

# **SAFETY, LIVENESS, AND CONSISTENCY**

Ellis Michael



# How Do We SPECIFY DISTRIBUTED SYSTEMS?

- **Execution:** Sequence of events (i.e., steps taken by the system), potentially infinite.
- **Property:** A predicate on executions.
- **Safety property:** Specifies the "bad things" that shouldn't happen in any execution.
- **Liveness property:** Specifies the "good things" that should happen in every execution.

(See paper for formal definitions.)



***THEOREM: EVERY PROPERTY IS EXPRESSIBLE AS THE  
CONJUNCTION OF A SAFETY PROPERTY AND A LIVENESS  
PROPERTY.***

[Alpern and Schneider. 1987]



***THEOREM: EVERY PROPERTY IS EXPRESSIBLE AS THE  
CONJUNCTION OF A SAFETY PROPERTY AND A LIVENESS  
PROPERTY.***

Neat automata theory!

[Alpern and Schneider. 1987]



# SOME PROPERTIES

- The system never deadlocks.
- Every client that sends a request eventually gets a reply.
- Both generals attack simultaneously.



# MORE PROPERTIES: CONSENSUS

$n$  processes, all of which have an input value from some domain.

Processes output a value by calling  $decide(v)$ . Non-faulty processes continue correctly executing protocol steps forever. We usually denote the number of faulty processes  $f$ .

- **Agreement:** No two correct processes decide different values.
- **Integrity:** Every correct process decides at most one value, and if a correct process decides a value  $v$ , some process had  $v$  as its input.
- **Termination:** Every correct process eventually decides a value.



# CONSISTENCY IS KEY!

**Consistency:** *the allowed semantics (return values) of a set of operations to a data store or shared object.*

Consistency properties specify the **interface**, not the **implementation**. The data might be replicated, cached, disaggregated, etc. "Weird" consistency semantics happen all over the stack!

**Anomaly:** violation of the consistency semantics



# TERMINOLOGY: STRENGTH AND WEAKNESS

- **Strong consistency:** the system behaves as if there's just a single copy of the data (or almost behaves that way).

The intuition is that things like caching and sharding are implementation decisions and shouldn't be visible to clients.

- **Weak consistency:** allows behaviors significantly different from the single store model.
- **Eventual consistency:** the aberrant behaviors are only temporary.



# WHY THE DIFFERENCE?

- **Performance**

- Consistency requires synchronization/coordination when data is replicated
- Often slower to make sure you always return right answer

- **Availability**

- What if client is offline, or network is not working?
- Weak/eventual consistency may be only option

- **Programmability**

- Weaker models are harder to reason against



# LAMPORT'S REGISTER SEMANTICS

Registers hold a single value. Here, we consider single-writer registers only supporting write and read.

Semantics defined in terms of the *real-time* beginnings and ends of operations to the object.

- **safe:** a read not concurrent with any write obtains the previously written value
- **regular:** safe + a read that overlaps a write obtains either the old or new value
- **atomic:** safe + reads and writes behave as if they occur in some definite order






# LAMPORT'S REGISTER SEMANTICS

Registers hold a single value. Here, we consider single-writer registers only supporting write and read.

Semantics defined in terms of the *real-time* beginnings and ends of operations to the object.

- **safe:** a read not concurrent with any write obtains the previously written value
- **regular:** safe + a read that overlaps a write obtains either the old or new value
- **atomic:** safe + reads and writes behave as if they occur in some definite order



**safe**  $\Rightarrow r_1 \rightarrow a$




# LAMPORT'S REGISTER SEMANTICS

Registers hold a single value. Here, we consider single-writer registers only supporting write and read.

Semantics defined in terms of the *real-time* beginnings and ends of operations to the object.

- **safe:** a read not concurrent with any write obtains the previously written value
- **regular:** safe + a read that overlaps a write obtains either the old or new value
- **atomic:** safe + reads and writes behave as if they occur in some definite order



**safe**  $\Rightarrow r_1 \rightarrow a$

**regular**  $\Rightarrow r_1 \rightarrow a \wedge (r_2 \rightarrow a \vee r_2 \rightarrow b) \wedge$   
 $(r_3 \rightarrow a \vee r_3 \rightarrow b)$



# LAMPORT'S REGISTER SEMANTICS

Registers hold a single value. Here, we consider single-writer registers only supporting write and read.

Semantics defined in terms of the *real-time* beginnings and ends of operations to the object.

- **safe:** a read not concurrent with any write obtains the previously written value
- **regular:** safe + a read that overlaps a write obtains either the old or new value
- **atomic:** safe + reads and writes behave as if they occur in some definite order



**safe**  $\Rightarrow r_1 \rightarrow a$

**regular**  $\Rightarrow r_1 \rightarrow a \wedge (r_2 \rightarrow a \vee r_2 \rightarrow b) \wedge$   
 $(r_3 \rightarrow a \vee r_3 \rightarrow b)$

**atomic**  $\Rightarrow r_1 \rightarrow a \wedge (r_2 \rightarrow a \vee r_2 \rightarrow b) \wedge$   
 $(r_3 \rightarrow a \vee r_3 \rightarrow b) \wedge$   
 $(r_2 \rightarrow b \Rightarrow r_3 \rightarrow b)$



# SEQUENTIAL CONSISTENCY


- Applies to arbitrary shared objects.
- Requires that a history of operations be *equivalent to a legal sequential history*, where a legal sequential history is one that respects the local ordering at each node.
- Called **serializability** when applied to transactions



# Is It SEQUENTIAL?




# Is It SEQUENTIAL?



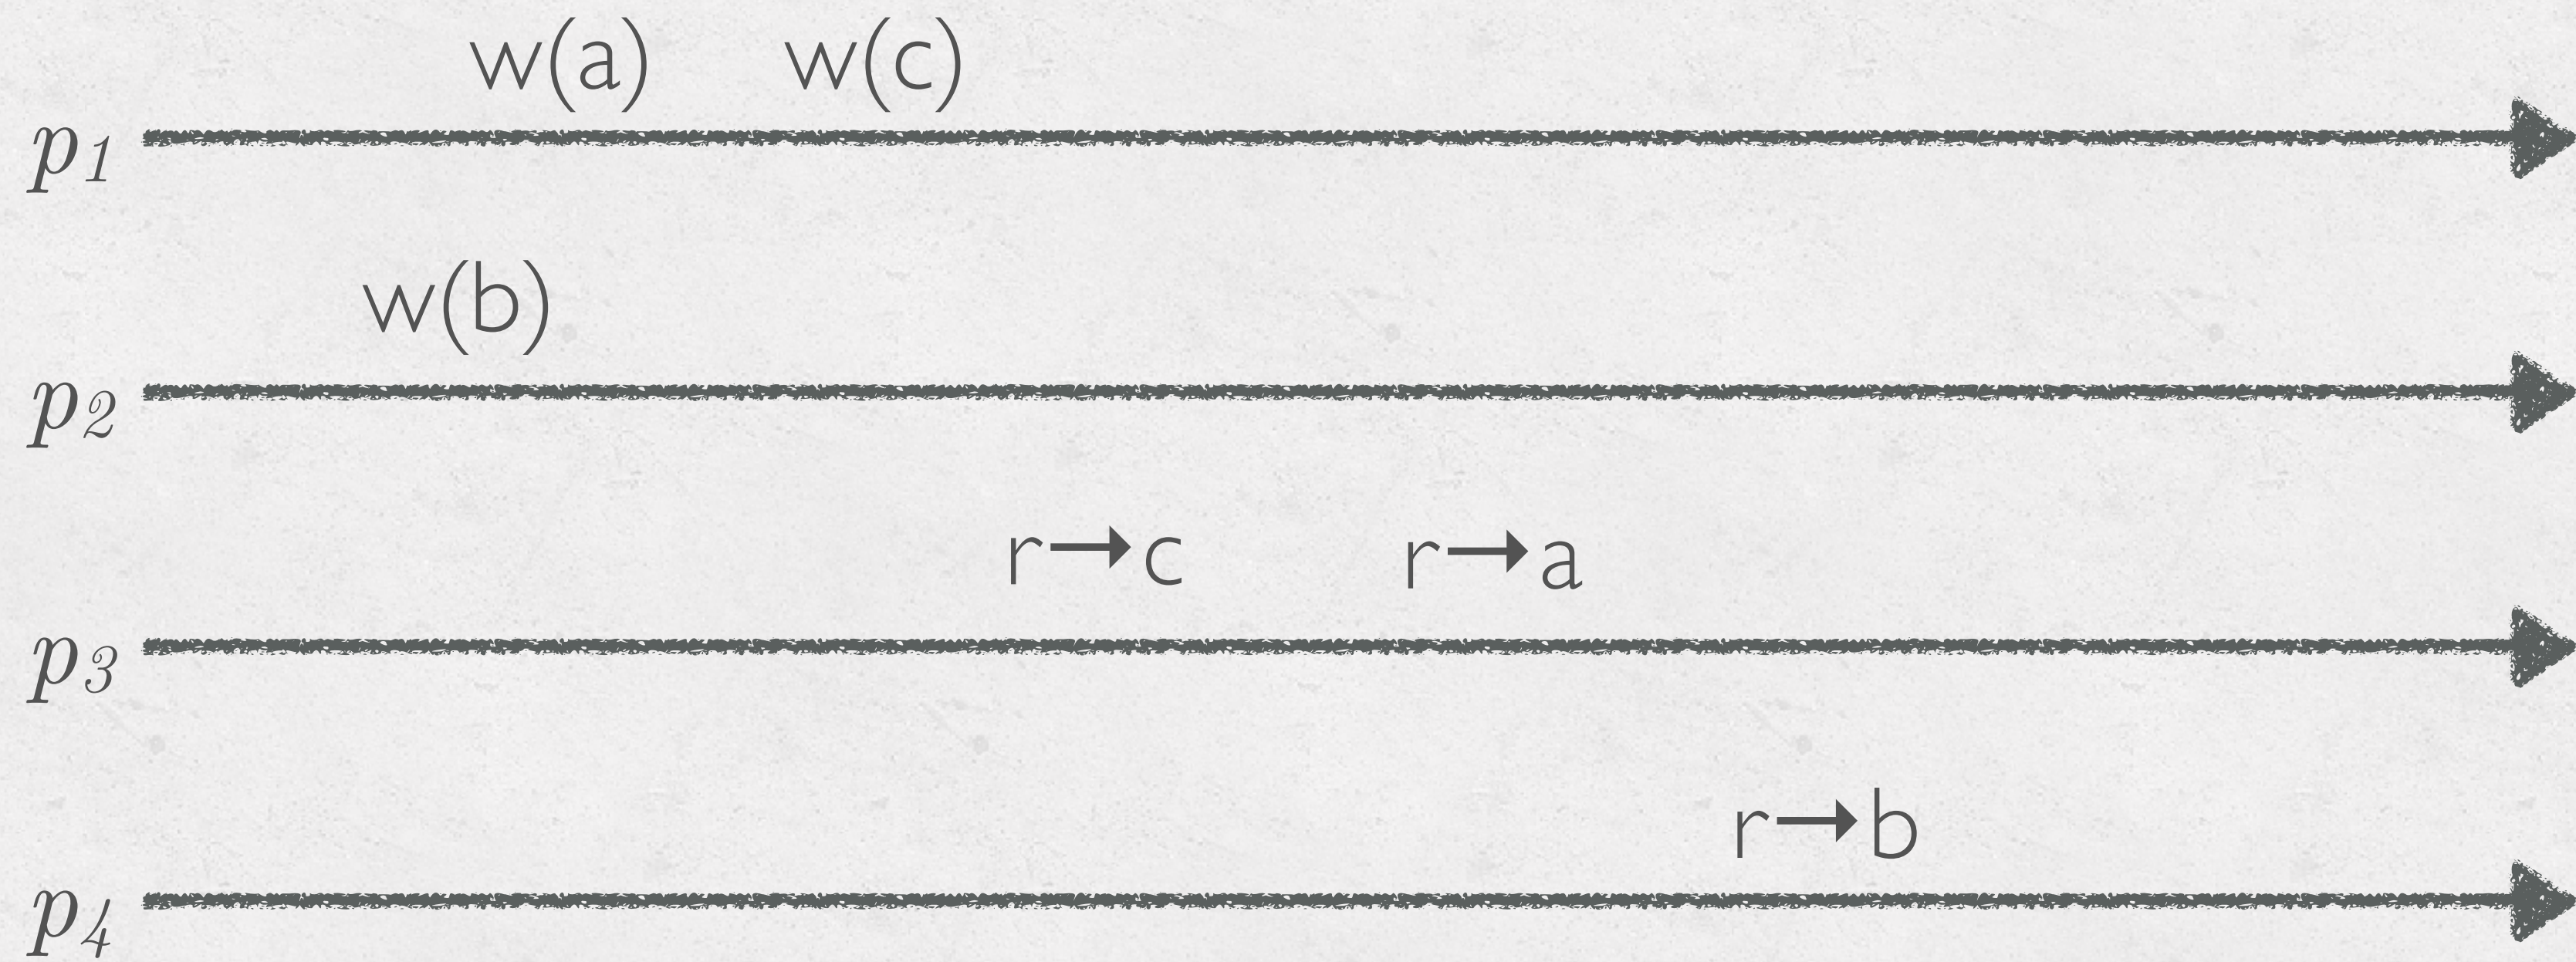


# Is It SEQUENTIAL?



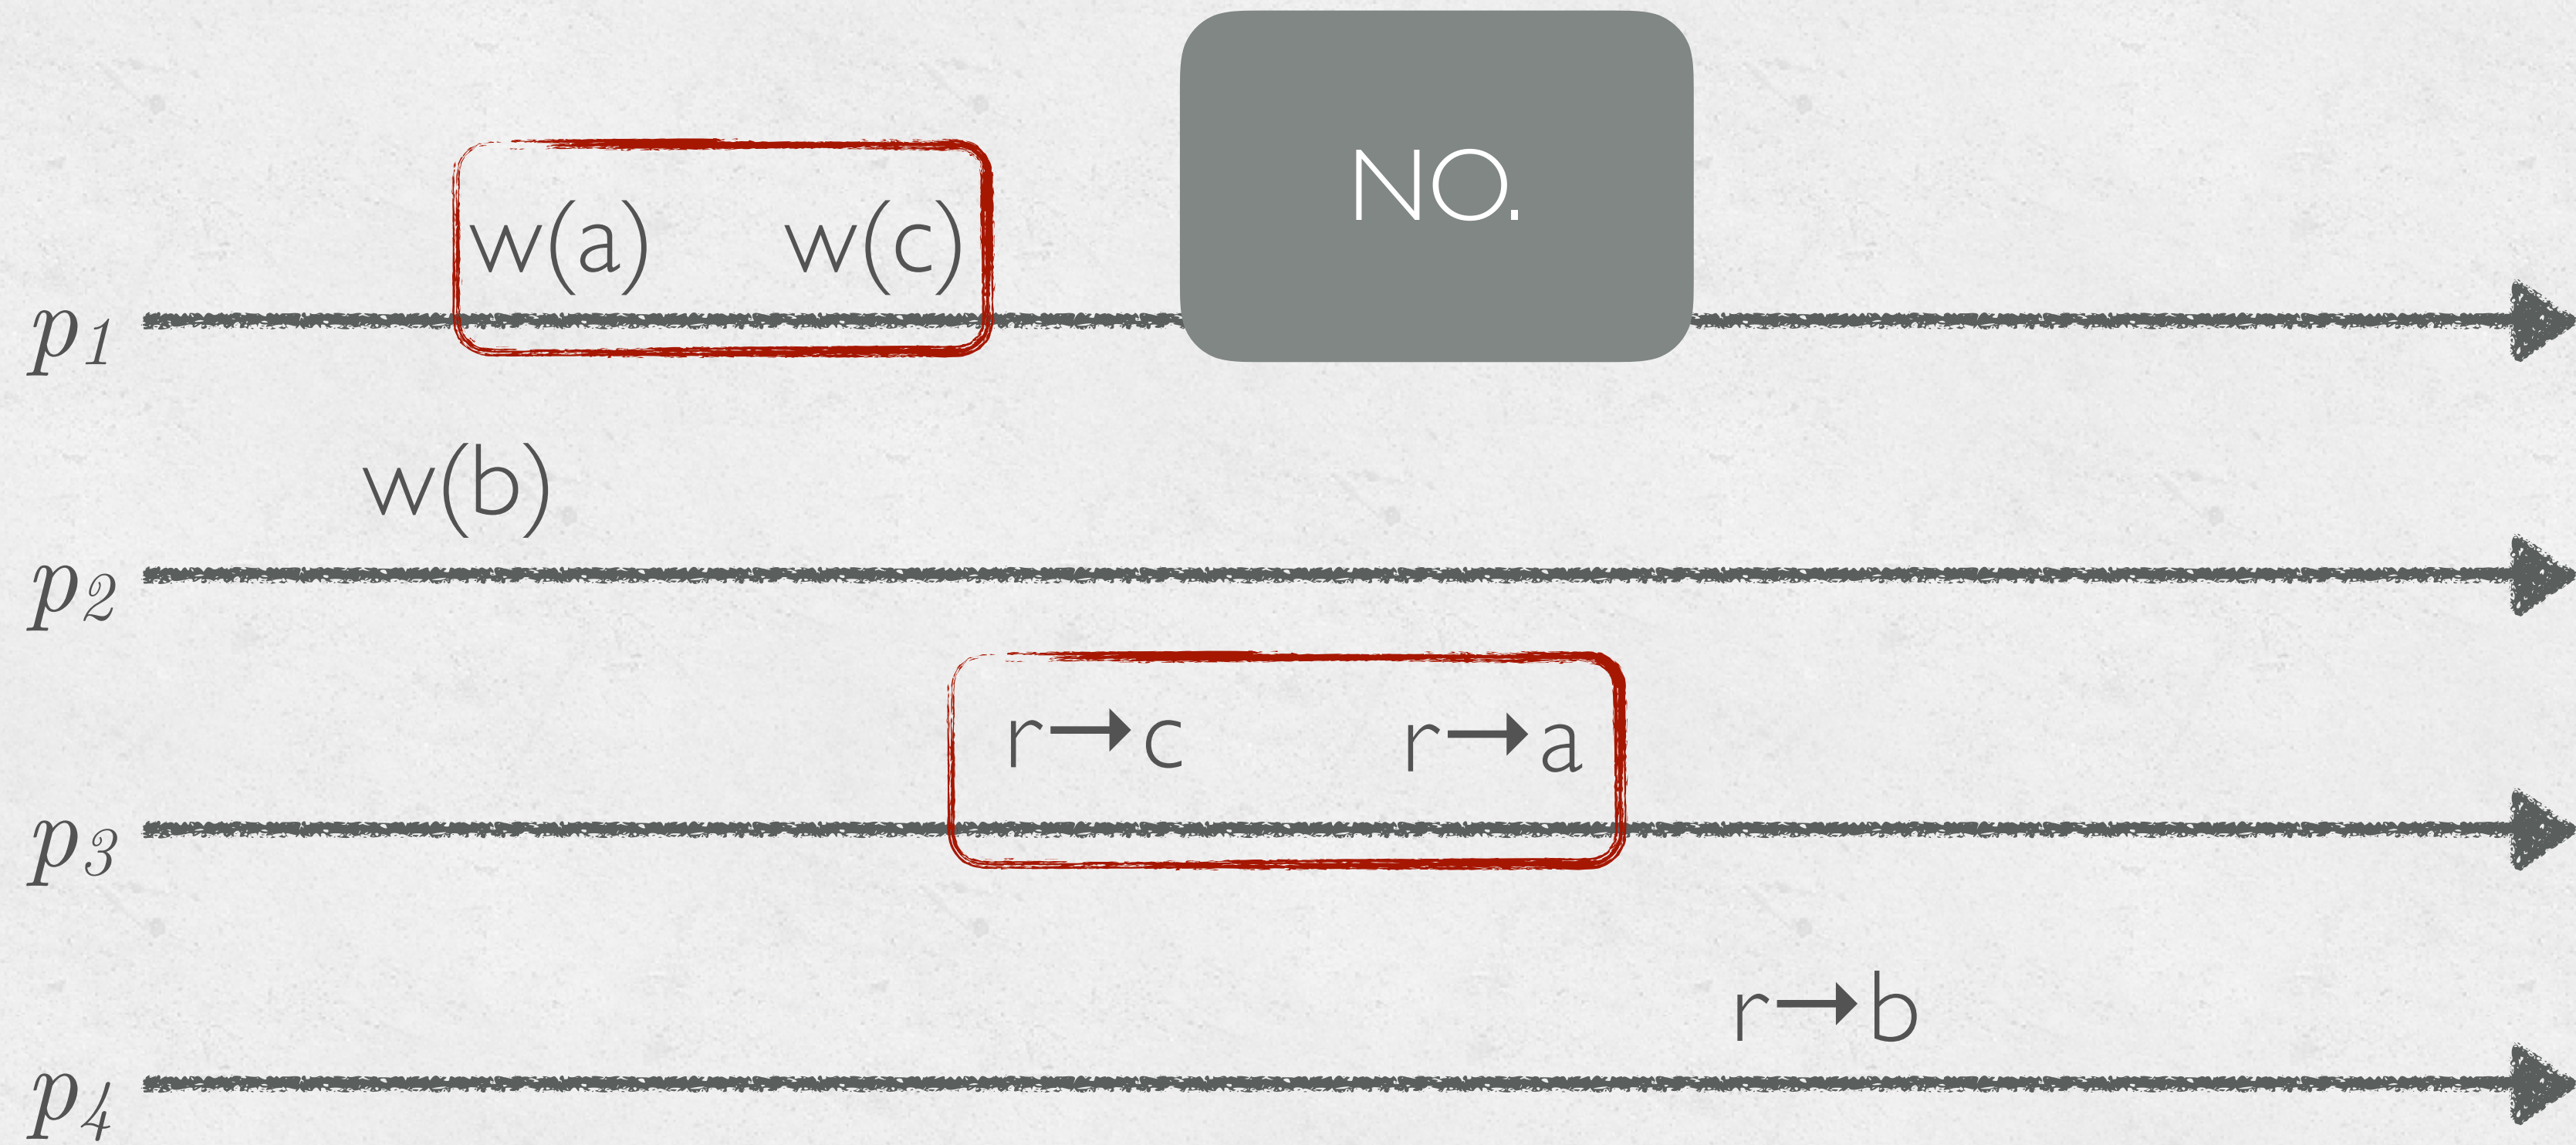


# Is It SEQUENTIAL?






# Is It SEQUENTIAL?






# Is It SEQUENTIAL?






# Is It SEQUENTIAL?

$w(a)$     $r \rightarrow a$     $r \rightarrow a$     $r \rightarrow a$     $w(b)$     $r \rightarrow b$



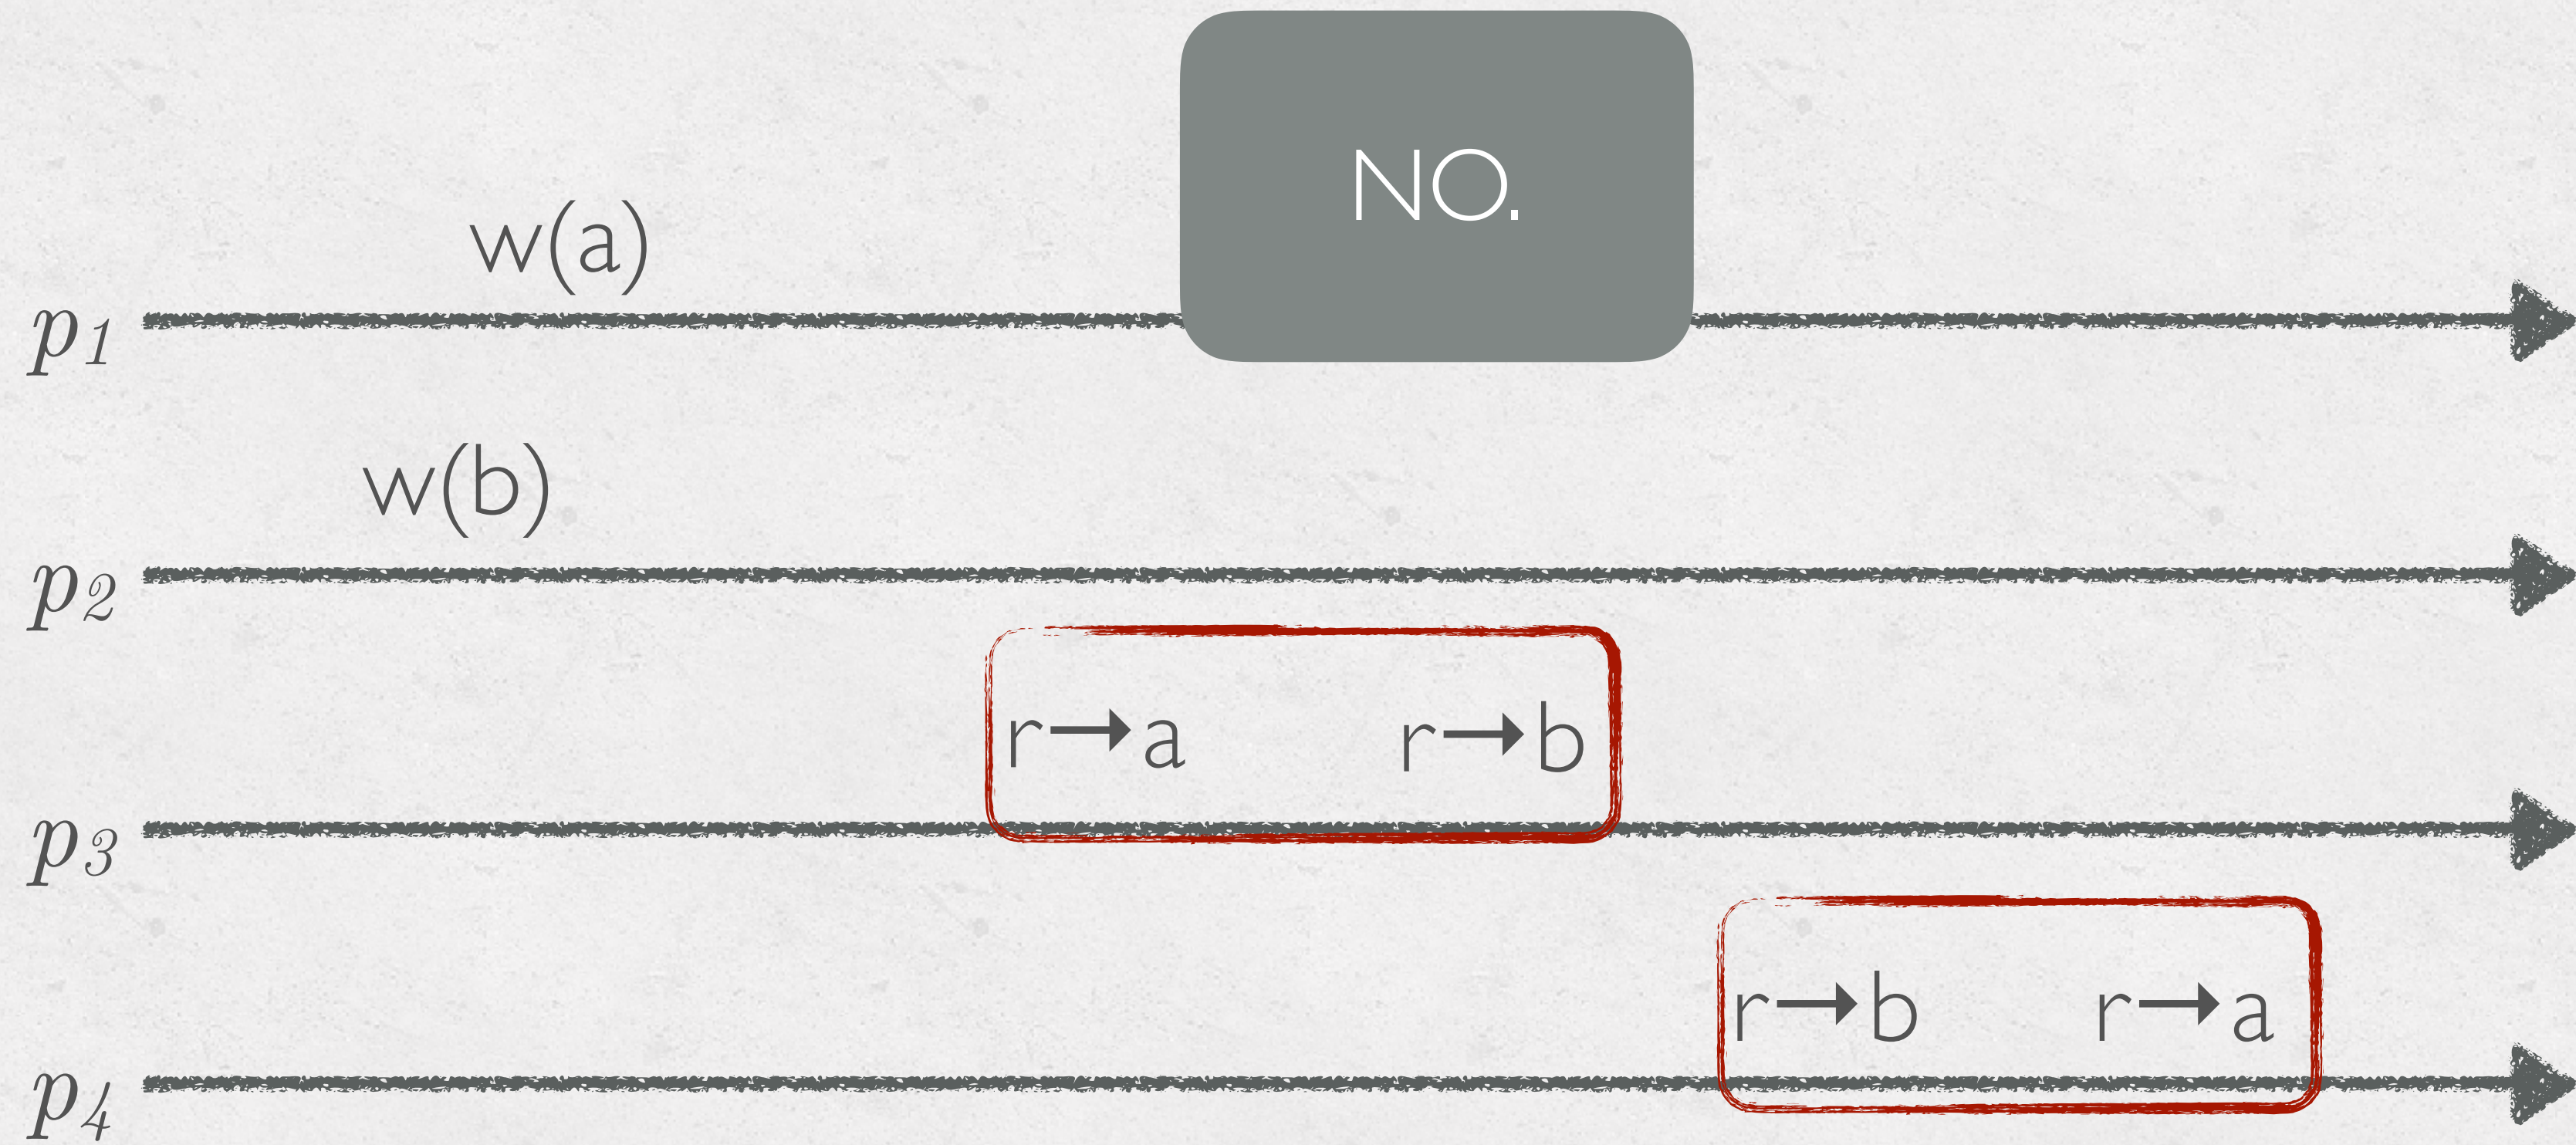


# Is It SEQUENTIAL?





# Is It SEQUENTIAL?





# LINEARIZABILITY

**Linearizability** = sequential consistency + respects real-time ordering.

If  $e_1$  **ends** before  $e_2$  **begins**, then  $e_1$  *appears before*  $e_2$  in the sequential history.


Linearizable data structures behave as if there's a single, correct copy.



Atomic registers are linearizable.

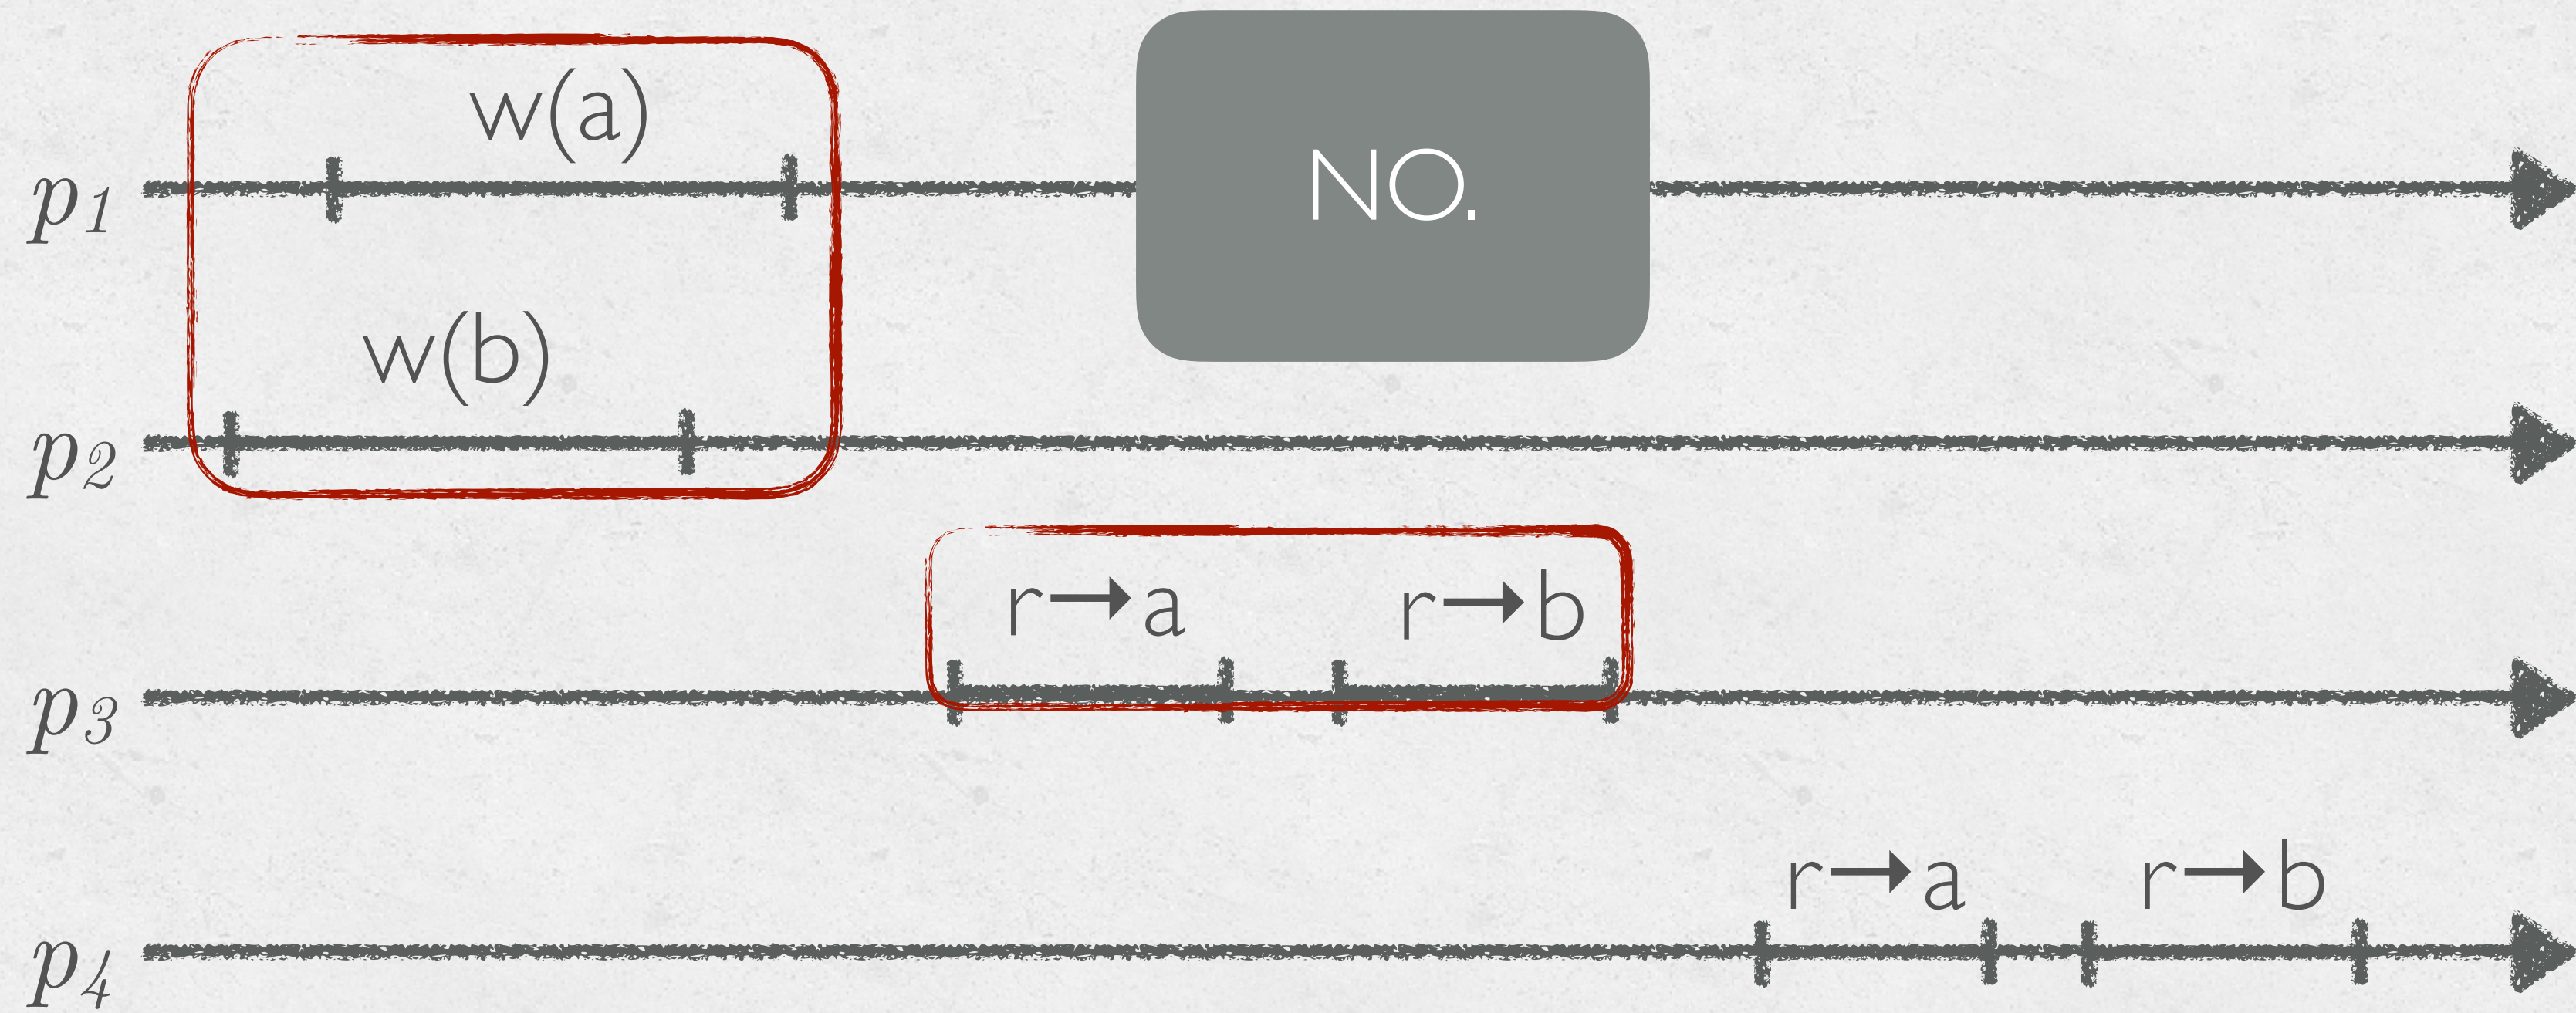


# Is It LINEARIZABLE?






# Is It LINEARIZABLE?






# Is It LINEARIZABLE?





# Is It LINEARIZABLE?

YES!





# LINEARIZABILITY VS. SEQUENTIAL CONSISTENCY

- Sequential consistency allows operations to appear out of real-time order. How could that happen in reality?



# LINEARIZABILITY VS. SEQUENTIAL CONSISTENCY

- Sequential consistency allows operations to appear out of real-time order. How could that happen in reality?
- The most common way systems are sequentially consistency but not linearizability is that they allow read-only operations to return **stale data**.



# STALE READS






# STALE READS



Primary Copy






# STALE READS



Primary Copy



Read-only Cache






# STALE READS



Primary Copy

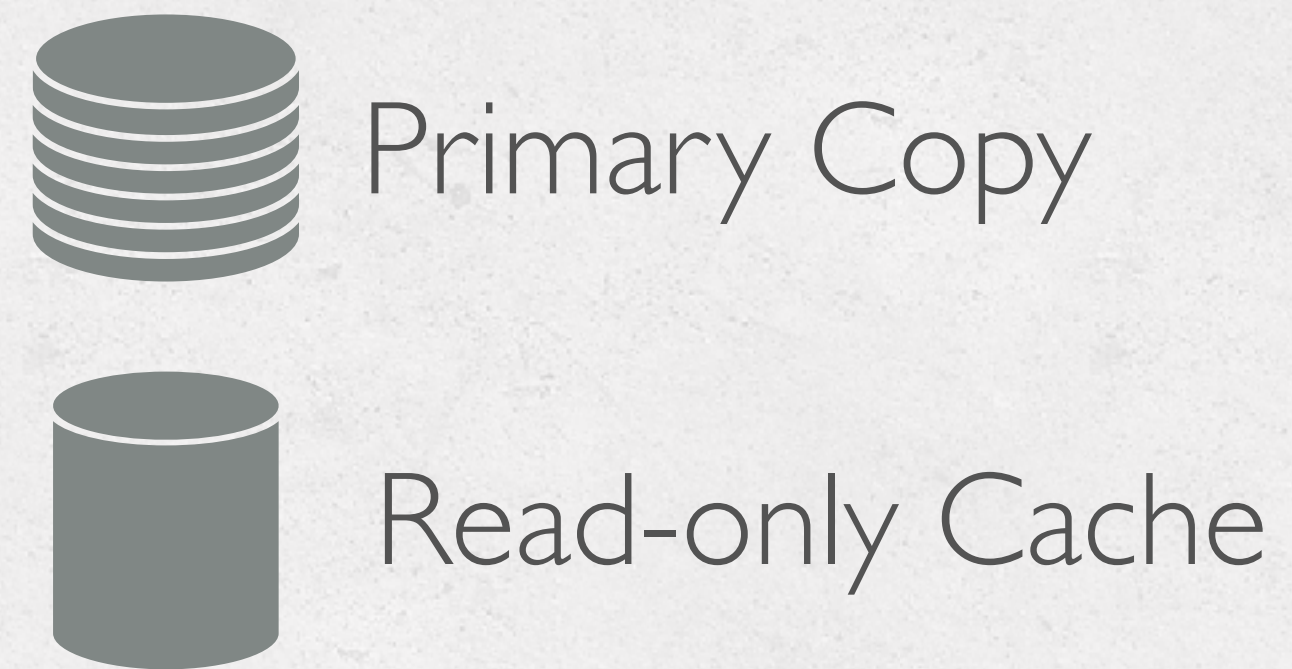


Read-only Cache





# STALE READS





# CAUSAL CONSISTENCY

- Writes that are not concurrent (i.e., writes related by the happens-before relation) must be seen in that order. Concurrent writes can be seen in different orders on different nodes.

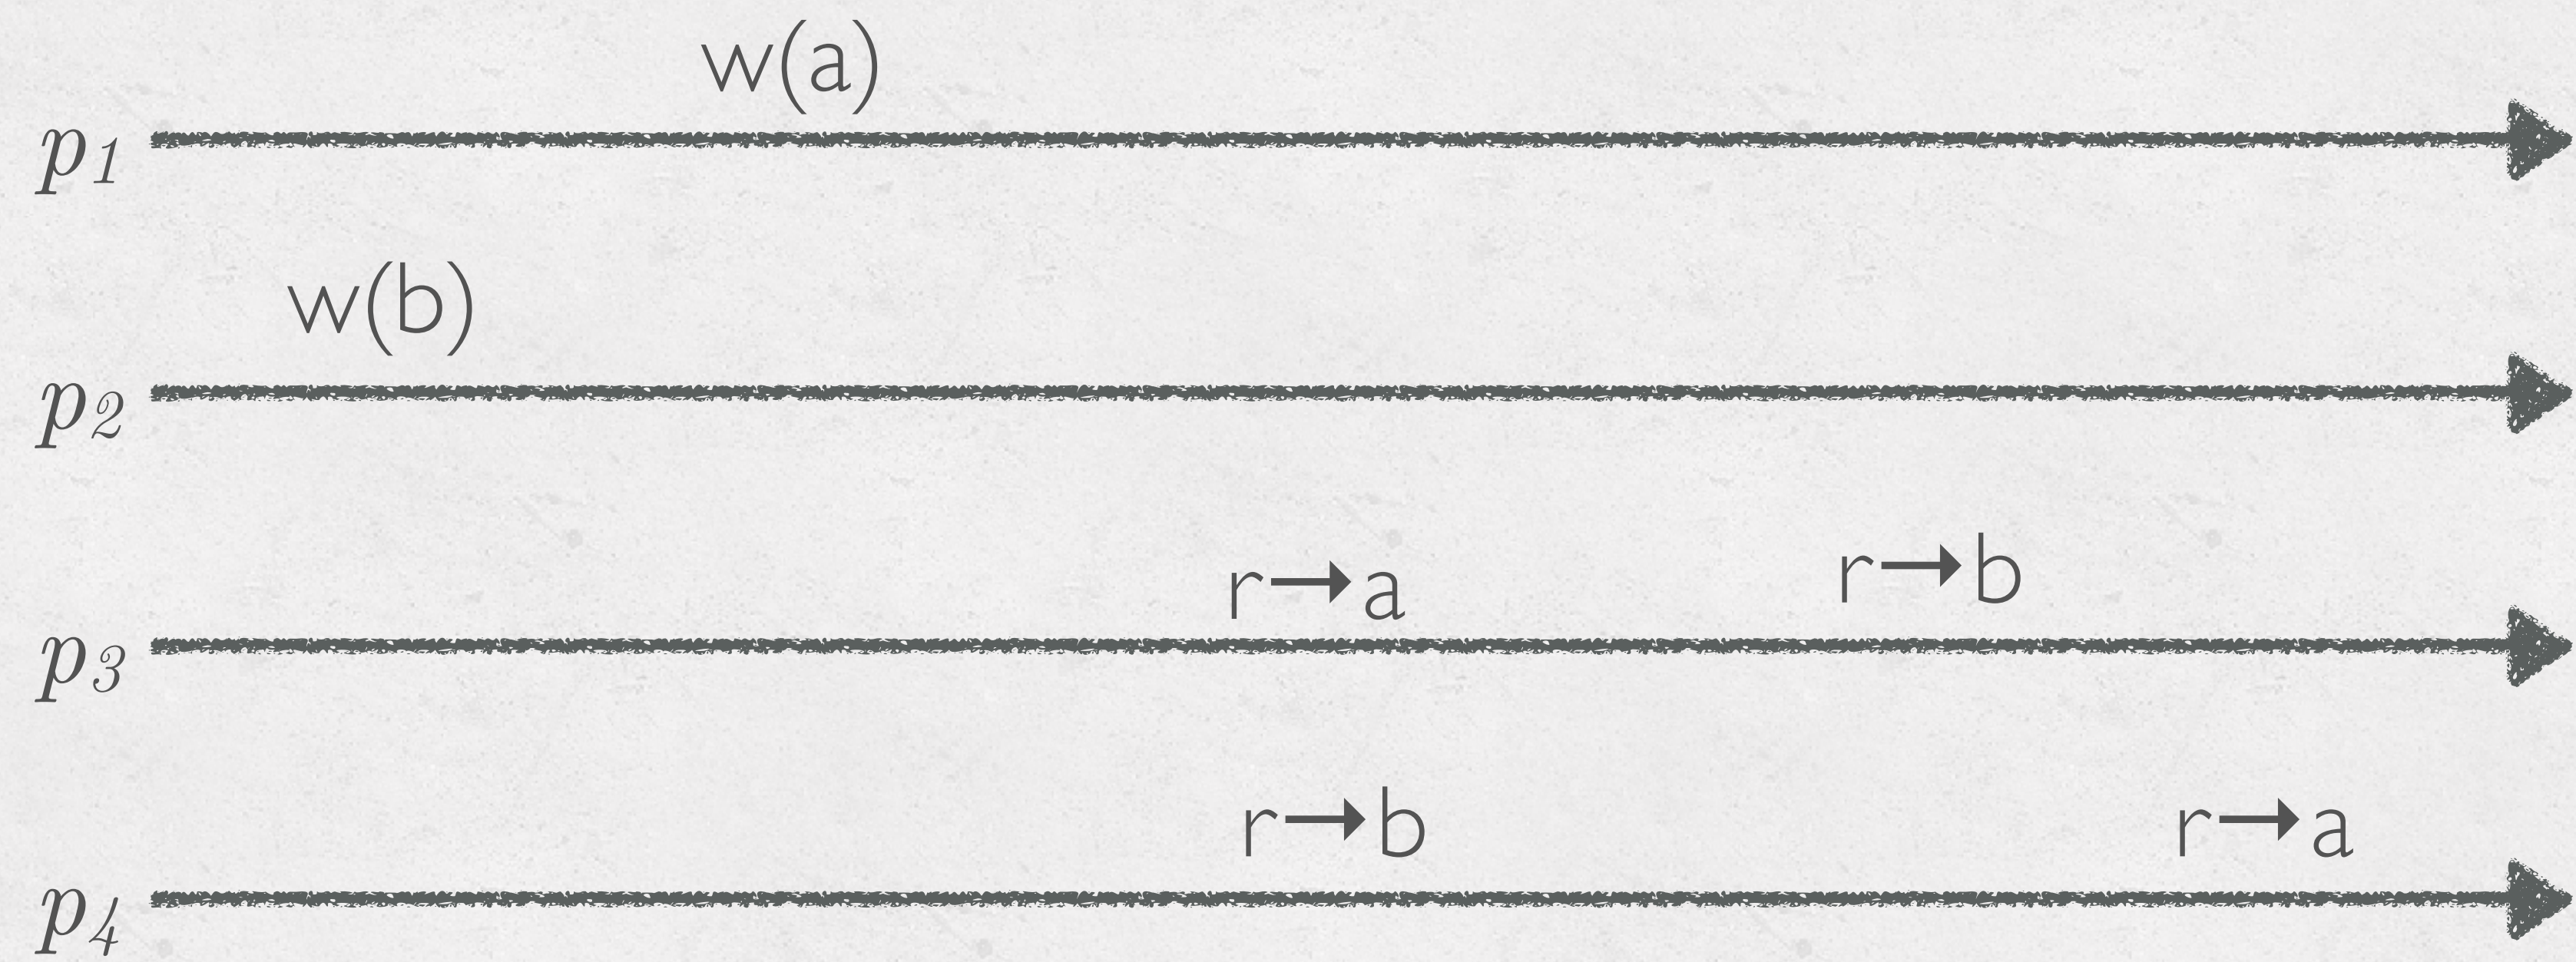


# CAUSAL CONSISTENCY

- Writes that are not concurrent (i.e., writes related by the happens-before relation) must be seen in that order. Concurrent writes can be seen in different orders on different nodes.
- Linearizability implies causal consistency.

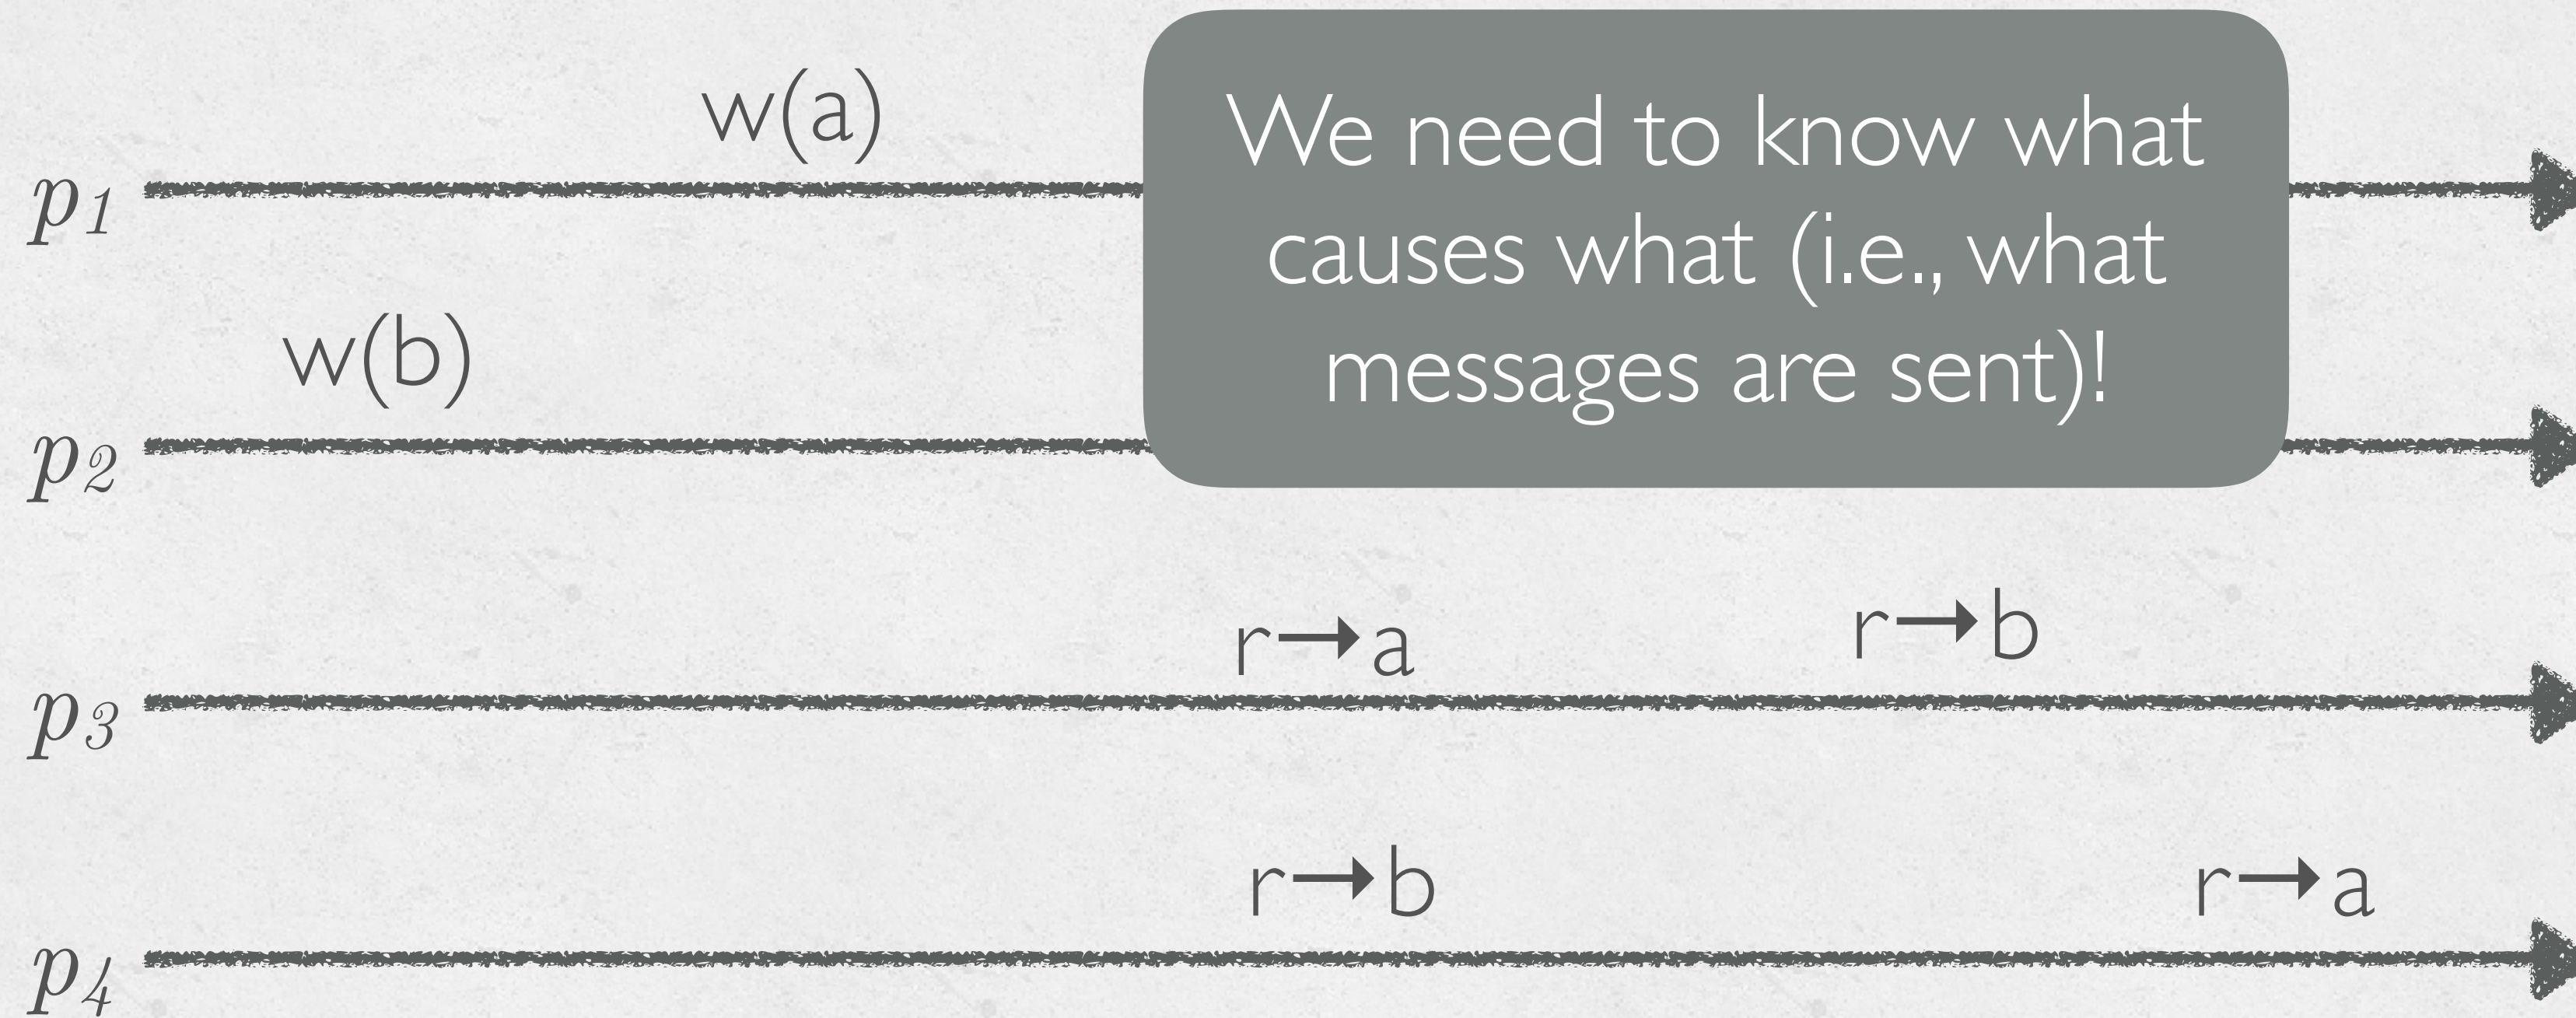


# Is It CAUSAL?






# Is It CAUSAL?






# Is It CAUSAL?






# Is It CAUSAL?






# Is It CAUSAL?





# Is It CAUSAL?





**Cool Theorem:** Causal consistency\* is the strongest form of consistency that can be provided in an always-available convergent system.

Basically, if you want to process writes even in the presence of network partitions and failures, causal consistency is the best you can do.




# WE CAN GET WEAKER!

- **FIFO Consistency:** writes done by the same process are seen in that order; writes to different processes can be seen in different orders. Equivalent to the PRAM model.
- **Eventual Consistency**  $\approx$  if all writes to an object stop, eventually all processes read the same value. (Not even a safety property! "Eventual consistency is no consistency.")




# Is It FIFO?






# Is It FIFO?






# Is It FIFO?





# Is It FIFO?

Not FIFO!





Lamport's register semantics, sequential consistency, linearizability, and causal consistency, and FIFO consistency are all *safety properties*.



# USING CONSISTENCY GUARANTEES

## Thread 1

```
a = 1  
print("b:" + b)
```

## Thread 2

```
b = 1  
print("a:" + a)
```

Initially, both a and b are 0.

What are the possible outputs of this program?



# USING CONSISTENCY GUARANTEES

## Thread 1

```
a = 1  
print("b:" + b)
```

## Thread 2

```
b = 1  
print("a:" + a)
```

Depends on  
memory  
consistency!

Initially, both a and b are 0.

What are the possible outputs of this program?



# USING CONSISTENCY GUARANTEES

## Thread 1

```
a = 1  
print("b:" + b)
```

## Thread 2

```
b = 1  
print("a:" + a)
```

Suppose both prints output 0.



# USING CONSISTENCY GUARANTEES

**Thread 1**

```
a = 1  
print("b:" + b)
```

**Thread 2**

```
b = 1  
print("a:" + a)
```

The diagram shows two threads, Thread 1 and Thread 2. Thread 1 has two lines of code: 'a = 1' and 'print("b:" + b)'. Thread 2 has two lines of code: 'b = 1' and 'print("a:" + a)'. Two arrows are drawn between the threads. One arrow starts from the 'a = 1' line of Thread 1 and points to the 'print("a:" + a)' line of Thread 2. The other arrow starts from the 'b = 1' line of Thread 2 and points to the 'print("b:" + b)' line of Thread 1. These two arrows cross each other, forming an 'X' shape, which represents a cycle in the happens-before graph.

Suppose both prints output 0.

Then there's a cycle in the happens-before graph.  
Not sequential!



# ASIDE: JAVA'S MEMORY MODEL

- Java is **not** sequentially consistent!
- It guarantees sequential consistency only when the program is *data-race free*.
- A **data-race** occurs when two threads access the same memory location concurrently, one of the accesses is a write, and the accesses are not protected by locks (or monitors etc.).



# A COMMON (INCORRECT) IDIOM

```
class Foo {  
    private Bar bar = null;  
  
    public void baz() {  
        if (bar == null) {  
            synchronized(this) {  
                if (bar == null) {  
                    bar = new Bar();  
                }  
            }  
        }  
        bar.doAThing();  
    }  
}
```



# A COMMON (INCORRECT) IDIOM, CORRECTED

```
class Foo {  
    private volatile Bar bar = null;  
  
    public void baz() {  
        if (bar == null) {  
            synchronized(this) {  
                if (bar == null) {  
                    bar = new Bar();  
                }  
            }  
        }  
        bar.doAThing();  
    }  
}
```

**volatile** = accesses are  
sequentially consistent



# A COMMON (INCORRECT) IDIOM, CORRECTED

```
class Foo {  
    private volatile Bar bar = null;
```

```
    public void baz()  
    {  
        if (bar == null)  
            synchronized (this)  
            {  
                if (bar == null)  
                    bar = new Bar();  
            }  
        }  
        bar.doAThing();  
    }  
}
```

Reminder: you don't need to worry about multi-threaded access for the labs! **volatile** accesses are sequentially consistent (except not grabbing locks in equals and hashCode)



# How to Use Weak Consistency?

- Separate operations with stronger semantics, weak consistency (and high performance) by default
- Application-level protocols, either using separate communication, or extra synchronization variables in the data store (not always possible)



# MAIN TAKEAWAYS

- The weaker the consistency model, the harder it is to program against (usually).
- The stronger the model, the harder it is to enforce (again, usually).