

Consistency

Robbert van Renesse

What is consistency?

- I know it when I see it...
 - US Supreme Court Justice Potter Stewart, 1964
 - praised as “realistic and gallant”
 - critiqued as “potentially fallacious, due to individualistic arbitrariness”
 - [https://en.wikipedia.org/wiki/I know it when I see it](https://en.wikipedia.org/wiki/I_know_it_when_I_see_it)
- An invariant?
 - What about “eventual consistency”?
- Many definitions and flavors
 - weak/relaxed consistency, strong consistency, entry consistency, lazy consistency, causal consistency, causal+ consistency, sequential consistency, FIFO (PRAM) consistency, serializability, strict serializability, and many more

Let's start with a simple sequential object

- E.g., an integer, a queue, a stack, etc.
- An object has
 - a state
 - a set of methods
- State has an initial value
- Each method may change the state and returns a value of some sort
 - easy to specify usually through a pre-condition and a post-condition
 - we will only consider deterministic methods

Example, a queue specification

- state: sequence of values, initially []
- methods:
 - enqueue(v: Value) -> r: ()
 - state := state :: [v]
 - dequeue() -> r: Value or ERROR
 - if state == []
 - r := ERROR
 - state := state
 - if state <> []:
 - r := head(state)
 - state := tail(state)

Example, a queue specification

- state: sequence of values, initially []
- methods:
 - enqueue(v: Value) -> r: ()
 - state := state :: [v]
 - dequeue() -> r: Value or ERROR
 - if state == []
 - state := state
 - r := ERROR
 - if state <> []:
 - r := head(state)
 - state := tail(state)

Can easily be translated
into TLA+ or Harmony, say

Some (nice) observations about sequential specifications

- State is meaningful only *between* method calls
- Methods only interacts through passing state
- Sequence of operations (method calls) defines a behavior
- Specification is “linear” in the number of methods
- Can add new methods without having to change the old ones
- “inconsistency” simply is violating the spec:
 - enqueue(1)
 - enqueue(2)
 - dequeue() → 2

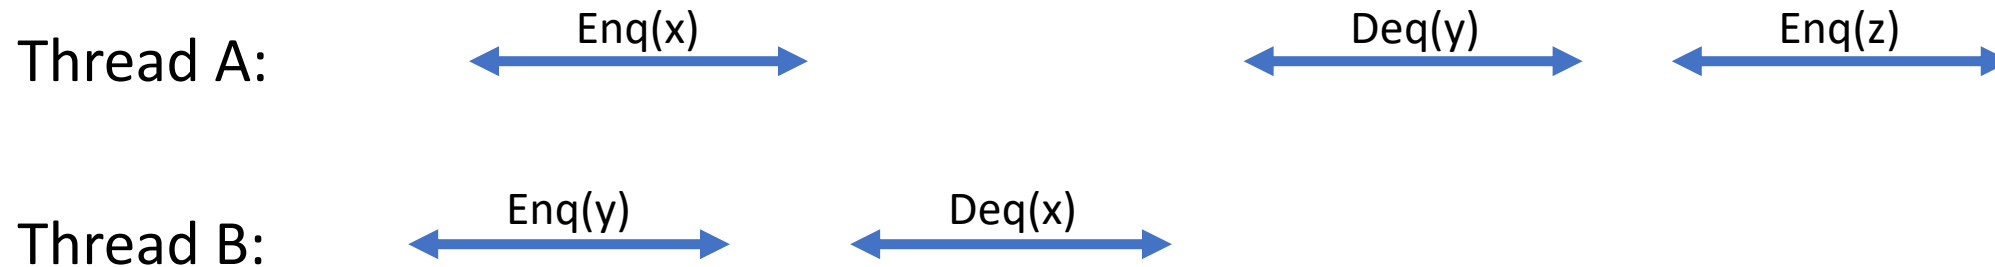
What happens when you add concurrency?

- E.g., what happens when two enqueue() operations are invoked at approximately the same time by two different threads (processes)?
- You have to consider what happens with state *during* an operation
 - all possible interactions...
- Operations take time! (Who knew?)
- Operations of different threads *overlap*
 - There many never be “between method calls”
- Many different cases to consider
 - Specification complicated and not linear in the number of methods 😞

What do we mean by consistency??

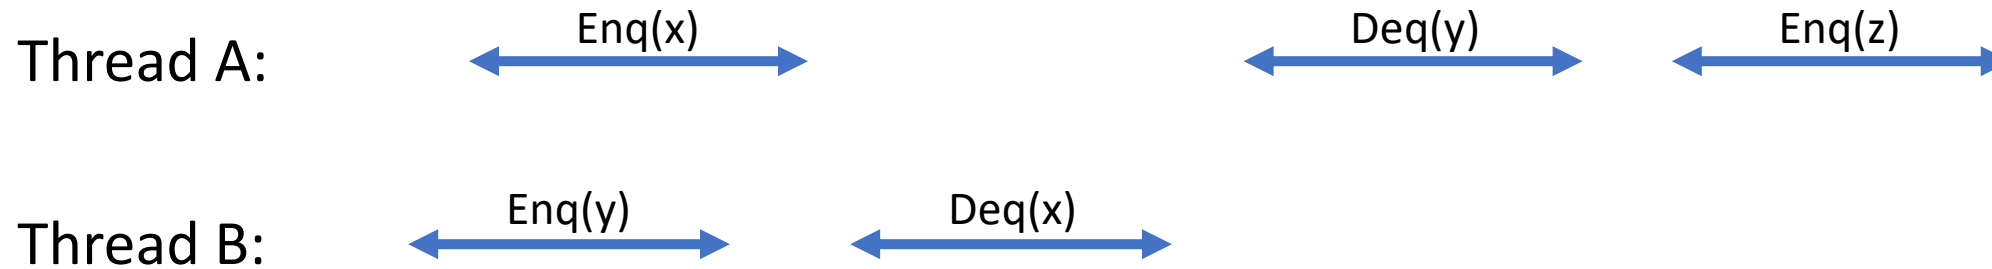
- Let's look at some examples

Example 1



“Consistent” or not?

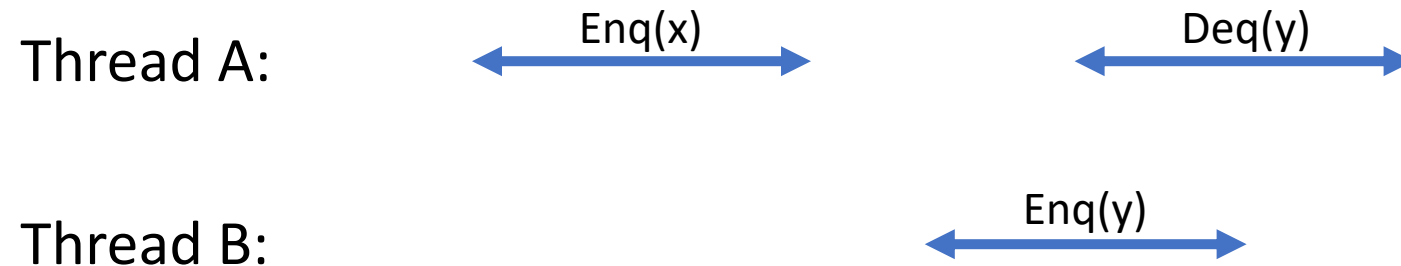
Example 1



“Consistent” or not?





Example 2



“Consistent” or not?


Example 2



Thread A:  



Thread B: 


“Consistent” or not?




Example 2




Thread A:  Enq(x)  Deq(y)

Thread B:  Enq(y)

“Consistent” or not?



or maybe 


Example 3



Thread A:



Thread B:



“Consistent” or not?


Example 3



Thread A:



Thread B:



“Consistent” or not?



What about simple register read/write?

- Operations: $R(x)$, $W(x)$
- Sequential spec:
 - read operation returns value of latest completed write operation
- But what if read and write operations can execute concurrently??

Example 4



Thread A: \longleftrightarrow W(0) \longleftrightarrow

\longleftrightarrow R(1) \longleftrightarrow

Thread B: \longleftrightarrow W(1) \longleftrightarrow

“Consistent” or not?

Example 4



Thread A: \longleftrightarrow W(0) \longleftrightarrow

\longleftrightarrow R(1) \longleftrightarrow

Thread B: \longleftrightarrow W(1) \longleftrightarrow

“Consistent” or not?



Example 5



Thread A: \longleftrightarrow W(0) \longleftrightarrow

\longleftrightarrow R(1) \longleftrightarrow

\longleftrightarrow R(0) \longleftrightarrow

Thread B:


\longleftrightarrow W(1) \longleftrightarrow


“Consistent” or not?

Example 5




Thread A: 





Thread B: 

“Consistent” or not?



Example 6



Thread A: \longleftrightarrow W(0) \longleftrightarrow

\longleftrightarrow R(1) \longleftrightarrow

\longleftrightarrow W(0) \longleftrightarrow

Thread B:

\longleftrightarrow W(1) \longleftrightarrow


\longleftrightarrow R(0) \longleftrightarrow


“Consistent” or not?

Example 6





Thread A: 





Thread B:






“Consistent” or not?



Example 7




Thread A: 

 R(1)

 W(0)

Thread B:

 W(1)


 R(1)


“Consistent” or not?

Example 7




Thread A: 






Thread B:





“Consistent” or not?



or maybe




Linearizability (Herlihy and Wing 1990)

- Each operation appears to have been executed atomically (instantaneously) at some time between its invocation and completion
 - known as “linearization point”
- Implementation is linearizable iff for every behavior you can find a corresponding sequential behavior of linearization points

Sequential Consistency (Lamport 1979)

- The result of any execution is the same as if the operations of all processes were executed in some sequential order and the operations of each process appear in this sequence in the order specified by its program




Example 2: sequentially consistent but not linearizable


Linearizability vs Sequential Consistency

- Linearizability implies Sequential Consistency (but not vice versa)
 - i.e., linearizability is a stronger consistency property than sequential consistency
 - sequential consistency allows more interleavings than linearizability
 - more concurrency, but harder to reason about
- Linearizability is a *local property*, but sequential consistency is not
 - Linearizability composes: a system of linearizable objects is linearizable
 - Vice versa: in a linearizable system each object is linearizable

Example 8



Thread A:  p.Enq(1)

 q.Enq(3)

 p.Dec(4)

Thread B:  q.Enq(2)

 p.Enq(4)

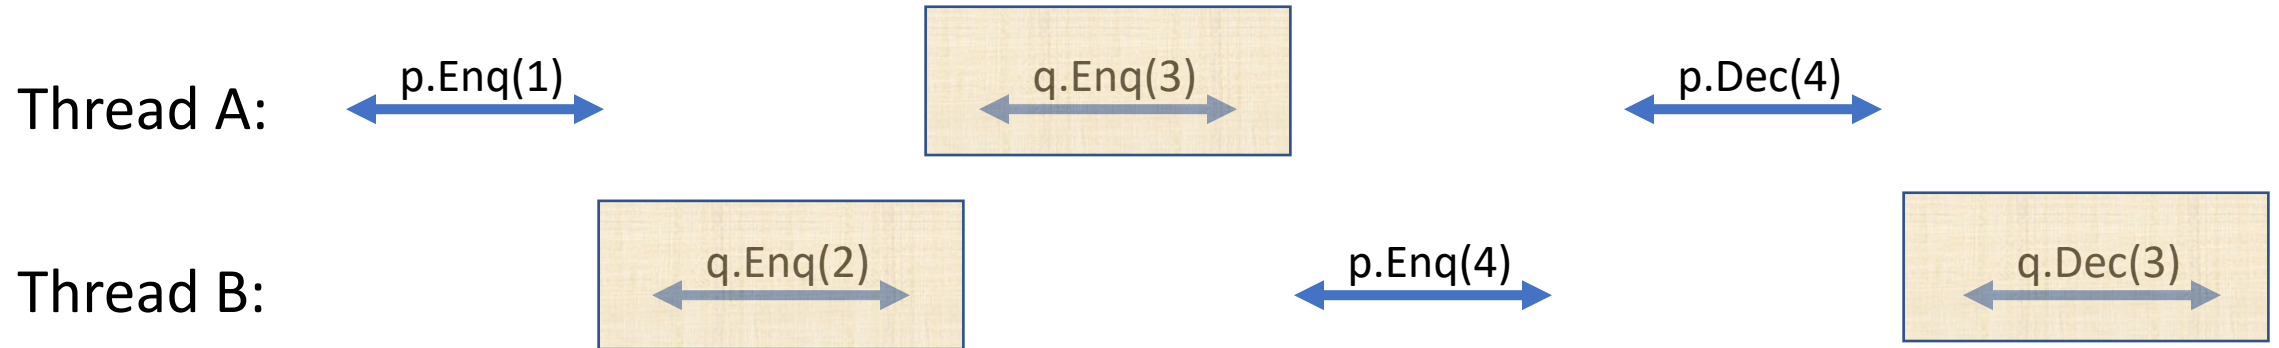
 q.Dec(3)

“Consistent” or not?

Linearizable?

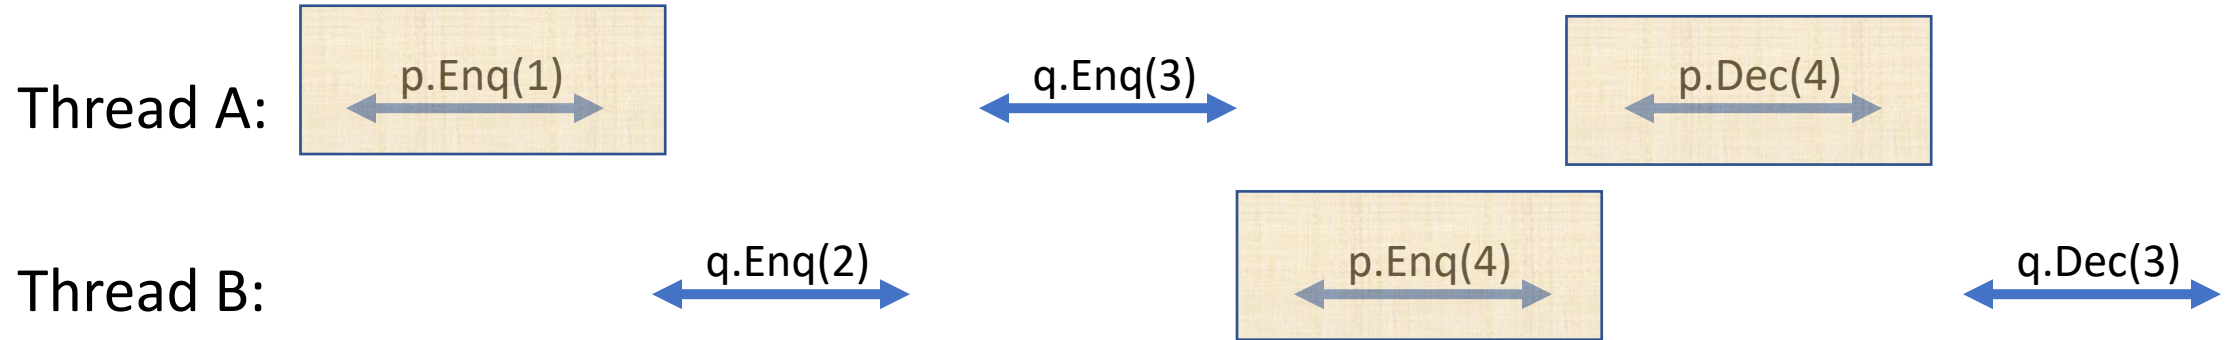
Sequentially consistent?

Example 8




Operations on p are sequentially consistent

Example 8




Operations on q are sequentially consistent

Example 8



But entire history is not sequentially consistent


Example 9



- Just Thread A: sequentially consistent
- Just Thread B: sequentially consistent
- Just location x: sequentially consistent
- Just location y: sequentially consistent
- Overall: not sequentially consistent

A model of linearizability

- Each object implemented by a sequential server
- Communication is through sending requests and receiving responses



Verifying Linearizability

- In general, need to identify linearization points and show that they form a legal sequential history of operations

Simple concurrent queue in Harmony

```
1  from synch import Lock, acquire, release
2  from alloc import malloc, free
3
4  def Queue():
5      result = { .head: None, .tail: None, .lock: Lock() }
6
7  def put(q, v):
8      let node = malloc({ .value: v, .next: None }):
9          acquire(?q→lock)
10         if q→tail == None:
11             q→tail = q→head = node
12         else:
13             q→tail→next = node
14             q→tail = node
15         release(?q→lock)
```

```
17  def get(q):
18      acquire(?q→lock)
19      let node = q→head:
20          if node == None:
21              result = None
22          else:
23              result = node→value
24              q→head = node→next
25              if q→head == None:
26                  q→tail = None
27              free(node)
28      release(?q→lock)
```

Seq queue spec in Harmony and seq test

```
1      import list
2
3      def Queue():
4          result = [ ]
5
6      def put(q, v):
7          !q = list.append(!q, v)
8
9      def get(q):
10         if !q == [ ]:
11             result = None
12         else:
13             result = list.head(!q)
14             !q = list.tail(!q)
```

```
1      import queue, queuespec
2
3      const NOPS = 4
4      const VALUES = { 1..NOPS }
5
6      implq = queue.Queue()
7      specq = queuespec.Queue()
8
9      for i in {1..NOPS}:
10         let op = choose({ .get, .put }):
11             if op == .put:
12                 let v = choose(VALUES):
13                     queue.put(?implq, v)
14                     queuespec.put(?specq, v)
15             else:
16                 let v = queue.get(?implq)
17                 let w = queuespec.get(?specq):
18                     assert v == w
```

Verifying linearizability

- Linearization points: any point between acquire and release lock

Specifying linearization points

```
1  from synch import Lock, acquire, release
2  from alloc import malloc, free
3
4  def Queue():
5      result = { .head: None, .tail: None, .lock: Lock(), .time: 0 }
6
7  def _linpoint(q):
8      atomic:
9          this.qtime = q→time
10         q→time += 1
11
12  def put(q, v):
13      let node = malloc({ .value: v, .next: None }):
14          acquire(?q→lock)
15          if q→tail == None:
16              q→tail = q→head = node
17          else:
18              q→tail→next = node
19              q→tail = node
20          _linpoint(q)
21          release(?q→lock)
```

```
23  def get(q):
24      acquire(?q→lock)
25      let node = q→head:
26          if node == None:
27              result = None
28          else:
29              result = node→value
30              q→head = node→next
31              if q→head == None:
32                  q→tail = None
33              free(node)
34      _linpoint(q)
35      release(?q→lock)
```

Testing linearizability

```
1  import queuelin, queuespec
2
3  const NOPS = 4
4  const VALUES = { 1..NOPS }
5
6  sequential qtime
7  qtime = 0
8
9  implq = queuelin.Queue()
10 specq = queuespec.Queue()
```

```
12  def thread():
13      let op = choose({ .get, .put }):
14          if op == .put:
15              let v = choose(VALUES):
16                  queuelin.put(?implq, v)
17                  await qtime == this.qtime
18                  queuespec.put(?specq, v)
19          else:
20              let v = queuelin.get(?implq):
21                  await qtime == this.qtime
22                  let w = queuespec.get(?specq):
23                      assert v == w
24      atomic:
25          qtime += 1
26
27  for i in {1..NOPS}:
28      spawn thread()
```

How about real memory?

- Registers: atomic (linearizable)
- Memory: not even sequentially consistent
 - write operations are buffered
 - processors, and even compilers, re-order operations in complex ways
 - or even remove operations that are deemed unnecessary but may not be
 - not usually a problem in sequential programs
 - big reads/big writes may be split across multiple instructions
 - i.e., 64-bit read/write on a 32-bit architecture
 - Note: Peterson's algorithm requires sequential consistency
- Modern processor has “memory barriers” or “fence” instructions to force data to memory
 - e.g. **mfence** instruction on x86

In high-level languages

- In Java, you can specify that a variable is “volatile”
 - Adds a memory barrier after each store
 - Inhibits compiler optimizations
- C++ offers various types of “atomic variables” with various consistency guarantees
 - The “volatile” tag inhibits optimizations but does not add a memory barrier

What is Eventual Consistency?

- It is not really consistency at all
- Think instead of anti-entropy protocols

Further reading

- “Art of Multiprocessor Programming” by Maurice Herlihy et al.
- “On Concurrent Programming” by Fred B. Schneider

Atomic Transactions

- From database community:
 - an atomic transaction is a group of operations
 - e.g.:

```
begin_transaction
  if x < y:
    x := y
end_transaction
```
 - ACID properties
 - Atomicity: all or nothing (transactions may commit or abort)
 - Consistency: satisfies application-level invariants
 - **Isolation: appears as if transactions are executed serially**
 - Durable: effects of successful transactions are permanent

Transactions: commit or abort

- If a transaction commits then all its actions are permanent
- If a transaction aborts then there are no (visible) actions
- Typical usage: try until successful

do

 begin_transaction()

 ...

 ...

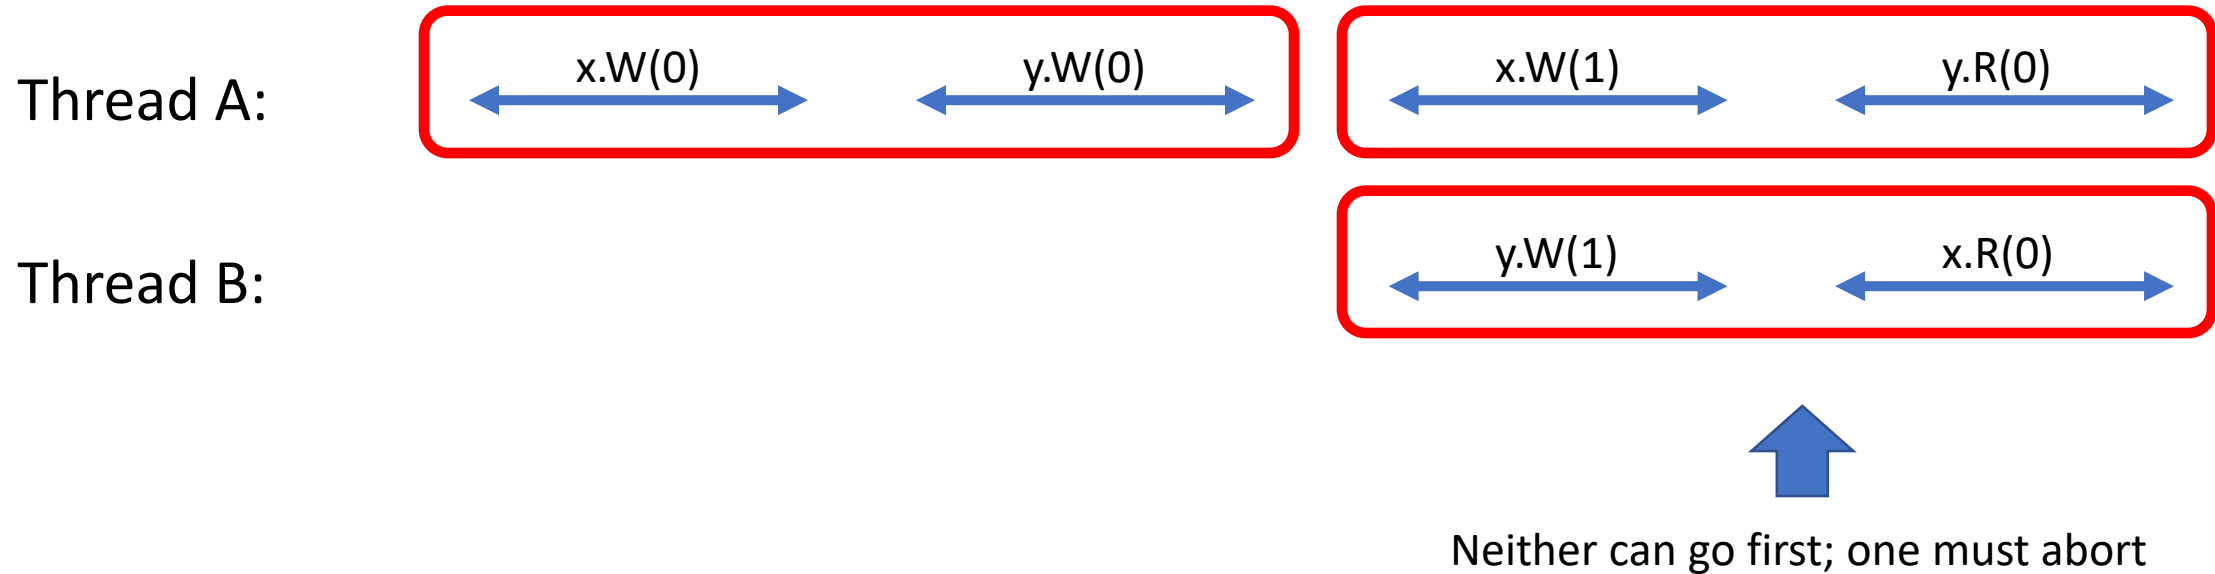
while (end_transaction() == ABORT)

Serializability

- Serializability: (successful) transactions appear to execute sequentially
 - i.e., *isolation*
- Strict serializability: Consistent with real-time order
 - if transaction B starts after transaction A finishes, then B must be ordered after A
- Linearizability is a special case of strict serializability
 - transactions with a single operation each

Strict serializability is not local!

(Similar to Example 9)



Concurrency Control

- Ways to guarantee serializability: isolation between transactions
 - Pessimistic: grab read/write locks as the transaction is progressing
 - 2 phase locking
 - don't release locks until end of transaction
 - acquire locks in some global order to prevent deadlock
 - Optimistic: keep track of read/write sets and check for conflicts at the end
 - abort transaction if its write set intersects with the read set or write set of a concurrent committed transaction or its read set intersects with the write set of a concurrent committed transaction

Atomic Commitment: 2 phase commit

- Actors: one coordinator and two or more participants
- Protocol:
 - 1A: coordinator broadcasts PREPARE to all participants
 - 1B: participants reply with either YES or NO
 - if YES, participant promises to remain ready to move forward
 - 2A: coordinator broadcasts COMMIT *only if* all participants responded YES
 - if some participant responds with NO, or does not respond, then coordinator broadcasts ABORT
 - 2B: upon receiving COMMIT, participant finalizes local operations
 - upon receiving ABORT, participant backs out of local operations
 - release locks if any

Example: bank

```
1  network = {}
2
3  def send(m):
4      atomic: network |= { m }
5
6  def bank(self, balance):
7      let status, received = (), {}:
8          while True:
9              select req in network - received where req.dst == self:
10                 received |= { req }
11                 if req.request == .withdraw:
12                     if (status != ()) or (req.amount > balance):
13                         send({ .dst: req.src, .src: self, .response: .no })
14                     else:
15                         status = balance
16                         balance -= req.amount
17                         send({ .dst: req.src, .src: self, .response: .yes, .funds: balance })
18                 elif req.request == .deposit:
19                     if status != ():
20                         send({ .dst: req.src, .src: self, .response: .no })
21                     else:
22                         status = balance
23                         balance += req.amount
24                         send({ .dst: req.src, .src: self, .response: .yes, .funds: balance })
25                 elif req.request == .commit:
26                     assert status != ()
27                     status = ()
28                 else:
29                     assert (status != ()) and (req.request == .abort)
30                     balance, status = status, ()
```

Transfer:

```
1
2
3  const NBANKS = 3
4  const NCOORDS = 2
5  const MAX_BALANCE = 1
6
7  def receive(self, sources):
8      let forme = { m for m in network where m.dst == self }:
9          result = { forme } if { m.src for m in forme } == sources else {}
10
11  def transfer(self, b1, b2, amt):
12      send({ .dst: b1, .src: self, .request: .withdraw, .amount: amt })
13      send({ .dst: b2, .src: self, .request: .deposit, .amount: amt })
14      select msgs in receive(self, { b1, b2 }):
15          if all(m.response == .yes for m in msgs):
16              possibly True
17              for m in msgs where m.response == .yes:
18                  send({ .dst: m.src, .src: self, .request: .commit })
19          else:
20              for m in msgs where m.response == .yes:
21                  send({ .dst: m.src, .src: self, .request: .abort })
```

ACID revisited

ACID properties

- Atomicity: all or nothing (transactions may commit or abort)
 - 2PC protocol or some other atomic commitment protocol
- Consistency: satisfies application-level invariants
 - That's up to the application (for example, prevent negative bank balances)
- **Isolation: appears as if transactions are executed serially**
 - Concurrency control protocol such as 2PL
- Durable: effects of successful transactions are permanent
 - Pragmatically speaking: store data on disk ideally before you commit