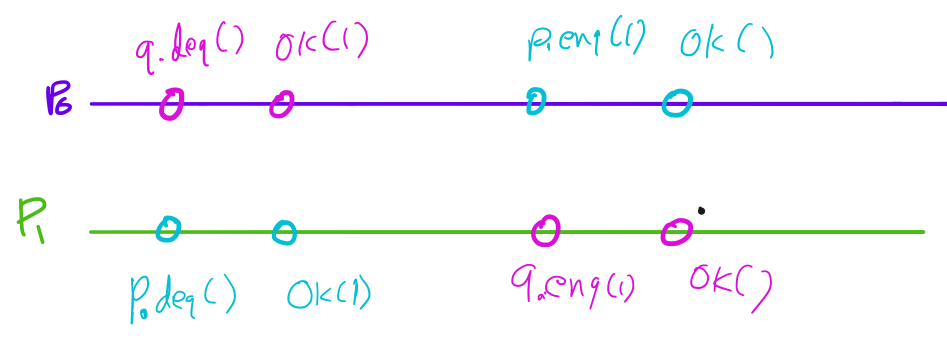
$$\prec_{P_i} \subseteq \prec_s$$

"No real-time ordering required"



Seq Cst ✓

Linearizability ✗



$H|_q$  is seq cst (see above)  
 $H|_p$  " " " (symmetry)  
 $H$  is **NOT** Seq Cst

$H|_O$ : linearizable  $\forall O_i \Rightarrow H$  linearizable

Sketch

$\forall O_i$ :  $H|_O$ : linearizable

$\Rightarrow$  Given  $H|_O$ : can find  $H'_i = H|_O + R_i$   
 s.t.  $\exists$  sequential history  $S_i$  where

$S_i$  equivalent to  $H'_i$

$$\Delta \quad \langle_{\text{complete}(H'_i)} \subseteq \langle_{S_i}$$

Proof by construction

$$H' = H + \sum_i R_i$$

$$\langle_{H'} = \langle_H \cup (\cup_i \langle_{H'_i}) \cup \text{Transitive edges}$$

$$\text{if } e_1 \langle_{H'} e_2 \wedge e_2 \langle_{H'} e_3 \text{ then } e_1 \langle_{H'} e_3$$

$\langle_{H'}$  IS A PARTIAL ORDER

$\hookrightarrow$  IRREFLEXIVE ( $(e, e) \notin \langle_{H'}$ ) By construction

$\rightarrow$  Transitive By CONSTRUCTION

$(a_1, a_2) \in S$

$(a_1, a_2) \in S$

= Antisymmetric

$(a, b) (b, c) \dots \Rightarrow$  Antisymmetric  
 $\hookrightarrow$  No cycles s.t.

$\Rightarrow (a, c) \in S$

$(a, b) (b, c) (a, c)$   
 $(c, a)$

$$e_1 <_{H'} e_2 <_{H'} e_3 \dots <_{H'} e_n$$

Proof by contradiction

ASSUME Otherwise

(I)

$e_1, e_2, \dots, e_n$  are not all operations on the same object

$<_{H'}$  partial order  $\neq$

(II)

PICK SMALLEST CYCLE s.t.  $e_1, e_2$  are ops on  $o_1 \neq o_2$

Lemma None of  $e_2, \dots, e_n$  are ops on  $o_1$   
 $\hookrightarrow$  Assume otherwise,  $e_i$  is first op after  $e_1$  on  $o_1$

$$e_2 <_{H'} e_i \Rightarrow \text{Ret}(e_2) \text{ Before Inv}(e_i)$$

$$e_1 <_{H'} e_2 \Rightarrow \text{Ret}(e_1) \text{ Before Inv}(e_2)$$

$\Rightarrow$  By construction

$$e_1 <_{H'} e_i$$

$\Rightarrow e_1, e_i, e_{i+1}, \dots, e_n$  is

Smaller cycle  
Contradiction

$\Rightarrow e_1 \& e_n$  are ops on different objects

$\Rightarrow e_1 <_H e_2 ; e_n <_H e_1$

But  $<_H$  is transitive

$\Rightarrow e_n <_H e_2$

$\Rightarrow e_2, \dots, e_n$  is smaller cycle.

No cycle exists


CLAIM Can construct sequential history  $S$   
consistent with  $complete(H')$   
where  $<_{complete(H')} \subseteq <_S$

Non-Blocking

$\hookrightarrow$  Why useful?



$<_I: Op_1 <_I Op_2 \iff inv(Op_1) \text{ before } inv(Op_2)$



→ Why true

# CAP theorem

