

Distributed Systems :- More Linearizability

Where we are

- Consistency models

↳ What order of operations
can be observed on a concurrent
data structure

(Given by an ADT with
a seq. spec.)

- History : $\langle \text{inv} \dots \rangle \langle \text{res} \dots \rangle$

- Linearizability :

An implementation is linearizable

\Leftrightarrow

any history H produced by impl.

can be extended to H' s.t

we can find a sequential history

S where $\langle \text{complete}(H') \rangle \subseteq \langle S \rangle$

- Linearizability - B

$\langle H \rangle \subseteq \langle S \rangle$

◦ Locality


H Linearizable $\Leftrightarrow H/O$ linearizable
 (or other consistency that provides locality)
 for all objects o appearing in H

◦ Non-blocking

H linearizable $\Leftrightarrow H$ contains $\langle inv\ o\ op \rangle$ w/o $\langle res\ o\ v \rangle \Rightarrow$ we can find v s.t. $H \circ \langle res\ o\ v \rangle$ is linearizable.

◦ Linearizable-B is not local & not non-blocking.

\hookrightarrow Saw two counter examples last class



H/O_1 : Lin-B ✓ } NOT
 H/O_2 : Lin-B ✓ } LOCAL
 H : Lin-B X

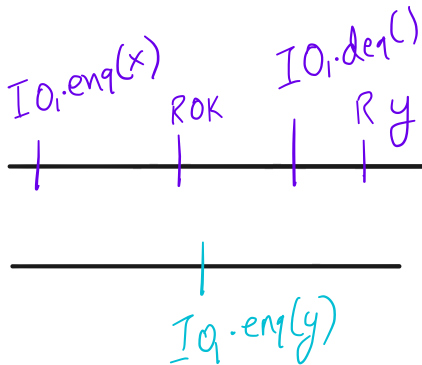
H deq R_x

H/O_1 : Lin X

eng(x)

HIO₂: Lin X

H' deg R_x ↓ OK



H: Lin-B ✓

H₀ <R OK>: Lin-BX

} Not non-blocking

H₀: Lin ?

Today

- A couple of other consistency models

↳ Seq Cst

→ Strict Serializability

Comparing the strength of consistency models.

- Building linearizable "things"

- CAP

Sequential Consistency / Seq Cst

- Often the strongest memory model supported by language specs:
C++ `std::memory_order::seq_cst`
Rust `Ordering::SeqCst`

- Guarantees

- Can find a total order for all processes
- Total order does not violate process order for any process.

Translating to our terminology

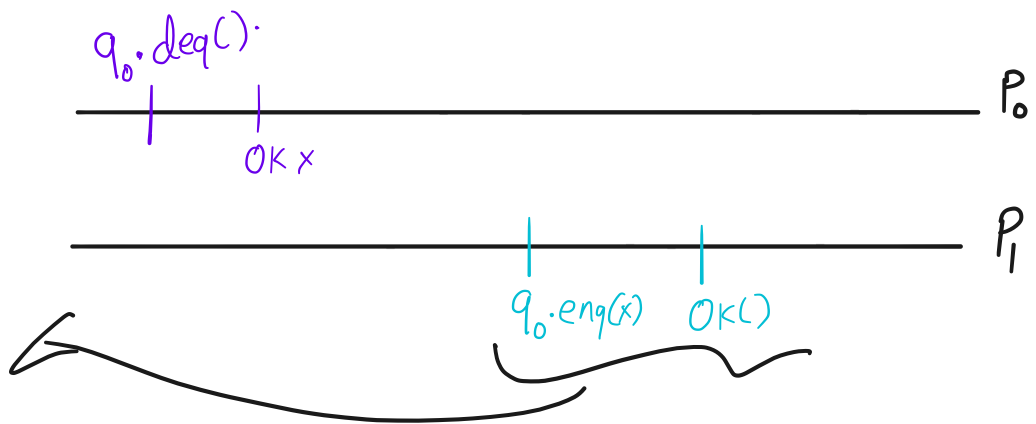
H: History. [But \prec_H is unused]

For any process $p \in \Pi$, \prec_p : order of operations.

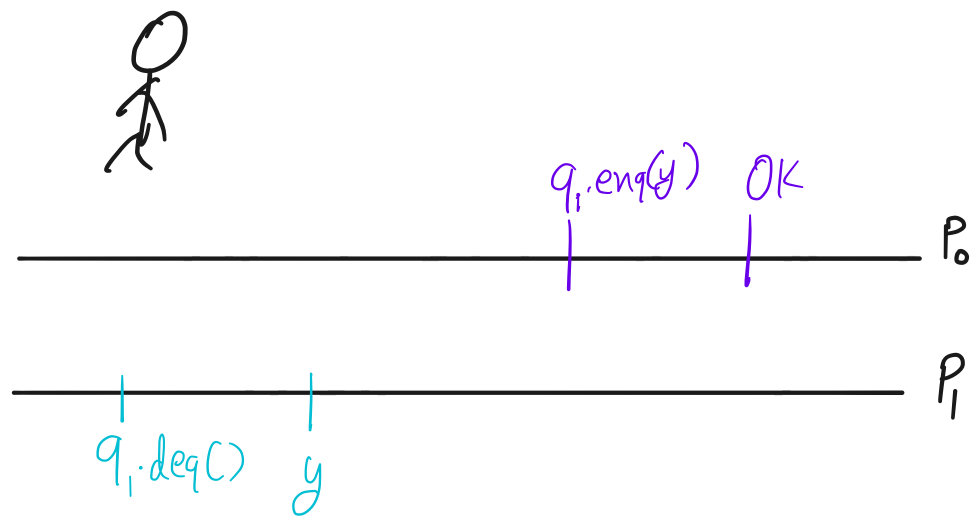
$H \upharpoonright_p \prec_p \equiv \prec_{H \upharpoonright_p}$

[Remember, we assumed processes are sequential]


$$H \rightarrow S \text{ s.t. } \prec_p \subseteq \prec_S \forall p \in \Pi$$




Seq cst ? ✓
 Linearizable? X



Seq cst ? ✓
 Linearizable? X






Seq Cst? $\leq R$

Is Seq Cst local? No

Strict Serializability

Transaction

- ↳ Sequence of operations issued by one process
- Operations might target different objects
- Must be ISOLATED



History strict serializable

- ↳ History linearizable
- s.t. in linearization **S**


cannot find O_1, O_2, O_3

s.t. \nexists $T_{KN}(O_1) = T_{KN}(O_3) \neq T_{KN}(O_2)$


$O_1 <_S O_2 \wedge O_2 <_S O_3 \wedge$

$O_1 <_S O_3$

Non-blocking?



Local?



Observe


Linearizable \Rightarrow Seq Cst

Why? - Processes are sequential

\Rightarrow For any history H , process P

$$\langle_p \subseteq \langle_H \subseteq \langle_S$$

LINEARIZABILITY



STRICT SERIALIZABILITY

if \dots

Notion of stronger/weaker consistency models

Implications for the readings?

RIFL

BUILDING LINEARIZABLE SYSTEMS

Why?

- The main reason we read the Herlihy & Wing paper.

- Argument: - Linearizability is easier for programmers to reason about
 - ↳ When do effects become visible?

Two Ends of the Spectrum

- Assume you are given a linearizable impl. of an ADT
 - Queue
 - Register - get } get: Returns value of last set (or \perp if no set has been called)
 - set }

Sequential Spec

| | |
|---|--|
| { | get() $\rightarrow \perp$ ✓ |
| | set(5) set(6) get() $\rightarrow 6$ ✓ |
| | set(5) set(6) set(7) get() $\rightarrow 6$ ✗ |


Types of systems

- Collection of independent lin. impl. of ADT

FOR EXAMPLE, KV Store

(No txns, only single key operations,

roughly what most people think
Redis provides.)



Sequential spec.

get(k) returns v from

last put(k, v) (or \perp)

Observe: Sequential spec. does not impose
any ordering requirement b/w keys

Just use a set of lin. register

Relation to locality

- Collection of dependent ADTs

Counter

- Init value \emptyset

- Inc \rightarrow increment value, return OK

- Dec \rightarrow

decrement value, return OK

- Get \rightarrow current value

Maybe track value in a lin. register

(get, set)

val

init:

val.set(0)

get:

return val.get()

inc:

c ← val.get()

val.set(c+1)

dec:

c ← val.get()

val.set(c-1)

Does this work?

Can we fix?

Practice

- One common approach in distributed system

- Maintain a totally ordered log
of operations

} R
S
M

- Apply operations in logical order)

- Why?

- Need to replicate data for F.T.

- Need to ensure replicas agree on order even for CONCURRENT OPS

[Failures cannot break abstraction]

CRITICALLY ANALYZING RIFL

- Assumptions

RPCs; Runtime

- System — Durable Storage & Rel. md

- Client

- Importance of exactly once to linearizability?

Does exactly once delivery \Rightarrow Linearizability?

Does linearizability \Rightarrow exactly once delivery?

- What does RIFL Provide?

CAP:

- Why?

- Need to replicate data for F.T.

- Need to ensure replicas agree on order
even for CONCURRENT OPS

[Failures cannot break abstraction]

Q Can replicas agree on order of concurrent ops without communicating?



Problem: Remember R0 cannot distinguish between - R1 has failed

- Messages from R1 are } Partition
delayed (indep.)

In the absence of messages from R1, R0 has two options

- Wait "Not available"

Processing might take unbounded time

- Process assuming R1 has failed

- process, assuming

R_0 & R_1 might disagree
on results

But R_0, R_1 are impl- details

↳ Violate linearizability

"Not consistent"