

RPCs & Linearizability

RPCs

- Why?

- The most common way that processes in implemented distributed systems interact

↳ Convenient way to define

① The structure of messages

- CONTENTS

- ENCODING

② HIDE NETWORKING DETAILS

- Some of the papers we read describe protocols in terms of RPC calls & responses

NOTE: The model in the paper & what we talk about closely corresponds to what GRPC, THRIFT, etc. provide

```

on recv (echo, m)
  from s :
  send s, m

```

send a, (echo, "hello")
 recv "hello" from a



```

msg echo(msg m) {
  return m;
}

```

ret = a.echo("hello")
 assert (ret == "hello")

SEEMS SIMPLE ENOUGH, BUT THERE ARE DIFFERENCES

① RE-ENTRANCY / ATOMICITY

on recv n from y:

```

a = a + n
b = 2 * a - 1
send(z, (up-a, a))
send(z, (up-b, b))
send(y, a + b)

```

both
 add-and-
 ret(1)

alice add-and-ret(s)

```

int add-and-ret(int n) {
  a = a + n;
  b = 2 * a - 1;
  z.up-a(a);
  z.up-b(b);
  return a + b;
}

```

Most of the protocols we encounter that are expressed in messages

- Assume 1 message processed at

ASSUME \perp busy
a time
→ ASSUME all state updates & accesses happen atomically.

RPC: Making a call \Rightarrow

Caller must wait (Block) for response

Runtime (usually) can process other calls while blocked.

↳ Handling other RPC calls was mandated by Birrell & Nelson

X
int $\pi A(int)$ {

...
int z = $\gamma \cdot \pi B(\dots)$;

...
}

}

int $\pi C(\dots)$ {

...

Y
int $\pi B(\dots)$ {

...
int h = $\chi \cdot \pi C(\dots)$;

...
}

}

Allowing concurrent RPC calls helps with

local reasoning

↳ Do not need to consider implementation of function being called.

② EFFECT OF REMOTE FAILURES

X
int πA(int) {

...
int z = Y.πB(...);
...
①

~~Y
int πB(...)~~ {

~~...
int h = X.πC(...);
...
②~~

Y crashes before returning to X
responding

Message model: X never receives from Y

- on recv... logic never executed

RPCs: Implementation dependent but

↳ Call will timeout &

fail
↳ Return value indicates failure
(of call not process)

→ Exception thrown

→ ...

Will show up in papers as


RPC request fails OR

NO RESPONSE IS RECEIVED

LINEARIZABILITY

Motivation/Goal:

One reason to build Distributed Systems
is to tolerate faults



behavior map

Q: How does program behavior
to dist. program behavior?

Differences

- Concurrency
- Failures

Linearizability: Consistency model

Given ADT: Queue
Stack
...

How specified: Sequential specification

Queue: FIFO container

enq(5)

deq() → 5

enq(6)

enq(7)

enq(8)


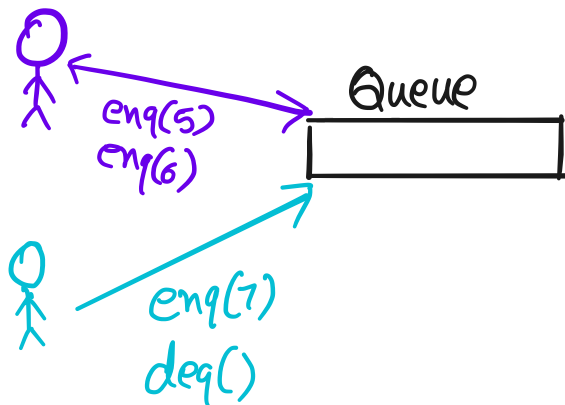
deq() → 6

deq() → 7

Stack: LIFO

push(1)
push(2)
pop()
...

Q: How does a concurrent queue or stack behave?



Many possibilities


- (a) Eve's observed order
- (b) Alice - Beth - Alice - Beth - ...
- (c) LINEARIZABILITY
- (d) SEQ CST
- (e) CAUSAL CONSISTENCY
- ...

Defining Linearizability

① ALICE ENQ(5) \rightarrow \langle ALICE Inv ENQ(5) \rangle
 \langle ALICE Res OK() \rangle

② History :- Sequence of inv, res interactions

Q. From whose perspective?



\prec_H : History partial order

$$C_1 \prec_H C_2 \Rightarrow Res C_1 < Inv C_2$$

[A Eng (5)]
①

②
[B Eng (7)]

[A Eng (6)]

[B Req (7)]

Some other useful definitions

H equivalent to H'

$$\Rightarrow H|_p = H'|_p \quad \forall p \in \Pi$$

Complete (H) \rightarrow Maximal subsequence
where every request has
a response.

H sequential \Rightarrow First event is invocation
 \triangleright

No concurrent requests ($<_H$ is total order)

Assume $H|p$ sequential $\forall p \in \Pi$

Defn

History H is linearizable if

H can be EXTENDED to H' s.t

① \exists Sequential history S , where

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

;
② $<_H \subseteq <_S$ \nrightarrow ③ S valid for ADT Seq. spec

Note, ② is wrong: Sela, Herlihy, Petrank
PODC '21 [Reading for next week]

Should be

$<_{\text{complete}(H')} \subseteq <_S$

PROPERTIES

① Locality: o_1, o_2, \dots, o_n objects

Linearizable $H \iff$

$H \upharpoonright_{o_1}, H \upharpoonright_{o_2}, \dots, H \upharpoonright_{o_n}$
linearizable

② Non-Blocking.

$S \iff H$ linearizable, contains
invocation $\langle A \text{ Inv Deq}(C) \rangle$

\exists response R s.t.

$H \cdot \langle A \text{ Inv Deq}(C) \rangle \langle A \text{ Res } R \rangle$

is linearizable.

Instructive to see examples where this is not

true



$I_{O_1} \text{ deq}(C)$

$R \text{ OK}(C)$

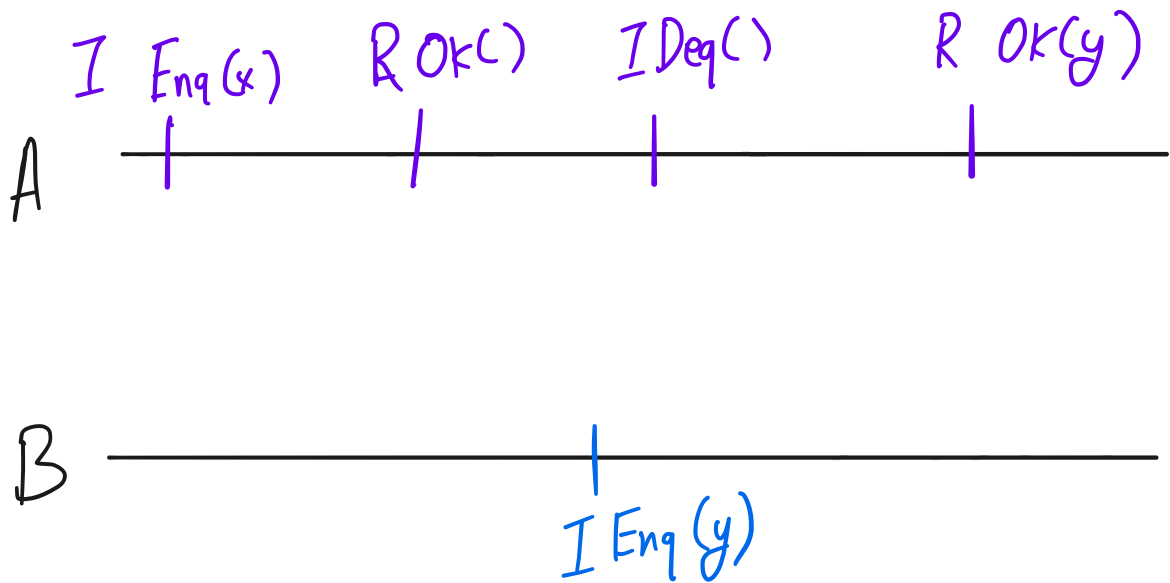
$I_{O_2} \text{ enq}(2)$




$R \text{ OK}(e)$





H/O₂





$O_2 \cdot \text{deg}$ OP_3

OP_4

$$\text{Complete}(H') = t|'$$

$$S|A = H'|A$$



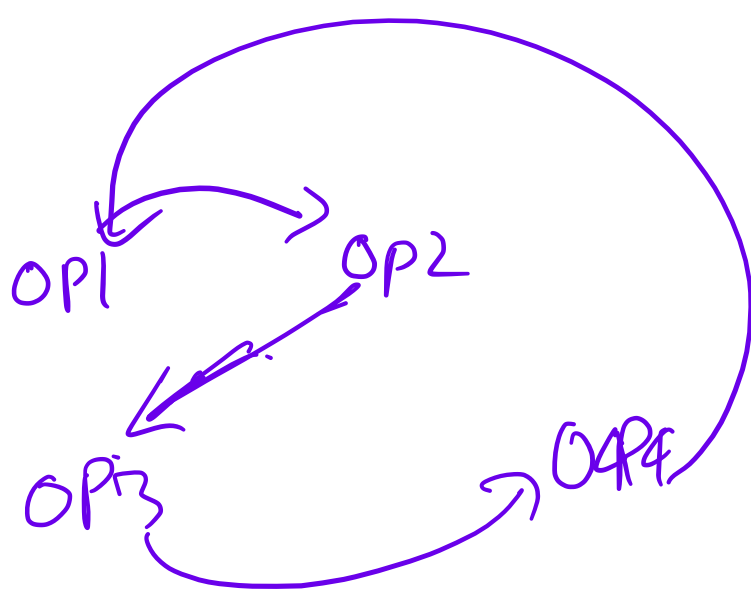
$$S|B = H'|B$$

$$OP_1 < OP_2$$

$$OP_3 < OP_4$$

$$OP_2 < OP_3$$

$$OP_4 < OP_1$$

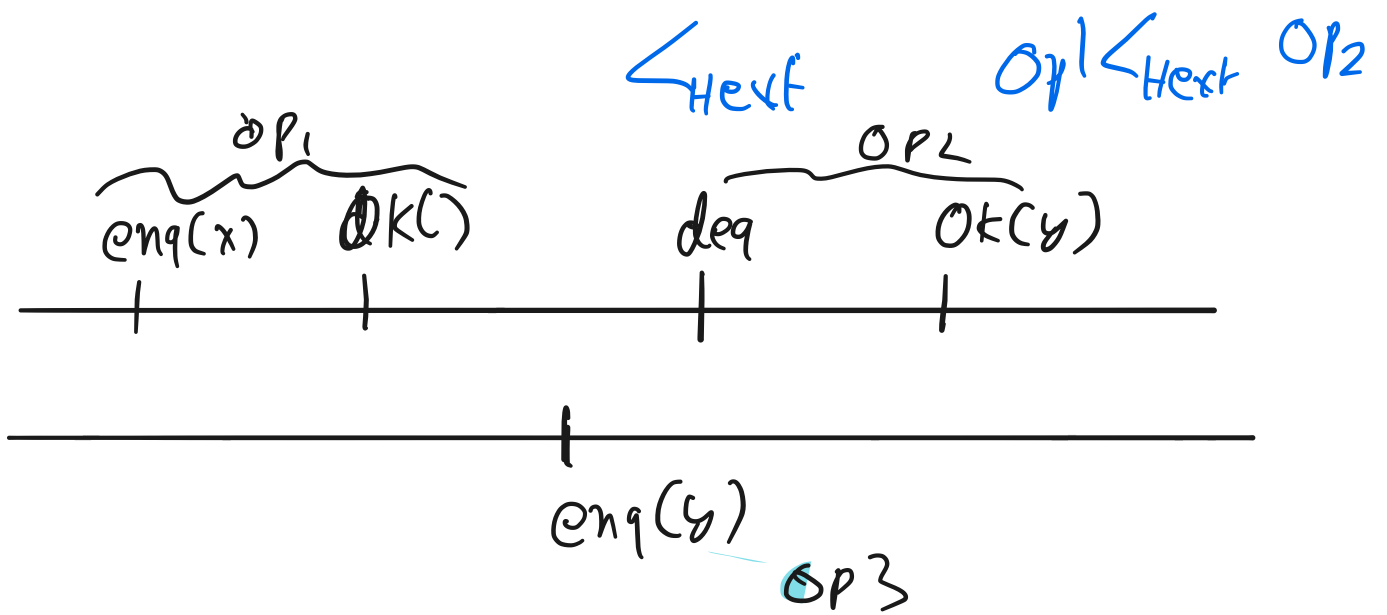


OP1

OP1 OP2 OP3

$$H_{ext} = H \cdot \langle R OK(C) \rangle$$

H



$$\langle_H : \{ (OP_1, OP_2) \}$$

H buggy-line

$$H \rightarrow H'$$

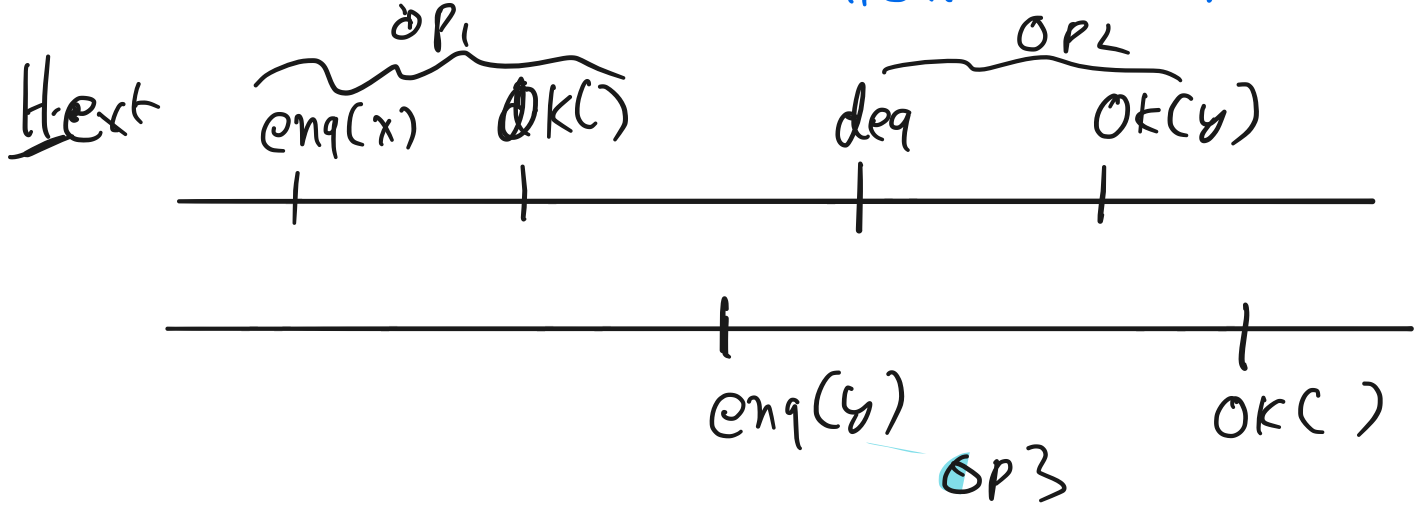
$$H \cdot \{ \cancel{OK(C)} \}$$

$$\text{Complete}(H') = H'$$

$$S = \{ OP_3, OP_1, OP_2 \}$$

Hext

7



$$\langle H_{ext} = \{ (OP_1, OP_2), (OP_1, OP_3) \}$$

S