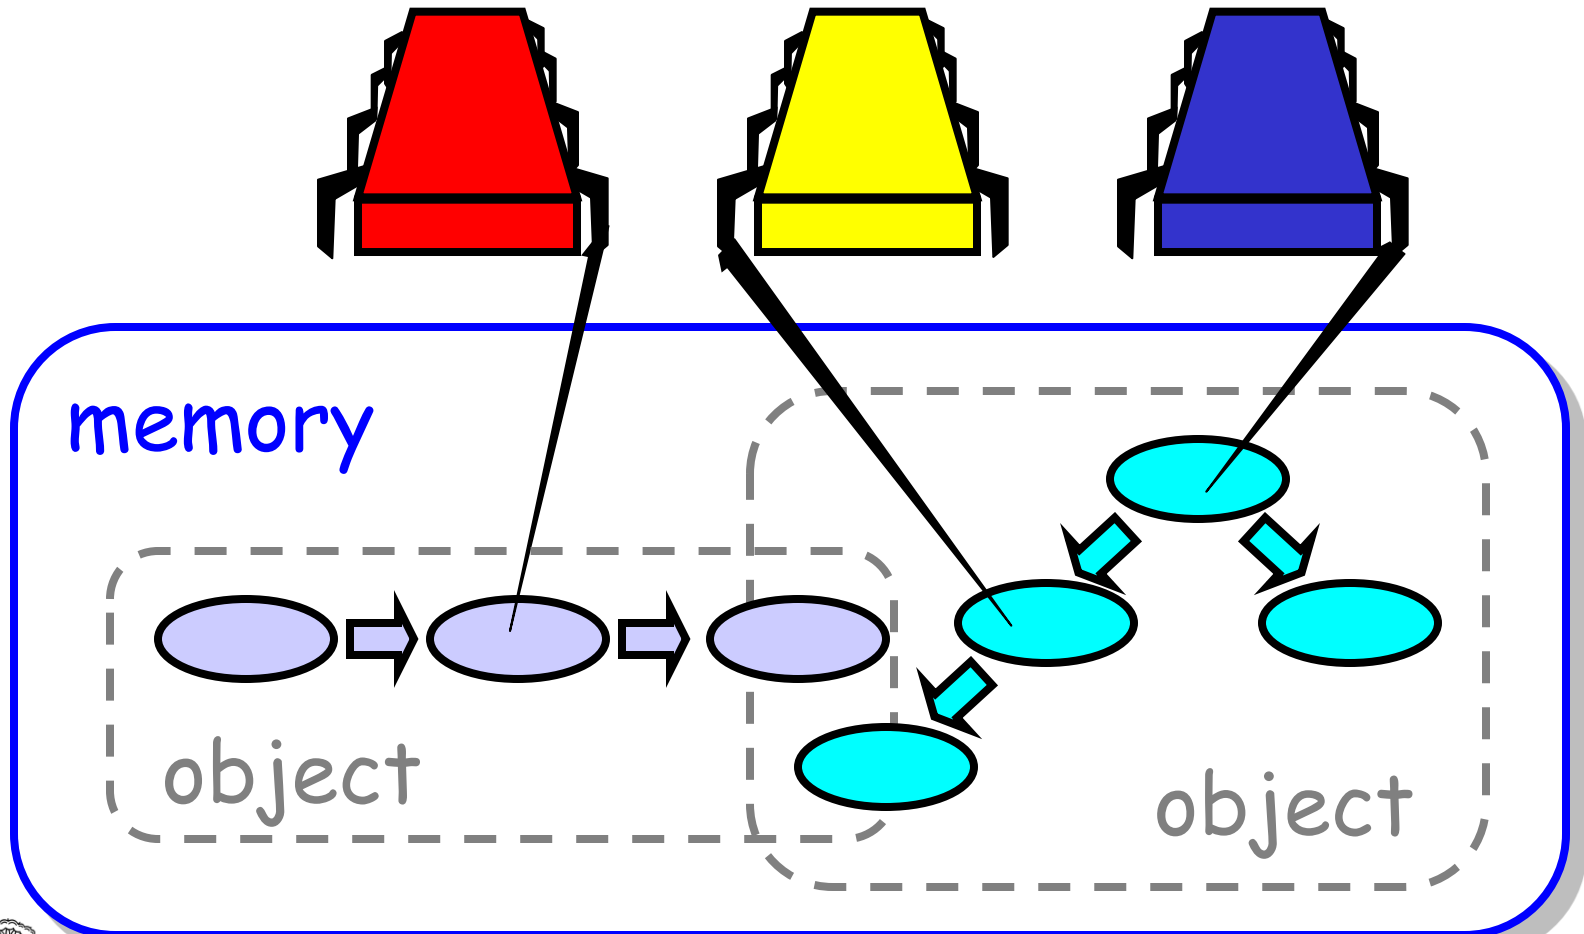


# Concurrent Objects

Companion slides for  
The Art of Multiprocessor  
Programming  
by Maurice Herlihy & Nir Shavit

# Concurrent Computaton




# Objectivism

- What is a concurrent object?
  - How do we describe one?
  - How do we implement one?
  - How do we tell if we're right?


# Objectivism

- What is a concurrent object?
  - How do we describe one?
  - How do we tell if we're right?

# FIFO Queue: Enqueue Method



# FIFO Queue: Dequeue Method



# Implementation: Deq

```
public class Queue<T> {  
  
    int head = 0, tail = 0;  
    T[QSIZE] items;  
  
    public synchronized T deq() {  
        while (tail - head == 0)  
            this.wait();  
        T result = items[head % QSIZE]; head++;  
        this.notifyAll();  
        return result;  
    }  
    ...  
}}
```



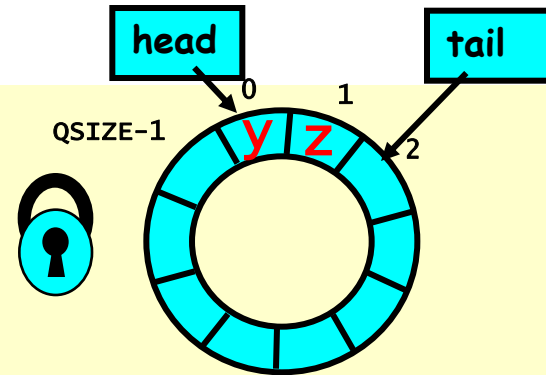
# Implementation: Deq

```
public class Queue<T> {
```

```
    int head = 0, tail = 0;  
    T[QSIZE] items;
```

```
    public synchronized T deq() {  
        while (tail - head == 0)  
            this.wait();  
        T result = items[head % QSIZE]; head++;  
        this.notifyAll();  
        return result;  
    }
```

```
    ...  
}}
```






# Implementation: Deq

```
public class Queue<T> {  
    int head = 0, tail = 0;  
    T[QSIZE] items;  
  
    public synchronized T deq() {  
        while (tail - head == 0)  
            this.wait();  
        T result = items[head % QSIZE]; head++;  
        this.notifyAll();  
        return result;  
    }  
    ...  
}
```

Method calls  
mutually exclusive



# Implementation: Deq

```
public class Queue<T> {  
    int head = 0, tail = 0; Is queue empty?  
    T[QSIZE] items;  
  
    public synchronized T deq() {  
        while (tail - head == 0)  
            this.wait();  
        T result = items[head % QSIZE]; head++;  
        this.notifyAll();  
        return result;  
    }  
    ...  
}
```



# Implementation: Deq

```
public class Queue<T> {  
    int head = 0, tail = 0;  
    T[QSIZE] items;  
  
    public synchronized T deq() {  
        while (tail - head == 0)  
            this.wait();  
        T result = items[head % QSIZE]; head++;  
        this.notifyAll();  
        return result;  
    }  
    ...  
}}
```

Release lock if  
need to wait



# Implementation: Deq

```
public class Queue<T> {  
  
    int head = 0, tail = 0;  
    T[QSIZE] items;  
  
    public synchronized T deq() {  
        while (tail - head == 0)  
            this.wait();  
        T result = items[head % QSIZE]; head++;  
        this.notifyAll();  
        return result;  
    }  
    ...  
}}
```

actual update



# Implementation: Deq

```
public class Queue<T> {  
    int head = 0, tail = 0;  
    T[QSIZE] items;  
  
    public synchronized T deq() {  
        while (tail - head == 0)  
            this.wait();  
        T result = items[head % QSIZE]; head++;  
        this.notifyAll();  
        return result;  
    }  
    ...  
}
```

Notify waiting threads  
that you put something  
in the queue



# Implementation: Deq

```
public class Queue<T> {  
  
    int head = 0, tail = 0;  
    T[QSIZE] items;  
  
    public synchronized T deq() {  
        while (tail - head == 0)  
            this.wait();  
        T result = items[head % QSIZE]; head++;  
        this.notifyAll();  
        return result;  
    }  
    ...  
}}
```

Should be correct because modifications are mutually exclusive...



# Now consider the following implementation

- The same thing without mutual exclusion
- For simplicity, only two threads
  - One thread enq only
  - The other deq only

# Lock-free 2-Thread Queue

```
public class LockFreeQueue {  
  
    int head = 0, tail = 0;  
    Item[QSIZE] items;  
  
    public void enq(Item x) {  
        while (tail-head == QSIZE); // busy-wait  
        items[tail % QSIZE] = x; tail++;  
    }  
    public Item deq() {  
        while (tail == head); // busy-wait  
        Item item = items[head % QSIZE]; head++;  
        return item;  
    }  
}
```



# Lock-free 2-Thread Queue


```
public class LockFreeQueue {
```

```
    int head = 0, tail = 0;  
    Item[QSIZE] items;
```

```
    public void enq(Item x) {  
        while (tail-head == QSIZE); // busy-wait  
        items[tail % QSIZE] = x; tail++;  
    }
```


```
    public Item deq() {  
        while (tail == head); // busy-wait  
        Item item = items[head % QSIZE]; head++;  
        return item;  
    }
```

```
}}
```



# Lock-free 2-Thread Queue

```
public class LockFreeQueue {  
    int head = 0, tail = 0;  
    Item[QSIZE] items;  
  
    public void enq(Item x) {  
        while (tail-head == QSIZE); // busy-wait  
        items[tail % QSIZE] = x; tail++;  
    }  
  
    public Item deq() {  
        while (tail == head);  
        Item item = items[head];  
        return item;  
    }  
}
```



`items[tail % QSIZE] = x; tail++;`

Queue is up at a lock!

How do we define "correct" when modifications are not exclusive?

# Defining correct concurrent queue implementations

- Need a way to specify a concurrent queue object
- Need a way to prove that an algorithm implements the object's specification
- Lets talk about object specifications
- ...

# Sequential Objects

- Each object has a **state**
  - Usually given by a set of **fields**
  - Queue example: sequence of items
- Each object has a set of **methods**
  - Only way to manipulate state
  - Queue example: **enq** and **deq** methods

# Sequential Specifications

- If (precondition)
  - the object is in such-and-such a state
  - before you call the method,
- Then (postcondition)
  - the method will return a particular value
  - or throw a particular exception.
- and (postcondition, con't)
  - the object will be in some other state
  - when the method returns,

# Pre and PostConditions for Dequeue

- Precondition:
  - Queue is non-empty
- Postcondition:
  - Returns first item in queue
- Postcondition:
  - Removes first item in queue

# Pre and PostConditions for Dequeue

- Precondition:
  - Queue is empty
- Postcondition:
  - Throws Empty exception
- Postcondition:
  - Queue state unchanged

# Why Sequential Specifications Totally Rock


- Interactions among methods captured by side-effects on object state
  - State meaningful between method calls
- Documentation size linear in number of methods
  - Each method described in isolation
- Can add new methods
  - Without changing descriptions of old methods




# What About Concurrent Specifications ?

- Methods?
- Documentation?
- Adding new methods?

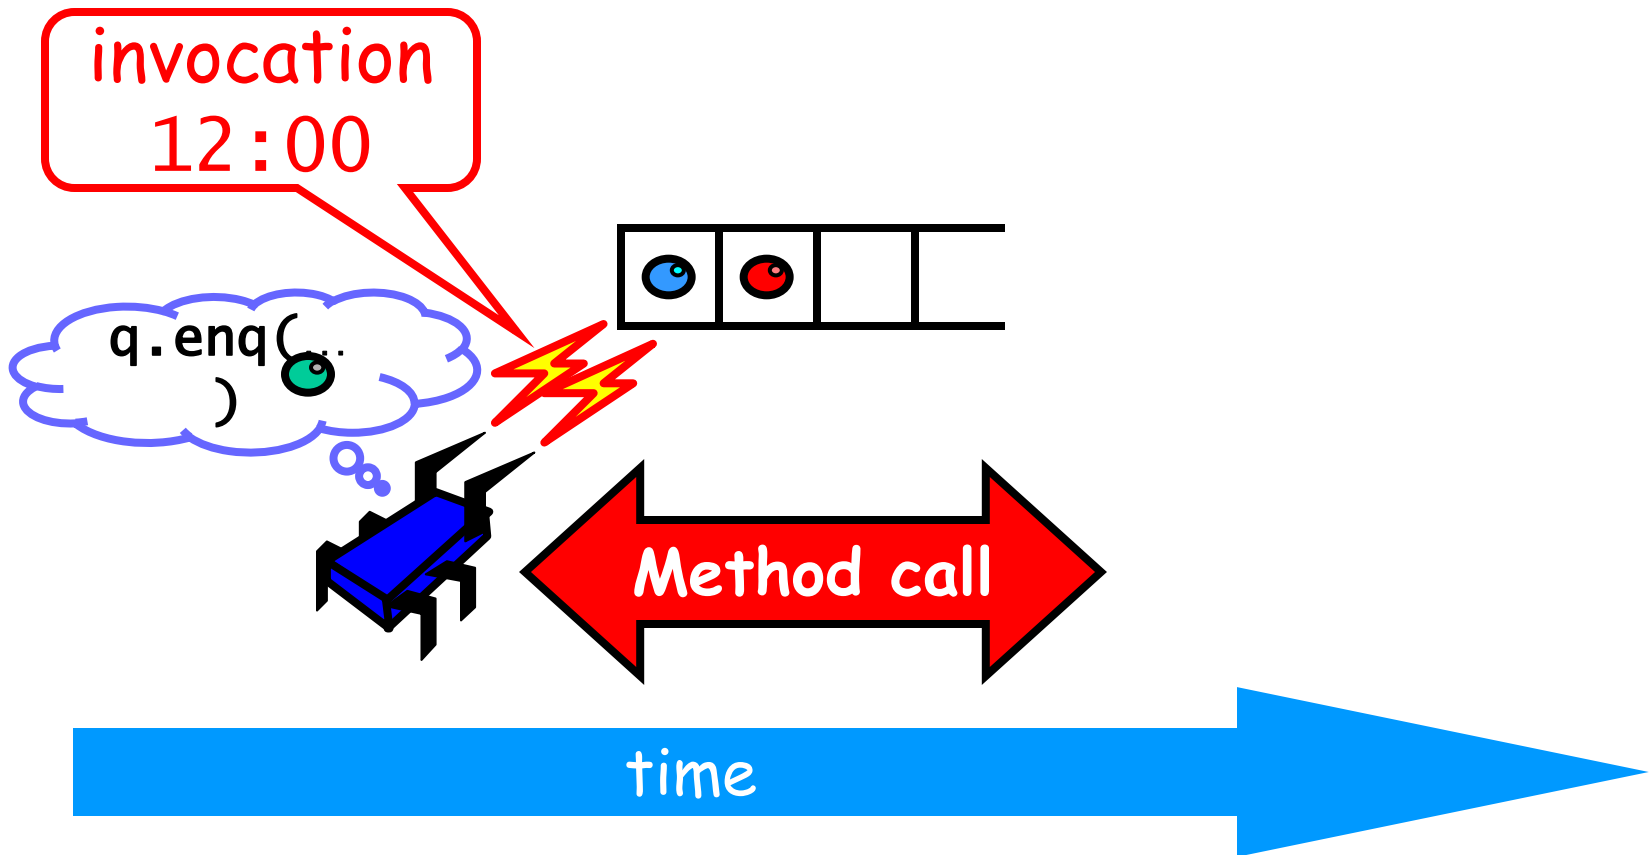
# Methods Take Time



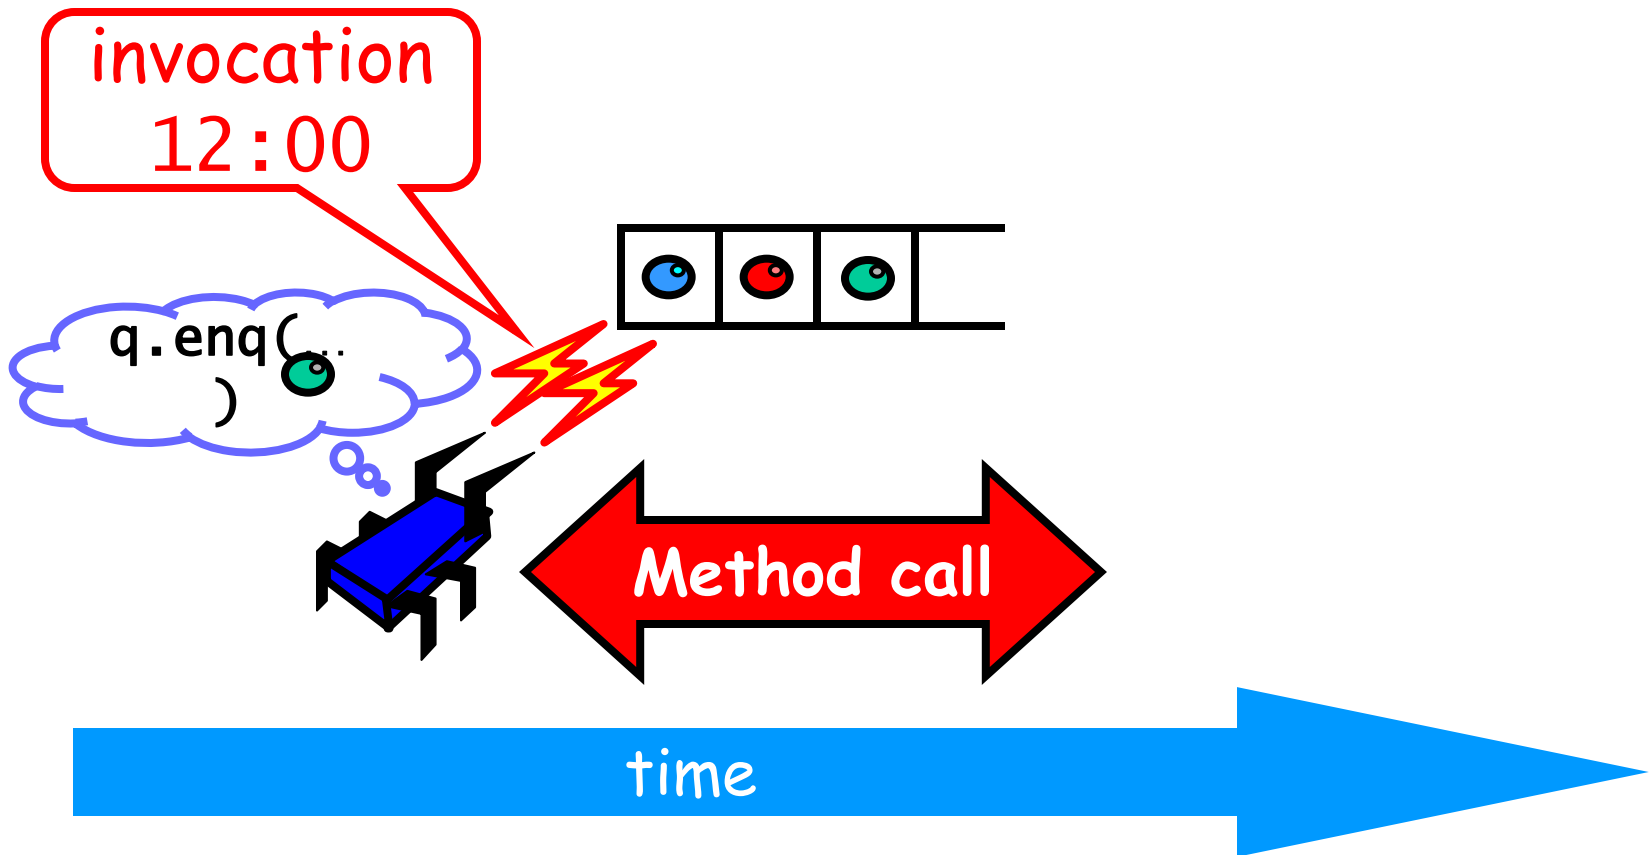
# Methods Take Time



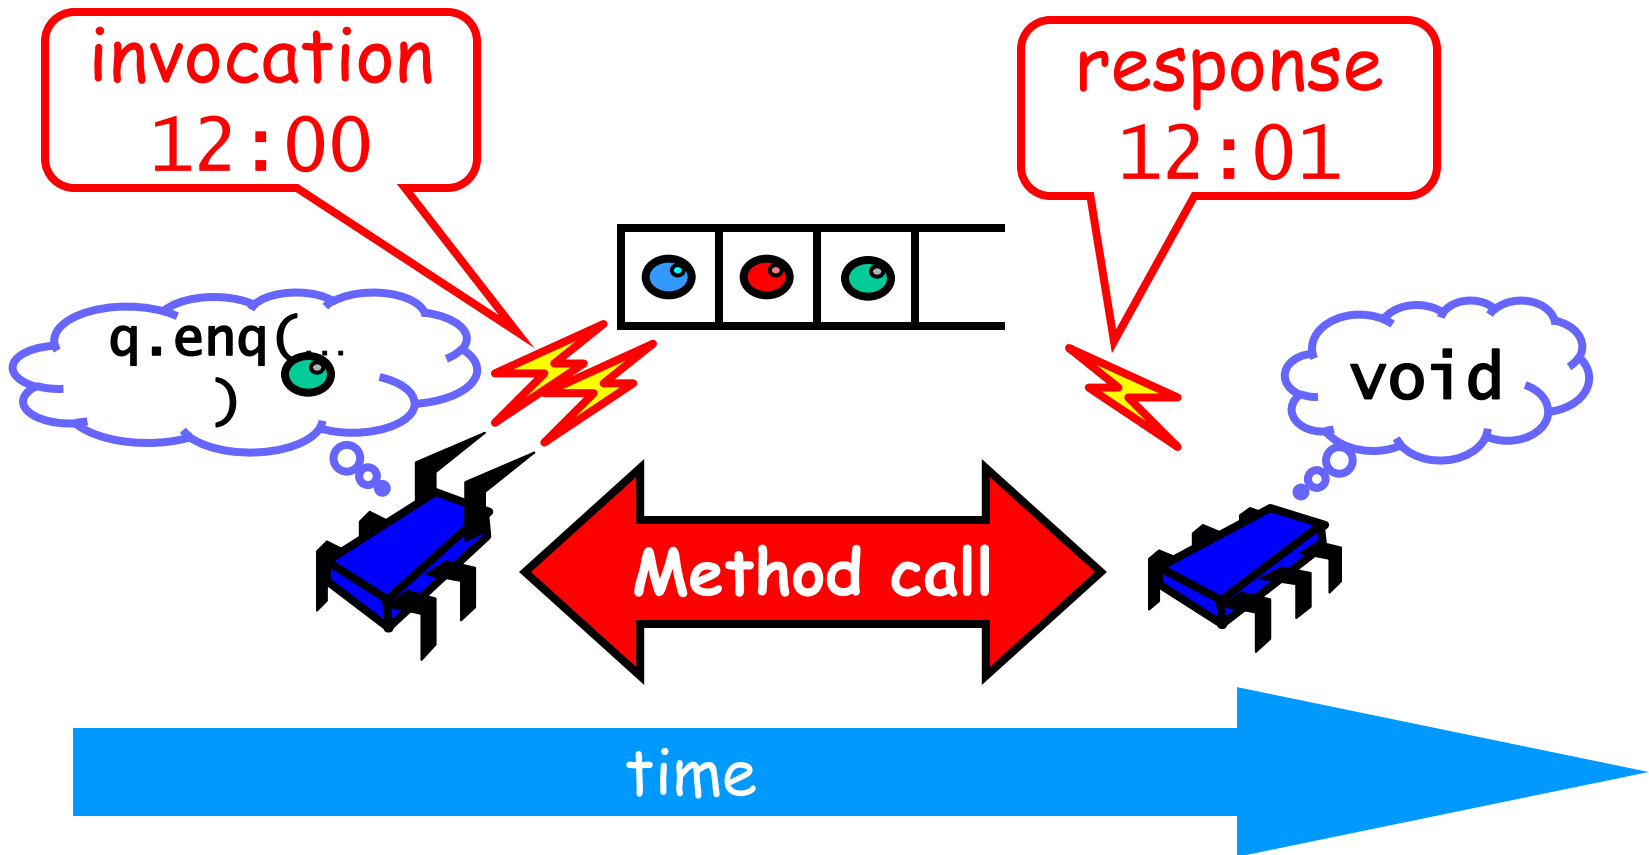
# Methods Take Time



# Methods Take Time




# Methods Take Time



# Sequential vs Concurrent

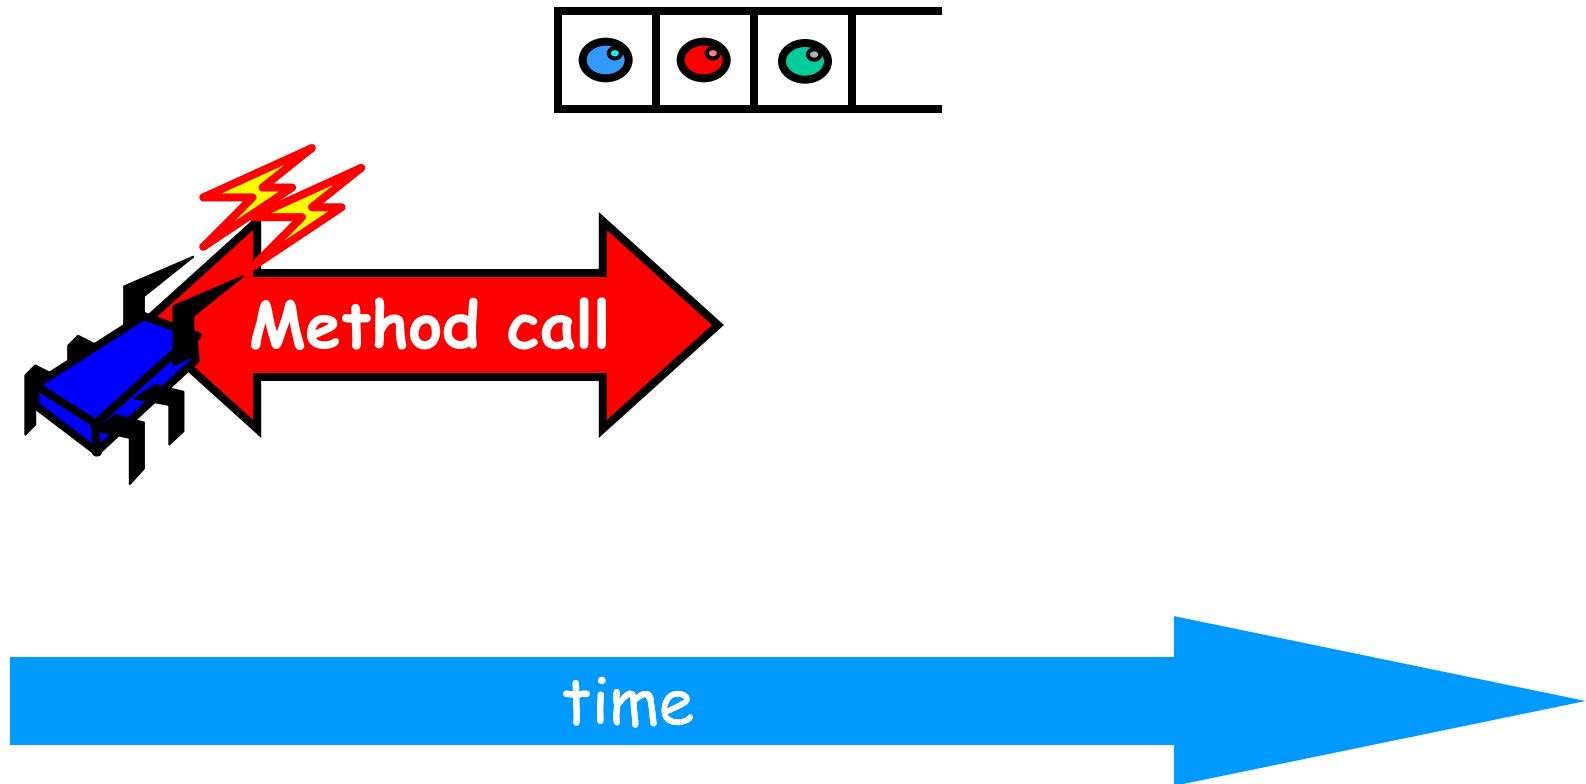
- Sequential
  - Methods take time? Who knew?
- Concurrent
  - Method call is not an event
  - Method call is an interval.

# Concurrent Methods Take Overlapping Time

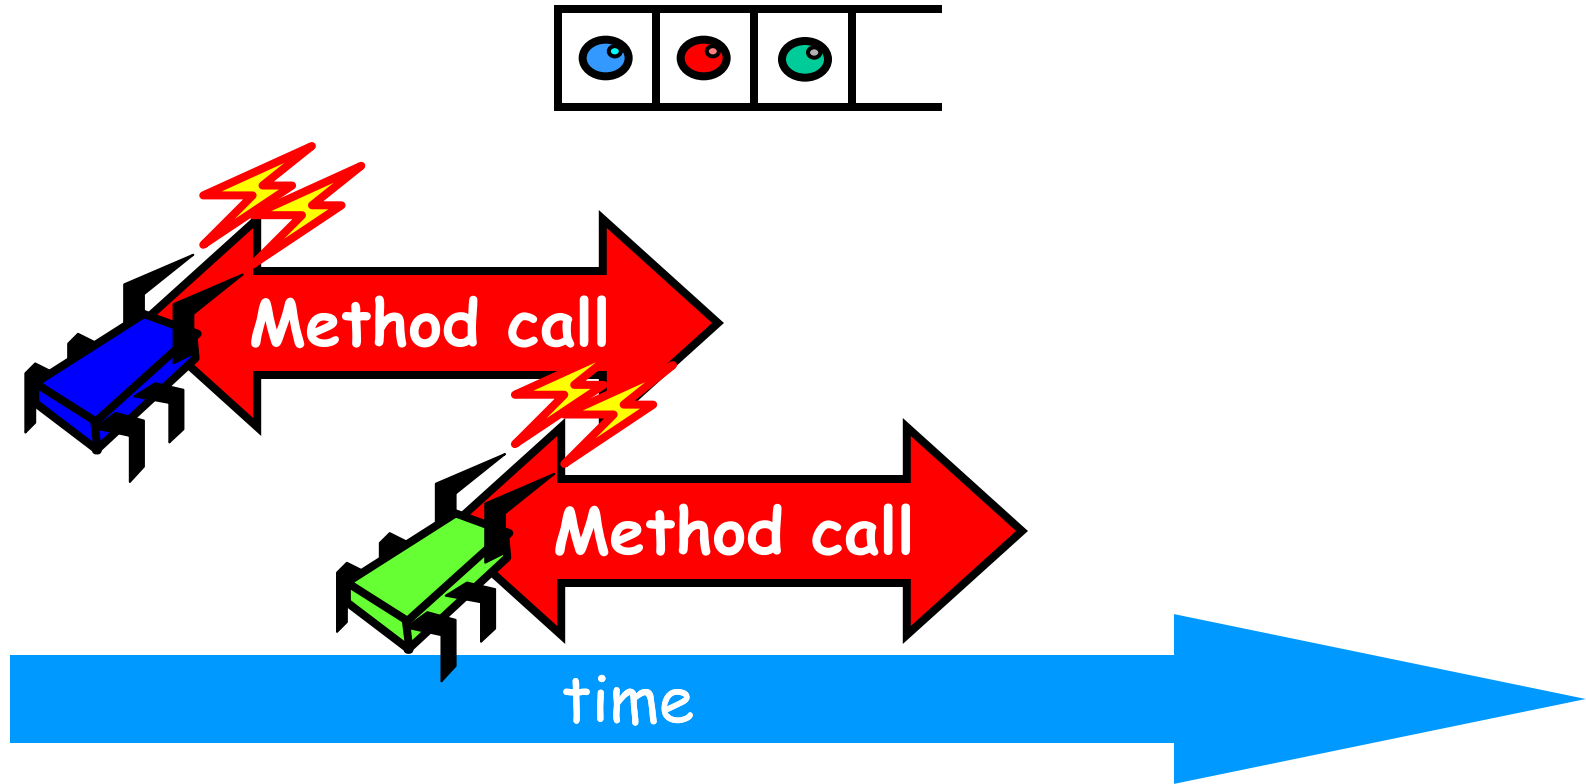





# Concurrent Methods Take Overlapping Time



# Concurrent Methods Take Overlapping Time



# Concurrent Methods Take Overlapping Time



# Sequential vs Concurrent

- Sequential:
  - Object needs meaningful state only between method calls
- Concurrent
  - Because method calls overlap, object might *never* be between method calls

# Sequential vs Concurrent

- Sequential:
  - Each method described in isolation
- Concurrent
  - Must characterize **all** possible interactions with concurrent calls
    - What if two enqs overlap?
    - Two deqs? enq and deq? ...

# Sequential vs Concurrent

- Sequential:
  - Can add new methods without affecting older methods
- Concurrent:
  - Everything can potentially interact with everything else

# Sequential vs Concurrent

- Sequential:
  - Can add new methods without affecting older methods
- Concurrent:
  - Everything can potentially interact with everything else

Panic!

# The Big Question

- What does it **mean** for a *concurrent* object to be correct?
  - What *is* a concurrent FIFO queue?
  - FIFO means strict temporal order
  - Concurrent means ambiguous temporal order



# Intuitively...

```
public class Queue<T> {  
  
    int head = 0, tail = 0;  
    T[QSIZE] items;  
  
    public synchronized void enq(T x) {  
        while (tail - head == QSIZE)  
            this.wait();  
        items[tail % QSIZE] = x; tail++;  
        this.notifyAll();  
    }  
    ...  
}}
```

# Intuitively...

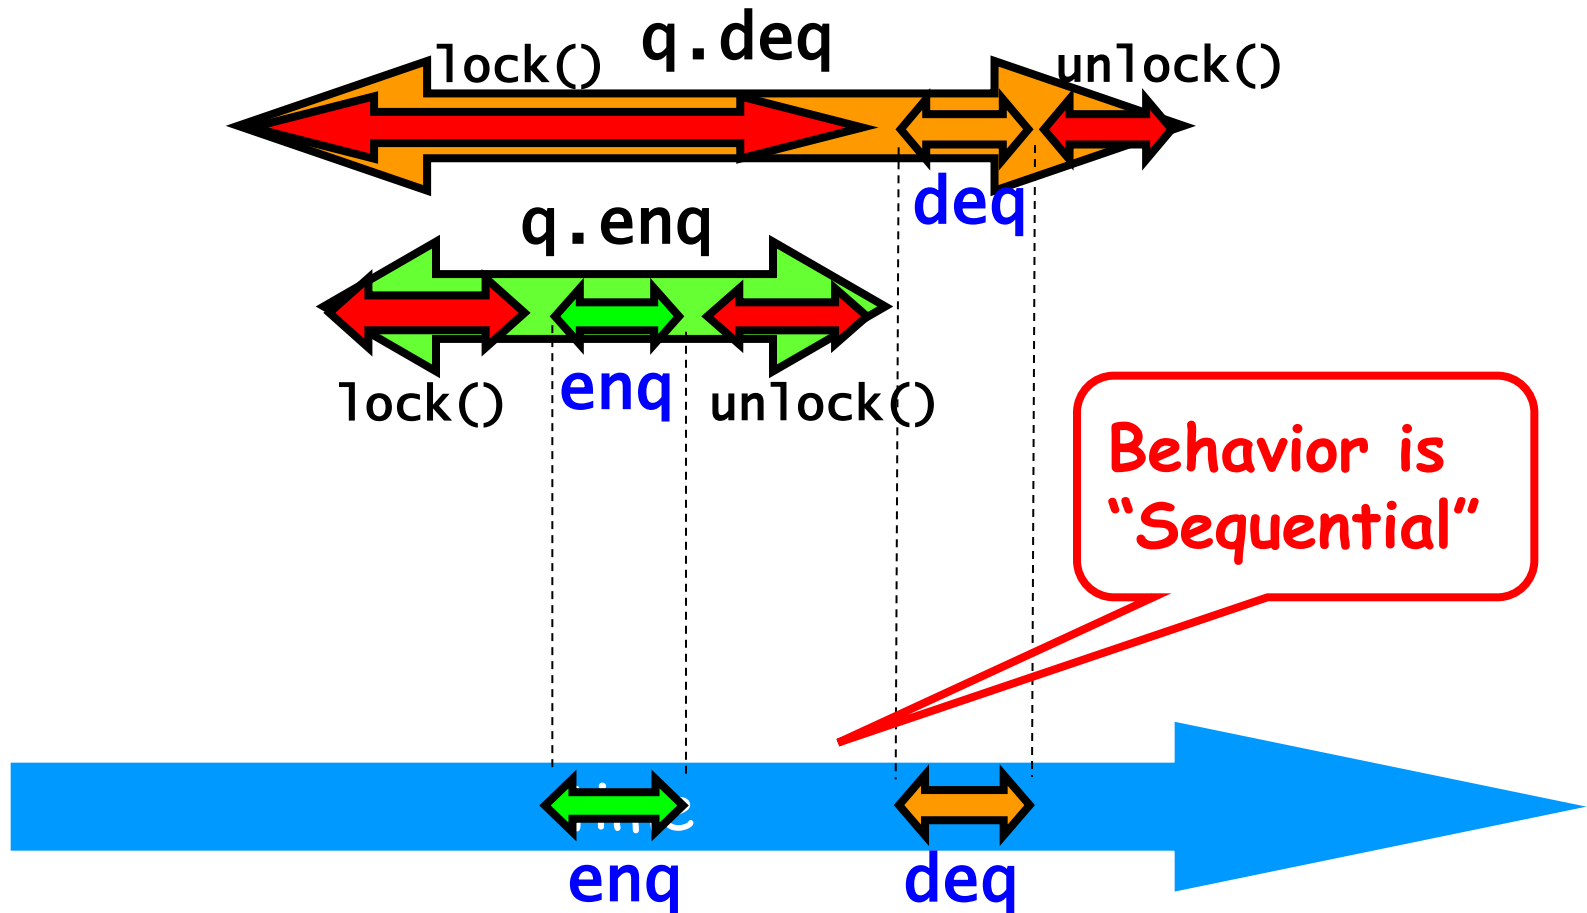
```
public class Queue<T> {  
  
    int head = 0, tail = 0;  
    T[QSIZE] items;  
  
    public synchronized void enq(T x) {  
        while (tail - head == QSIZE)  
            this.wait();  
        items[tail % QSIZE] = x; tail++;  
        this.notifyAll();  
    }  
    ...  
}}
```

Queue is updated while holding lock  
(mutually exclusive)



# Testimonials

Lets capture the idea of describing the concurrent via the sequential




# Linearizability

- Each method should
  - “take effect”
  - Instantaneously
  - Between invocation and response events
- Object is correct if this “sequential” behavior is correct
- Any such concurrent object is
  - **Linearizable**<sup>TM</sup>



# Is it really about the object?

- Each method should
  - “take effect”
  - Instantaneously
  - Between invocation and response events
- Sounds like a property of an execution...
- A linearizable object: one all of whose possible executions are linearizable

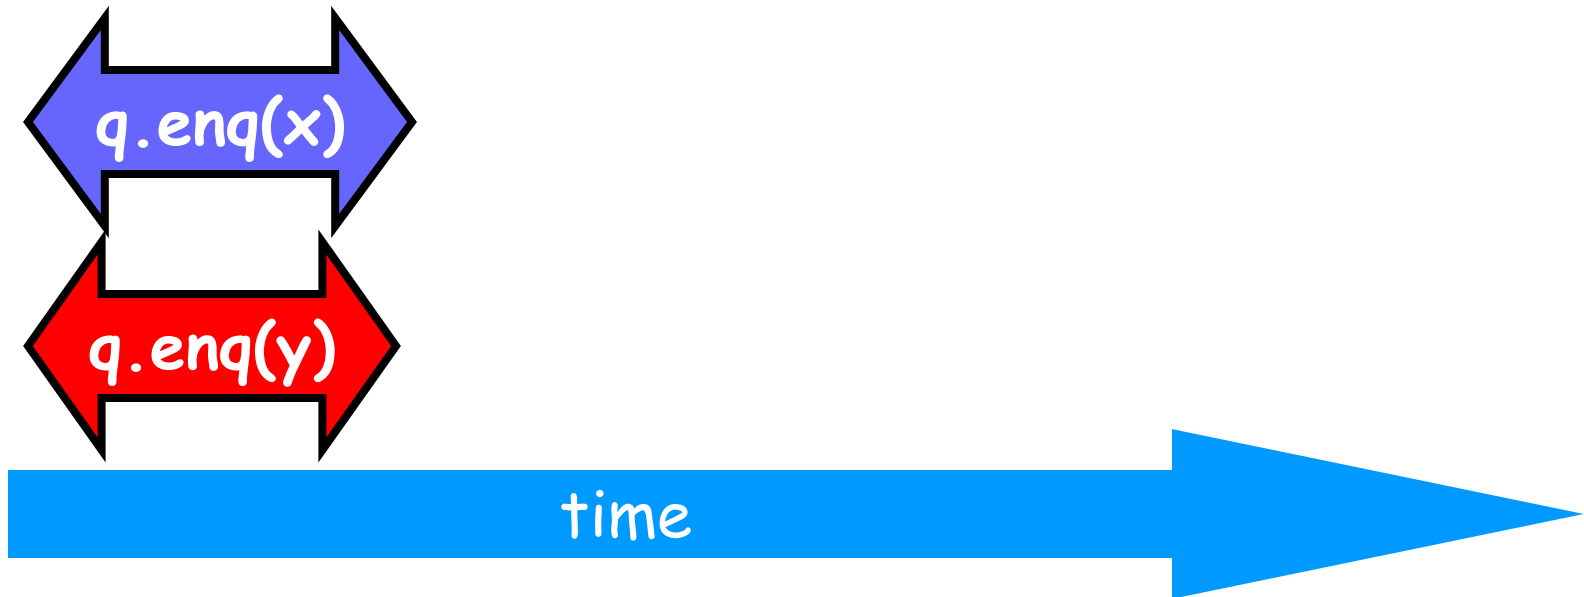

# Example



# Example

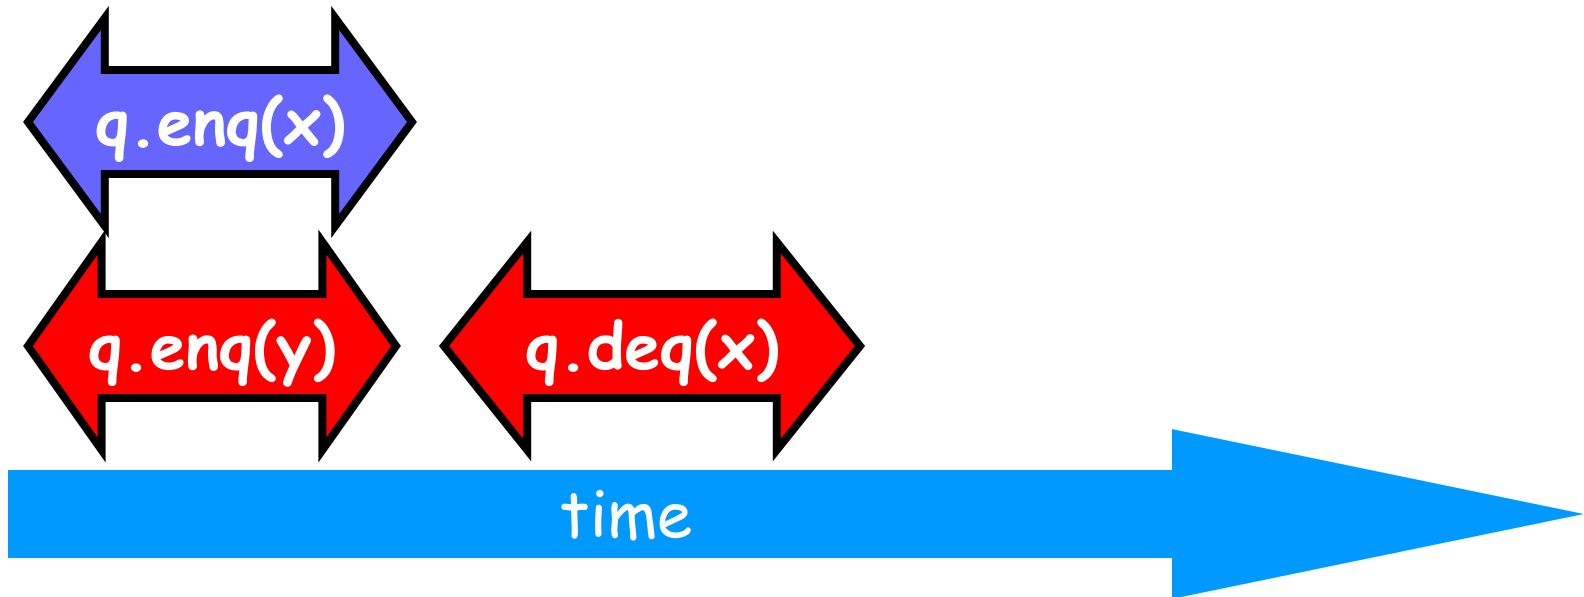



# Example



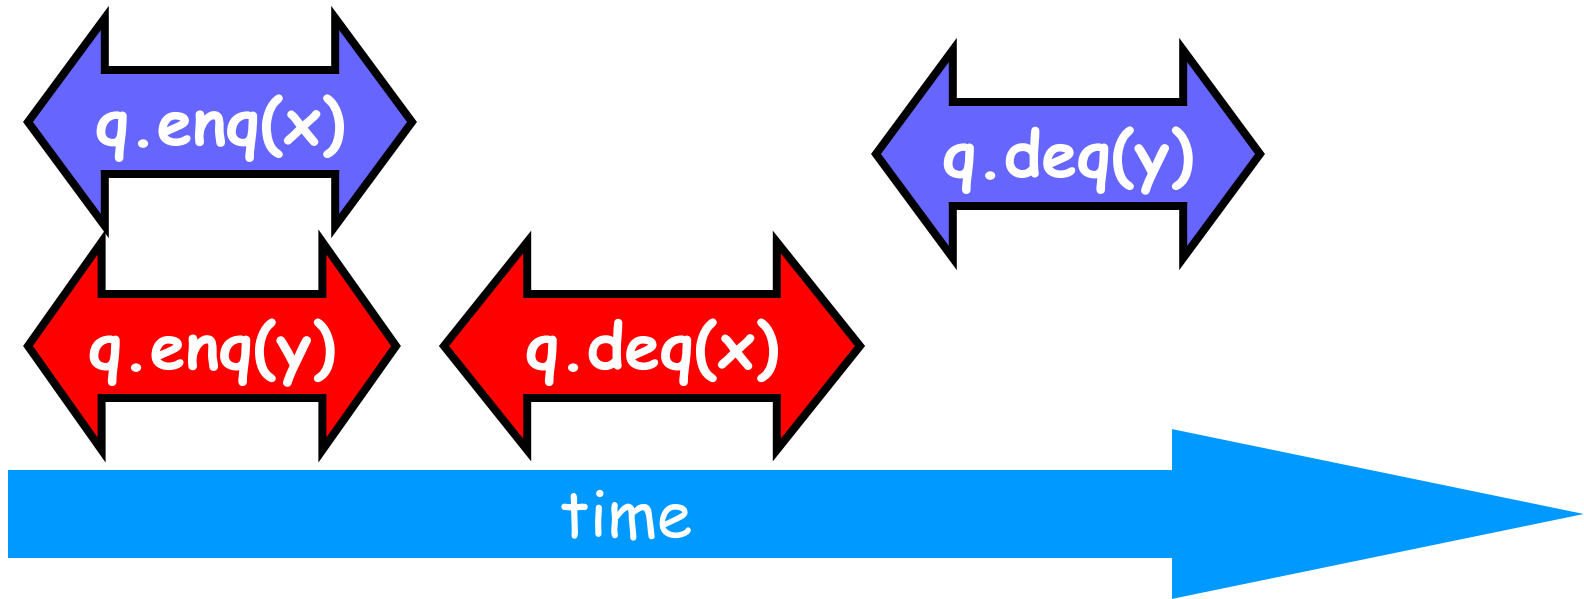



# Example

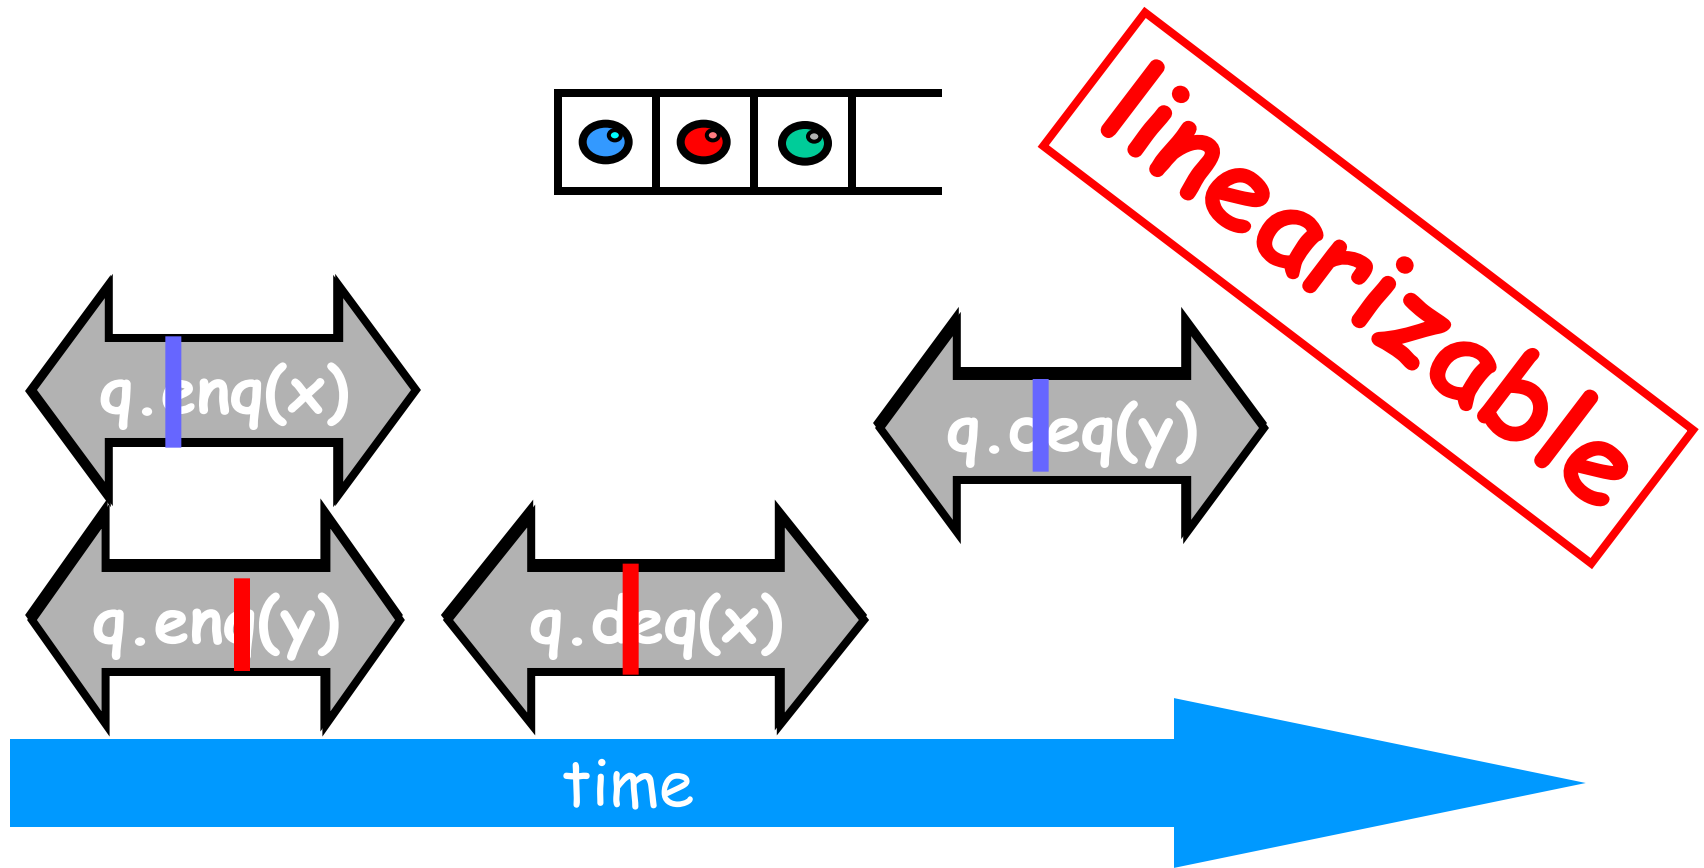




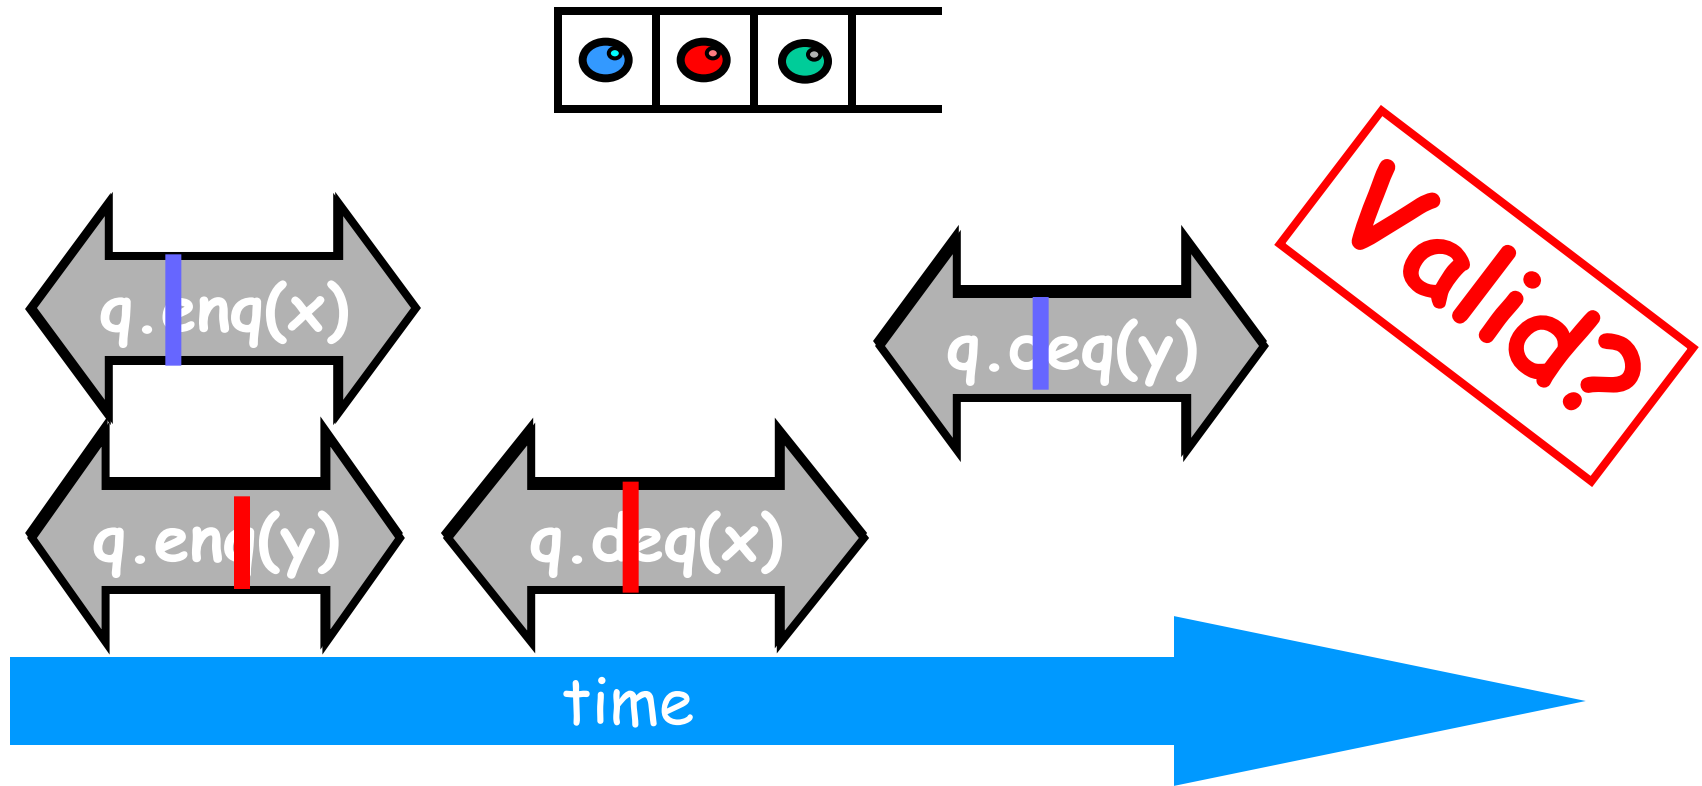
# Example




# Example





# Example






# Example



# Example








# Example









# Example











# Example






# Example



**not linearizable**



# Example

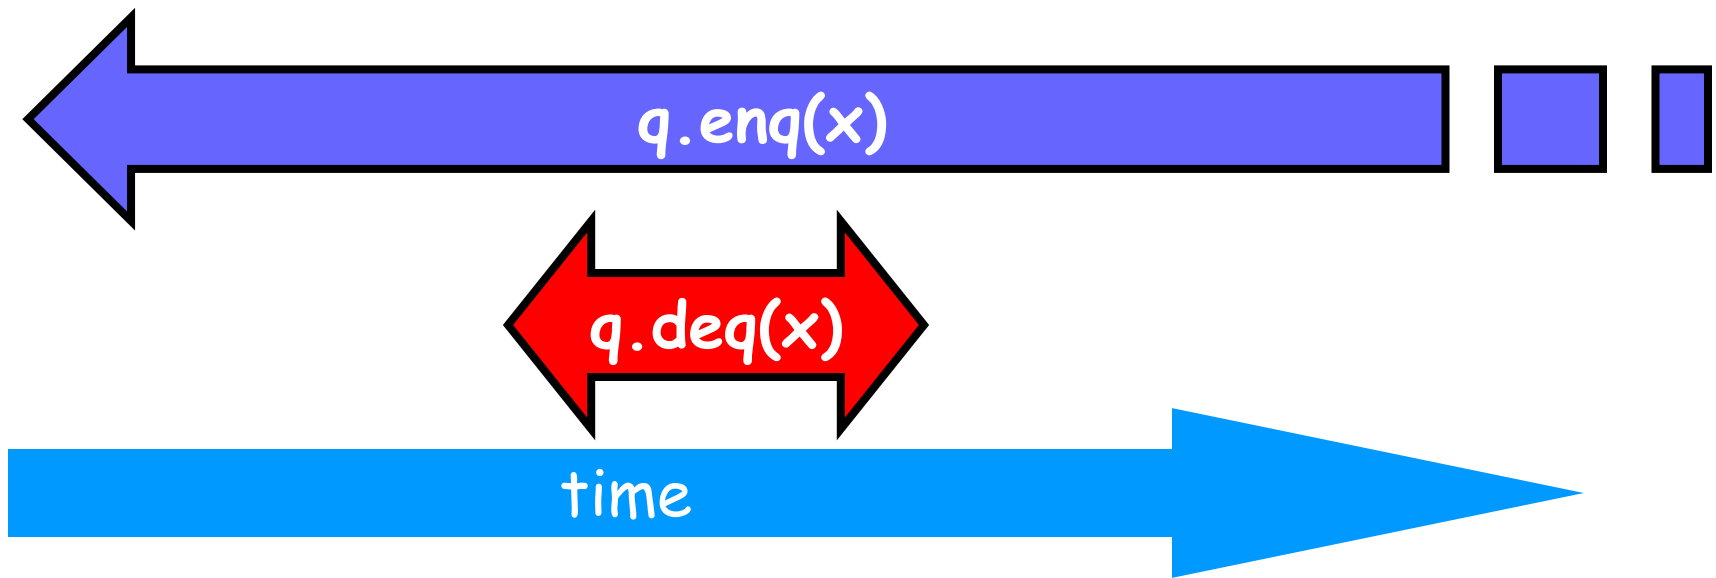



# Example



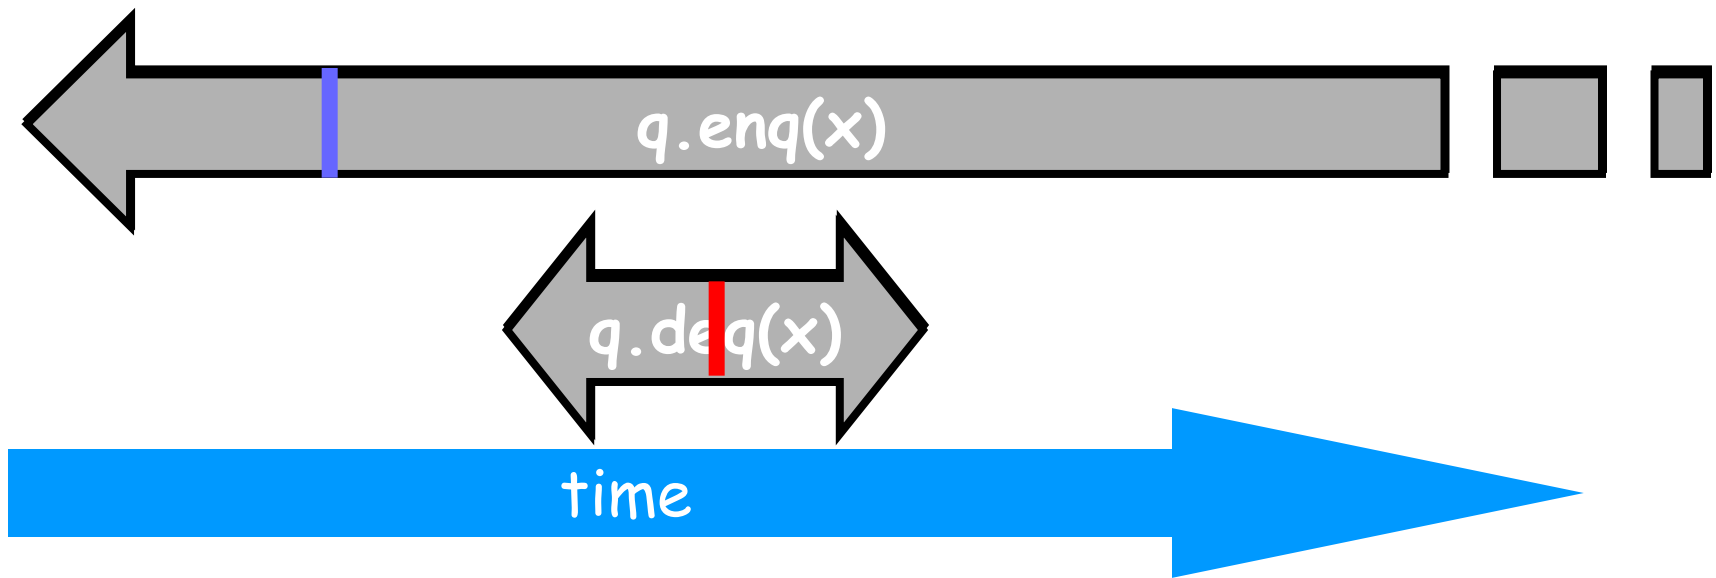



# Example






# Example

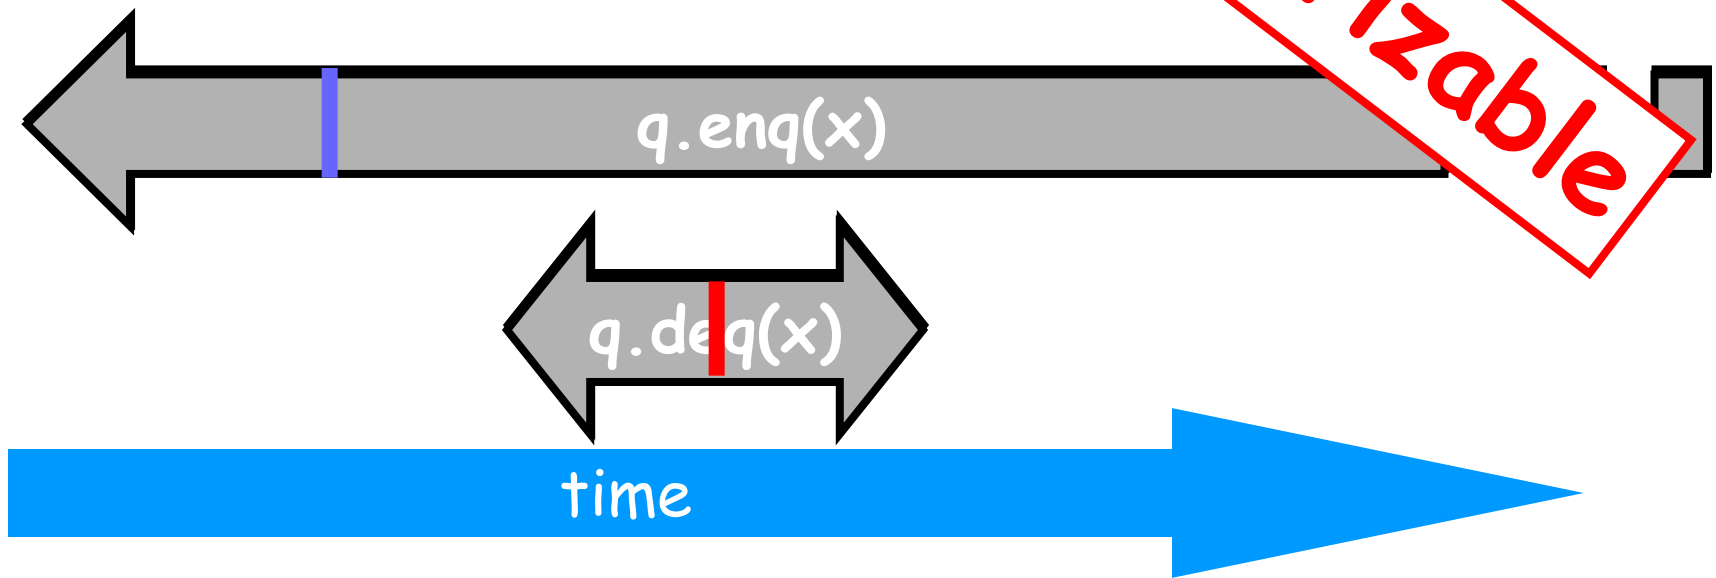






# Example



**linearizable**





# Example









# Example





# Example



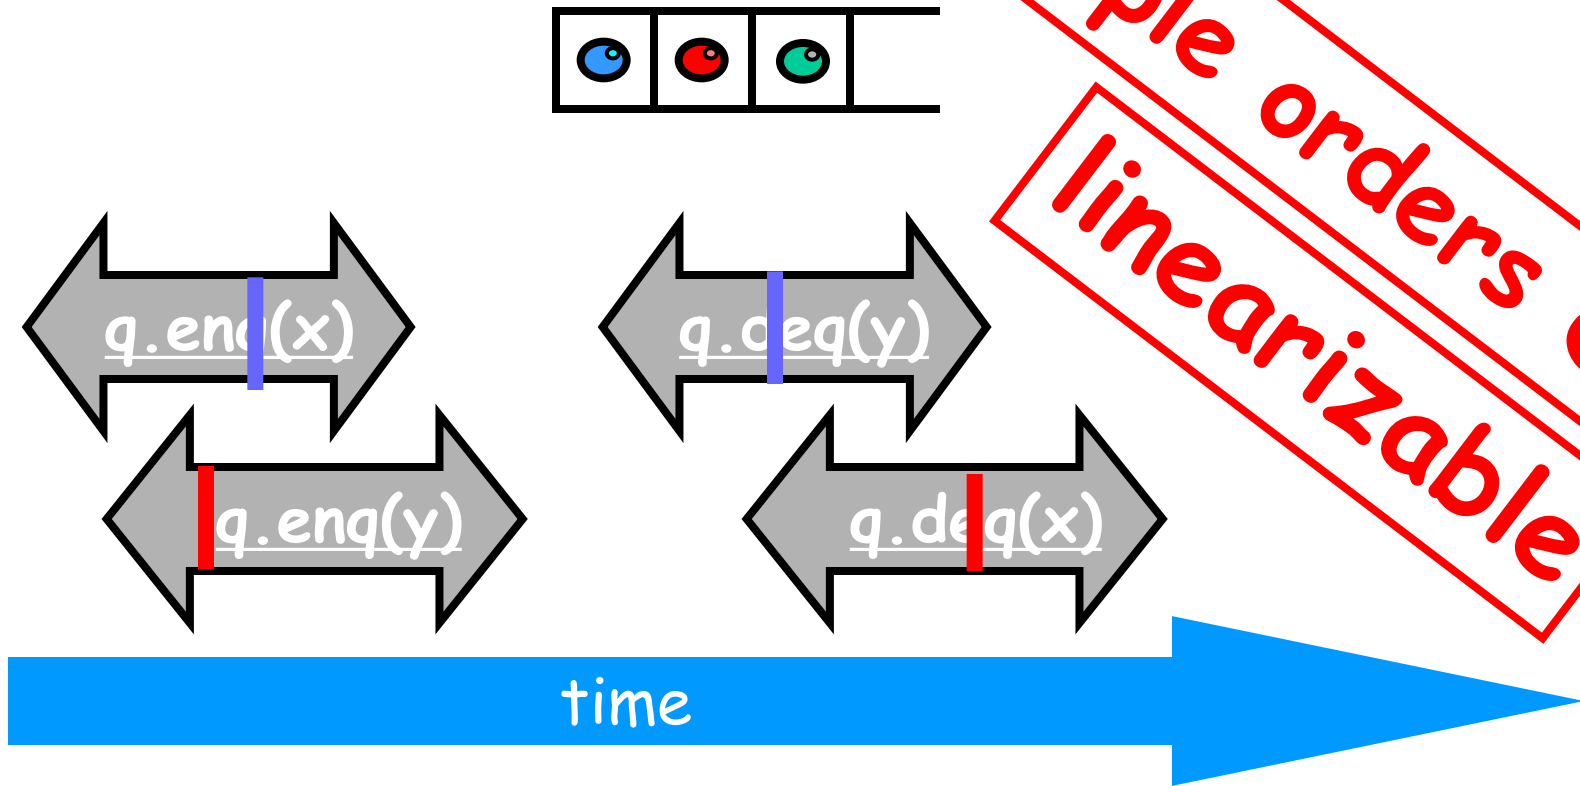


# Example



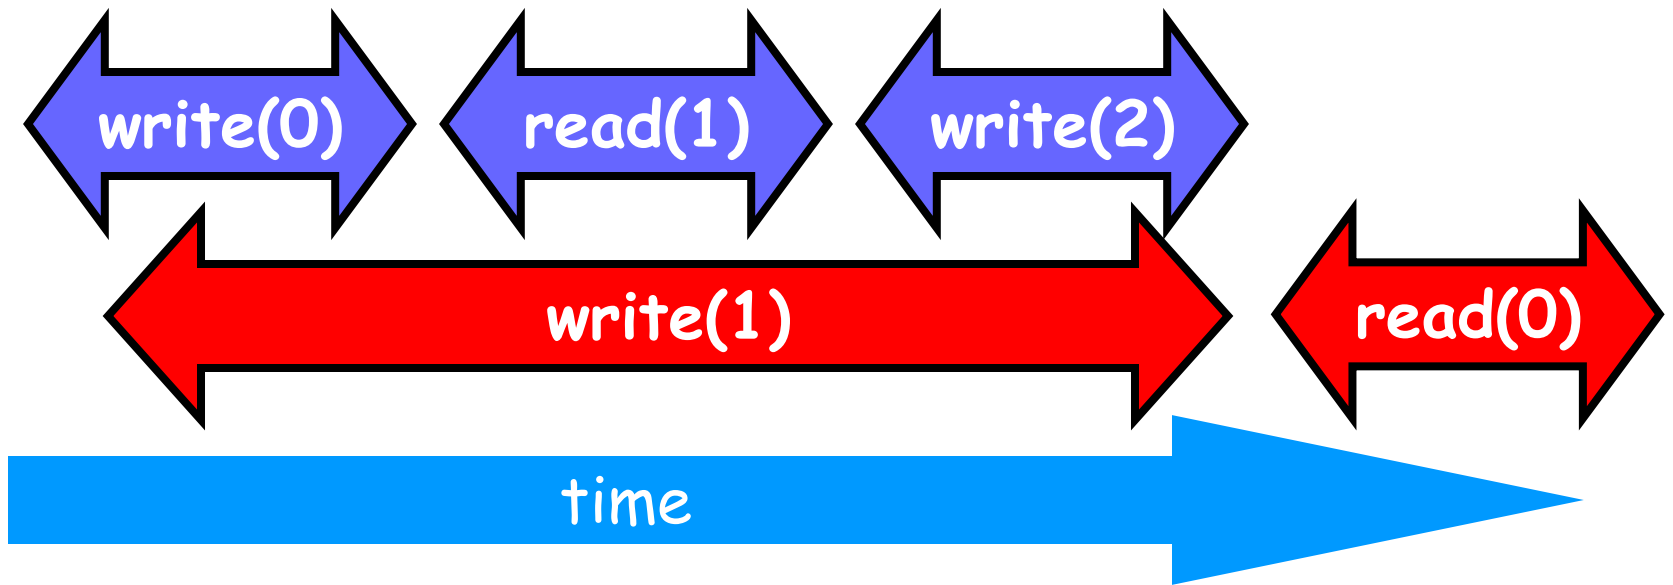
Comme ci  
Comme ça

Example

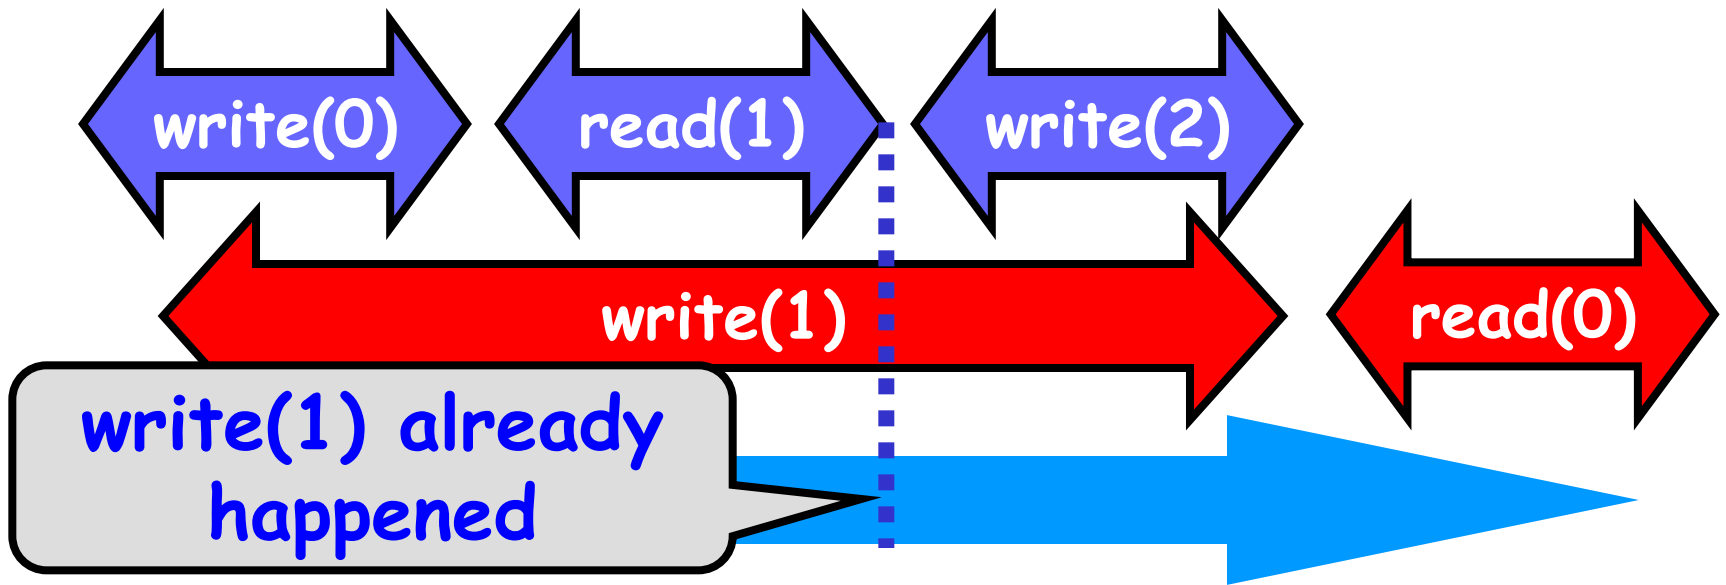


multiple orders OK  
linearizable

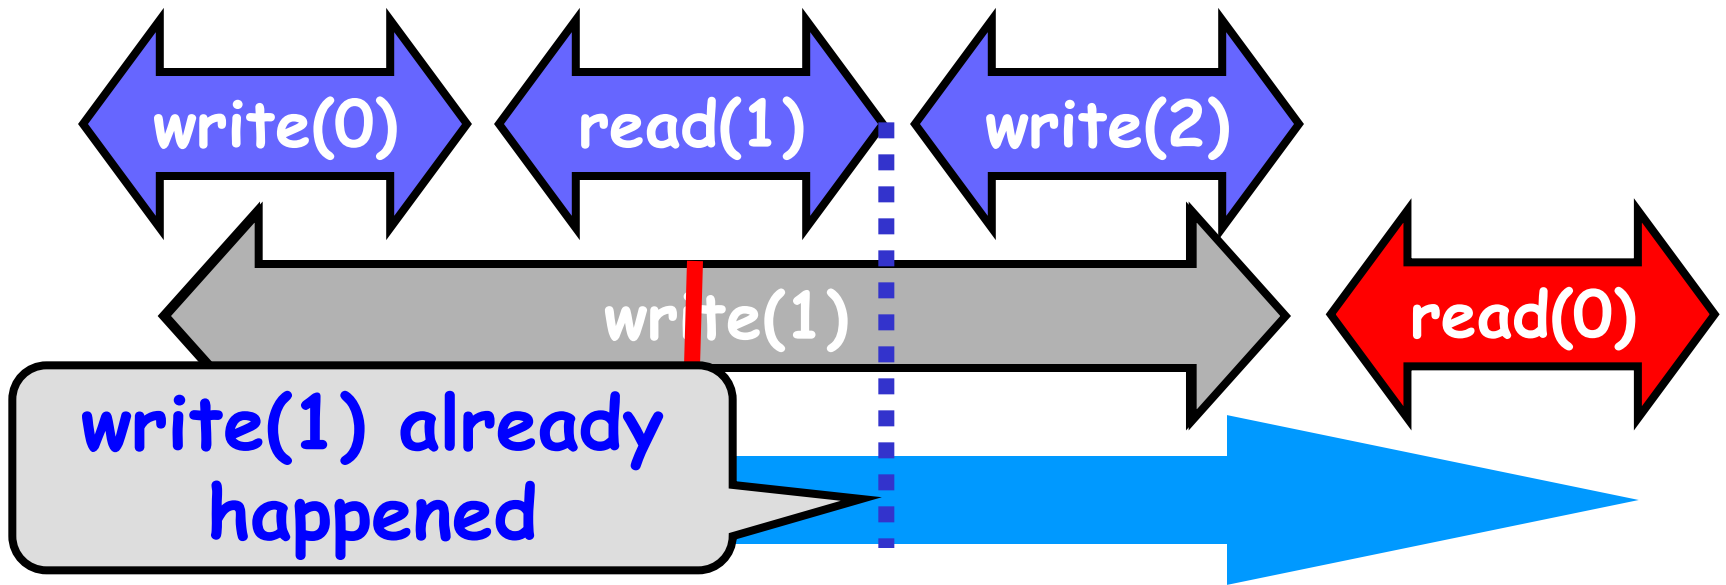
# Read/Write Register Example



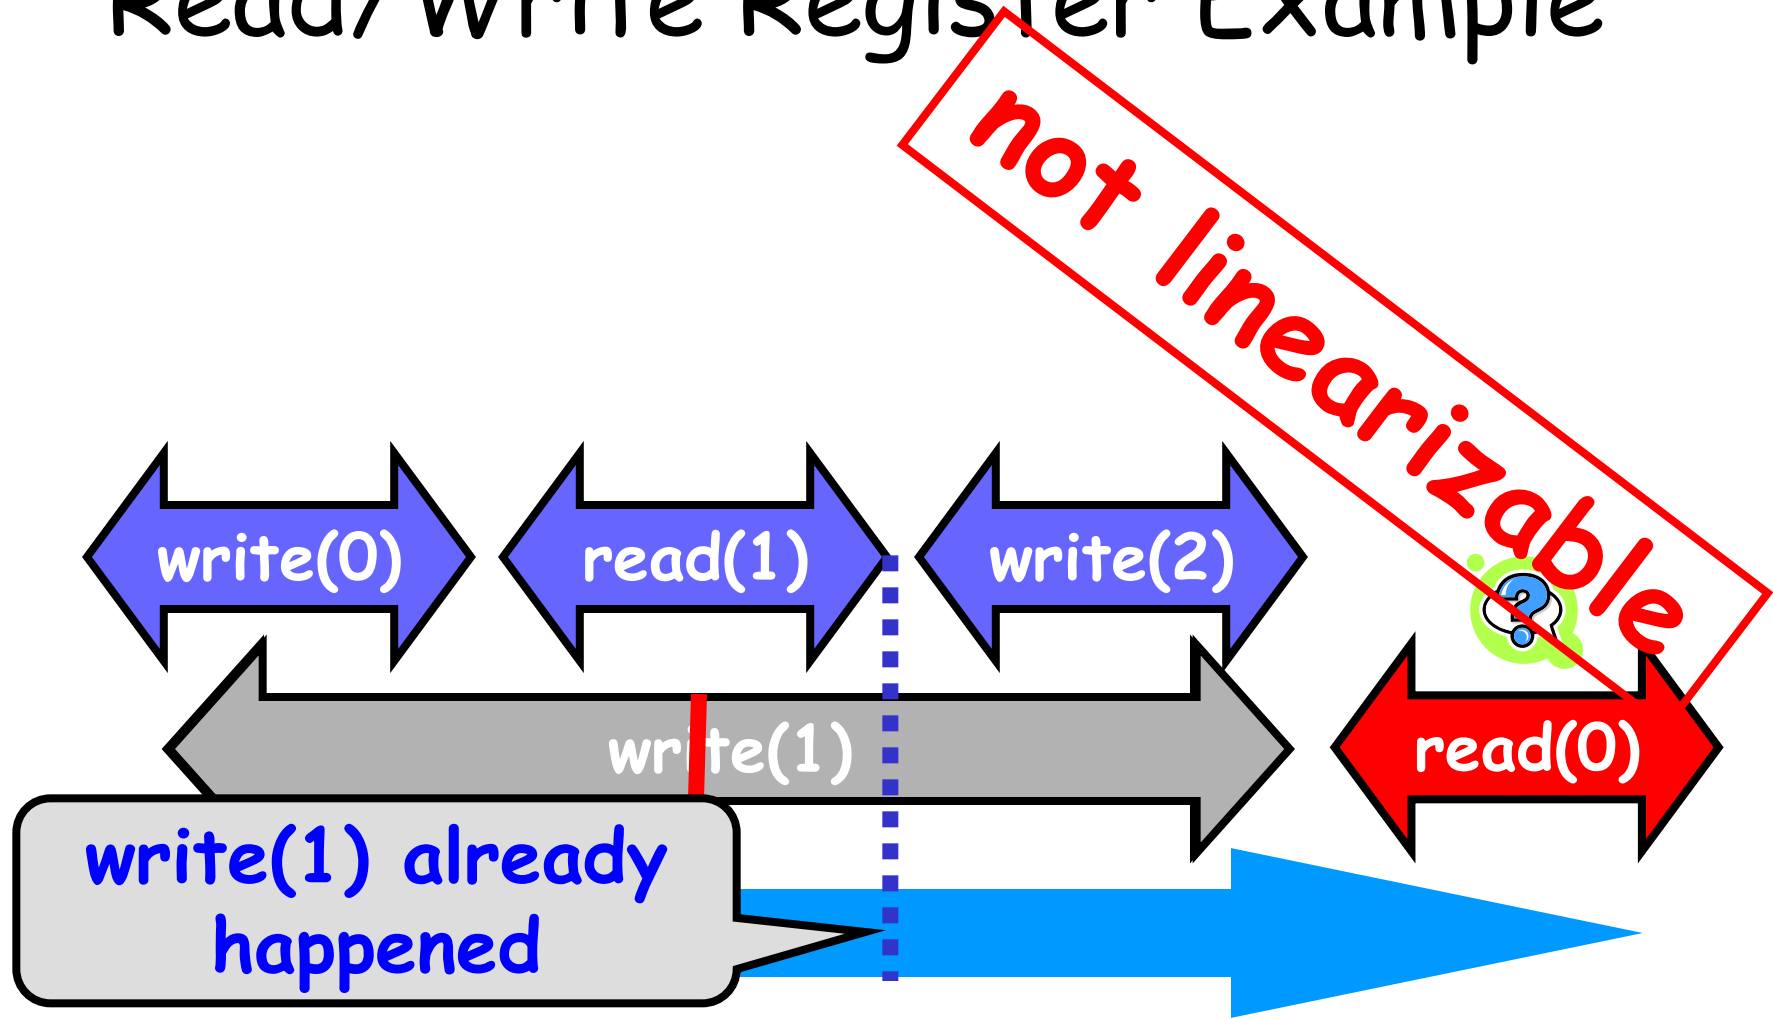
# Read/Write Register Example



# Read/Write Register Example

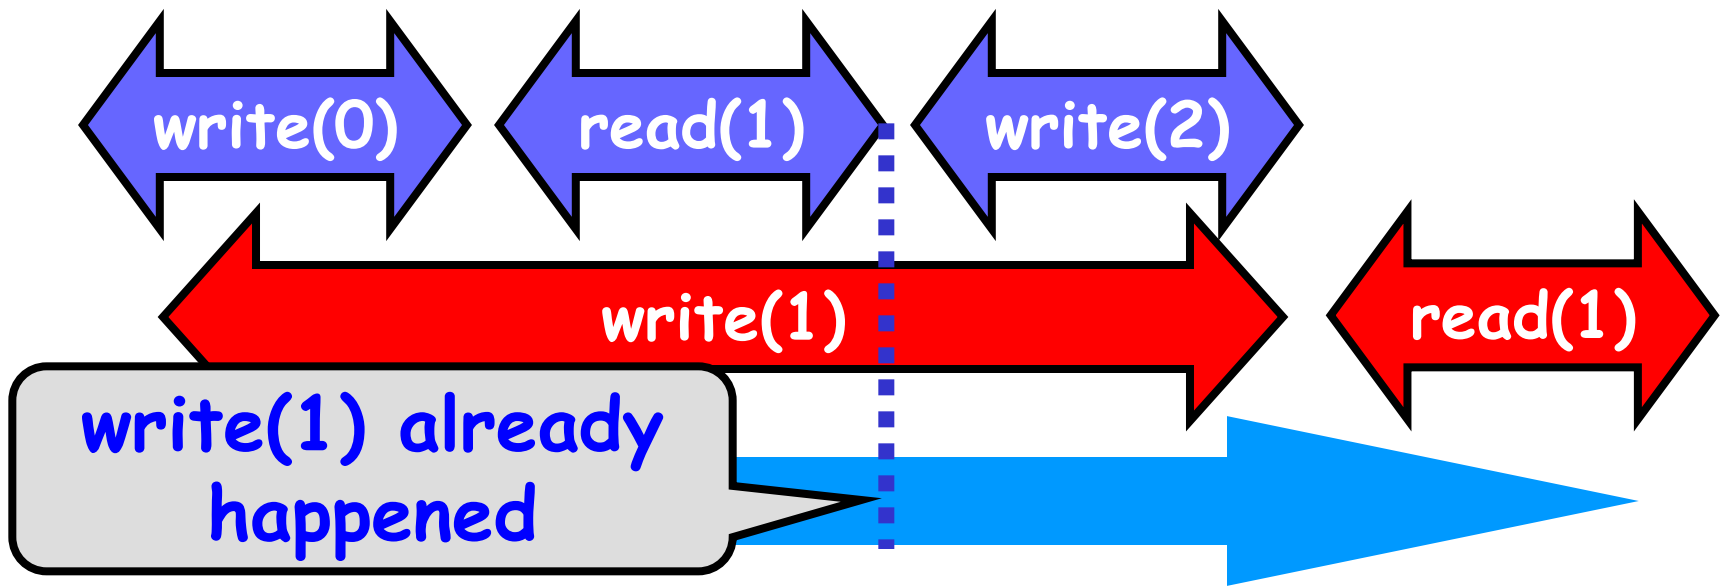


# Read/Write Register Example

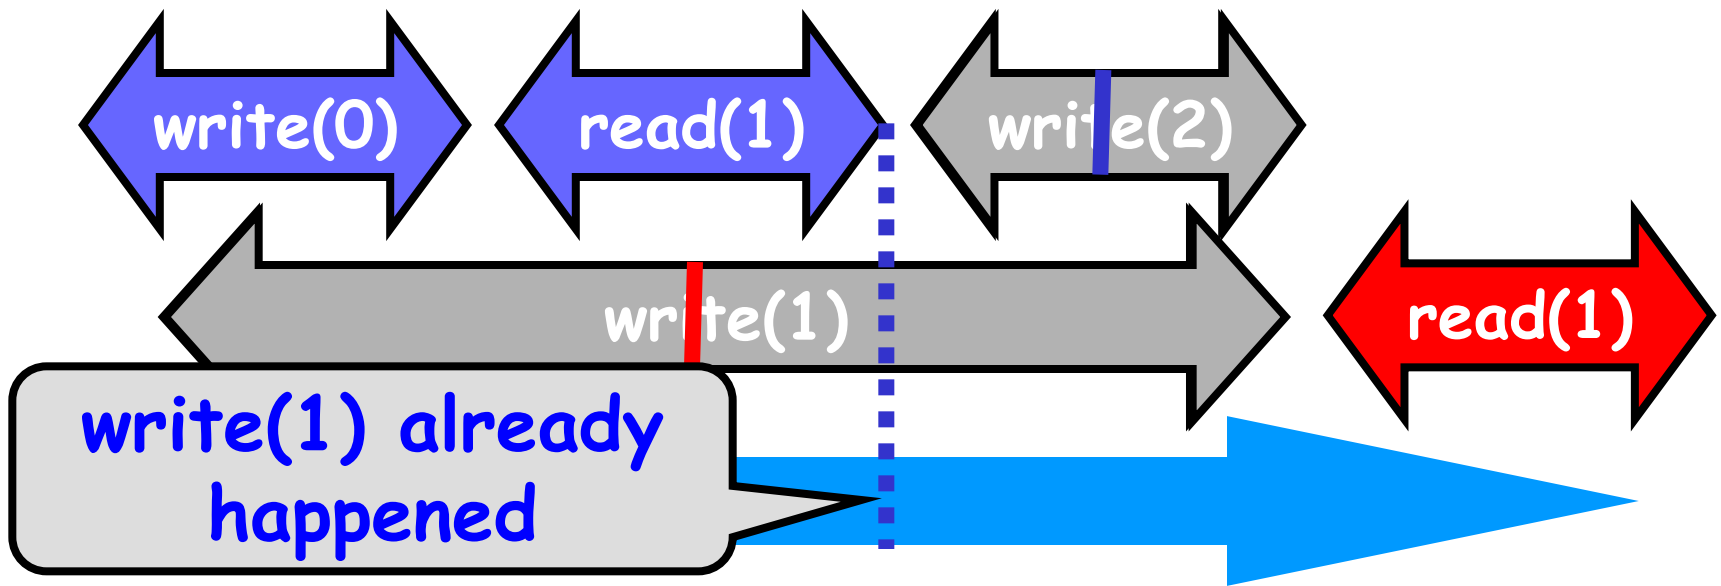




# Read/Write Register Example

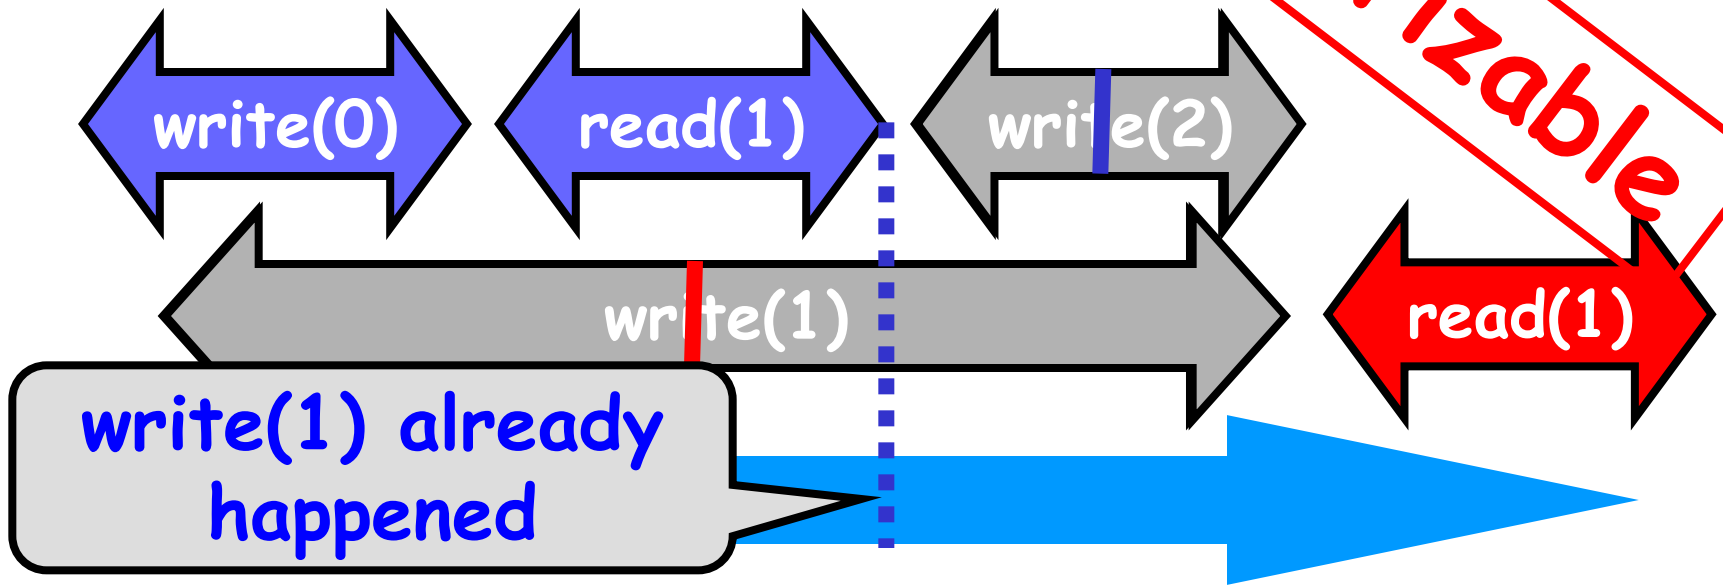


# Read/Write Register Example

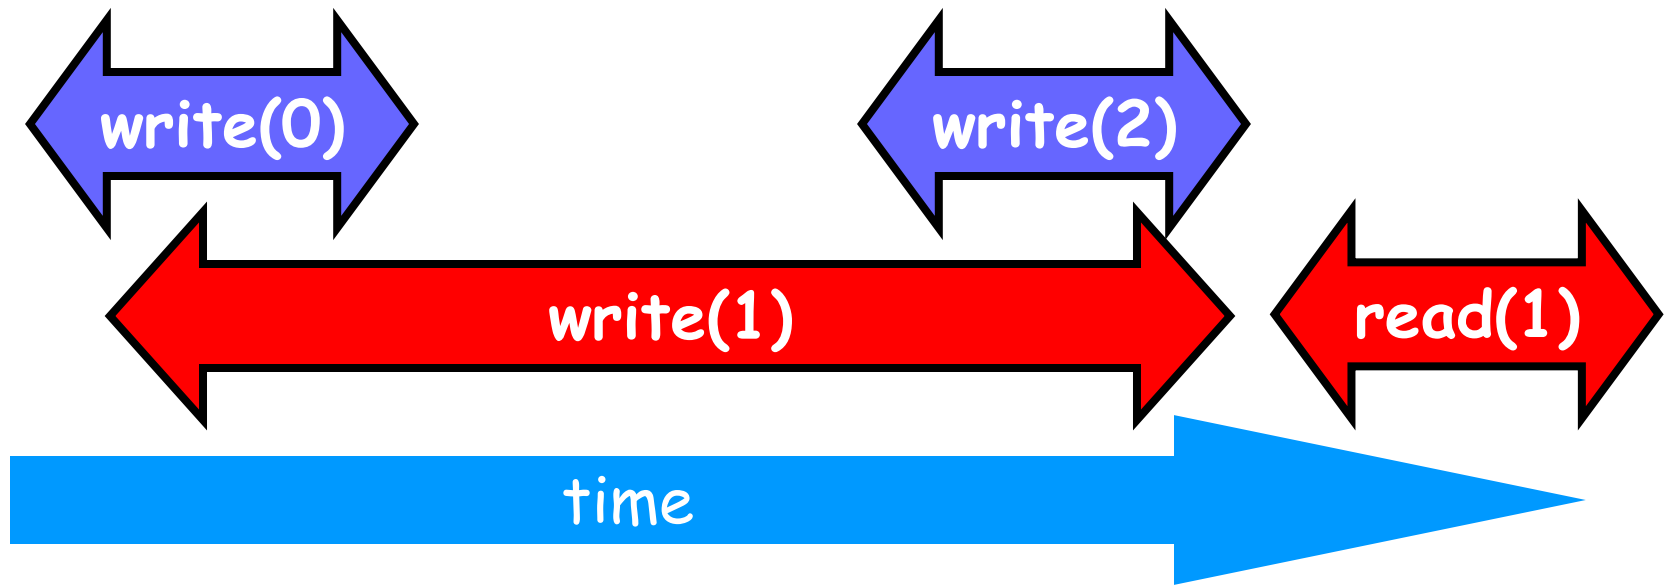


# Read/Write Register Example

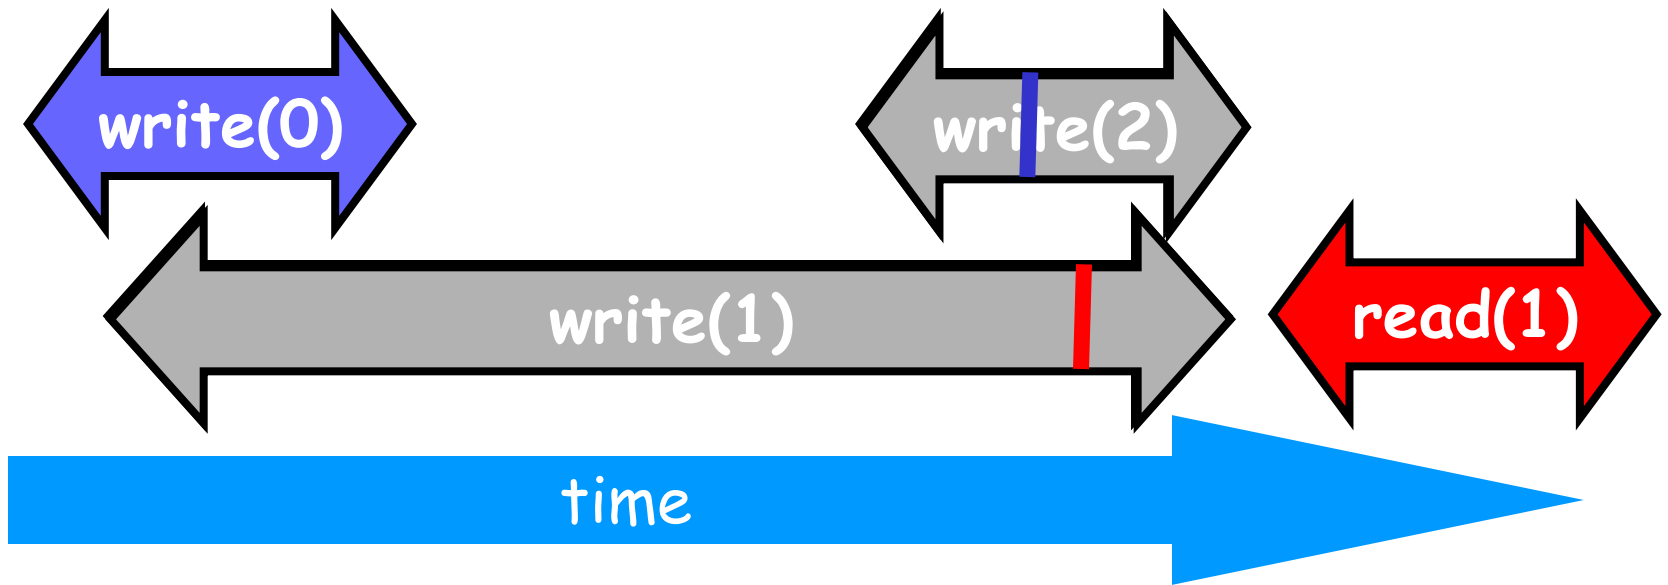
**not linearizable**



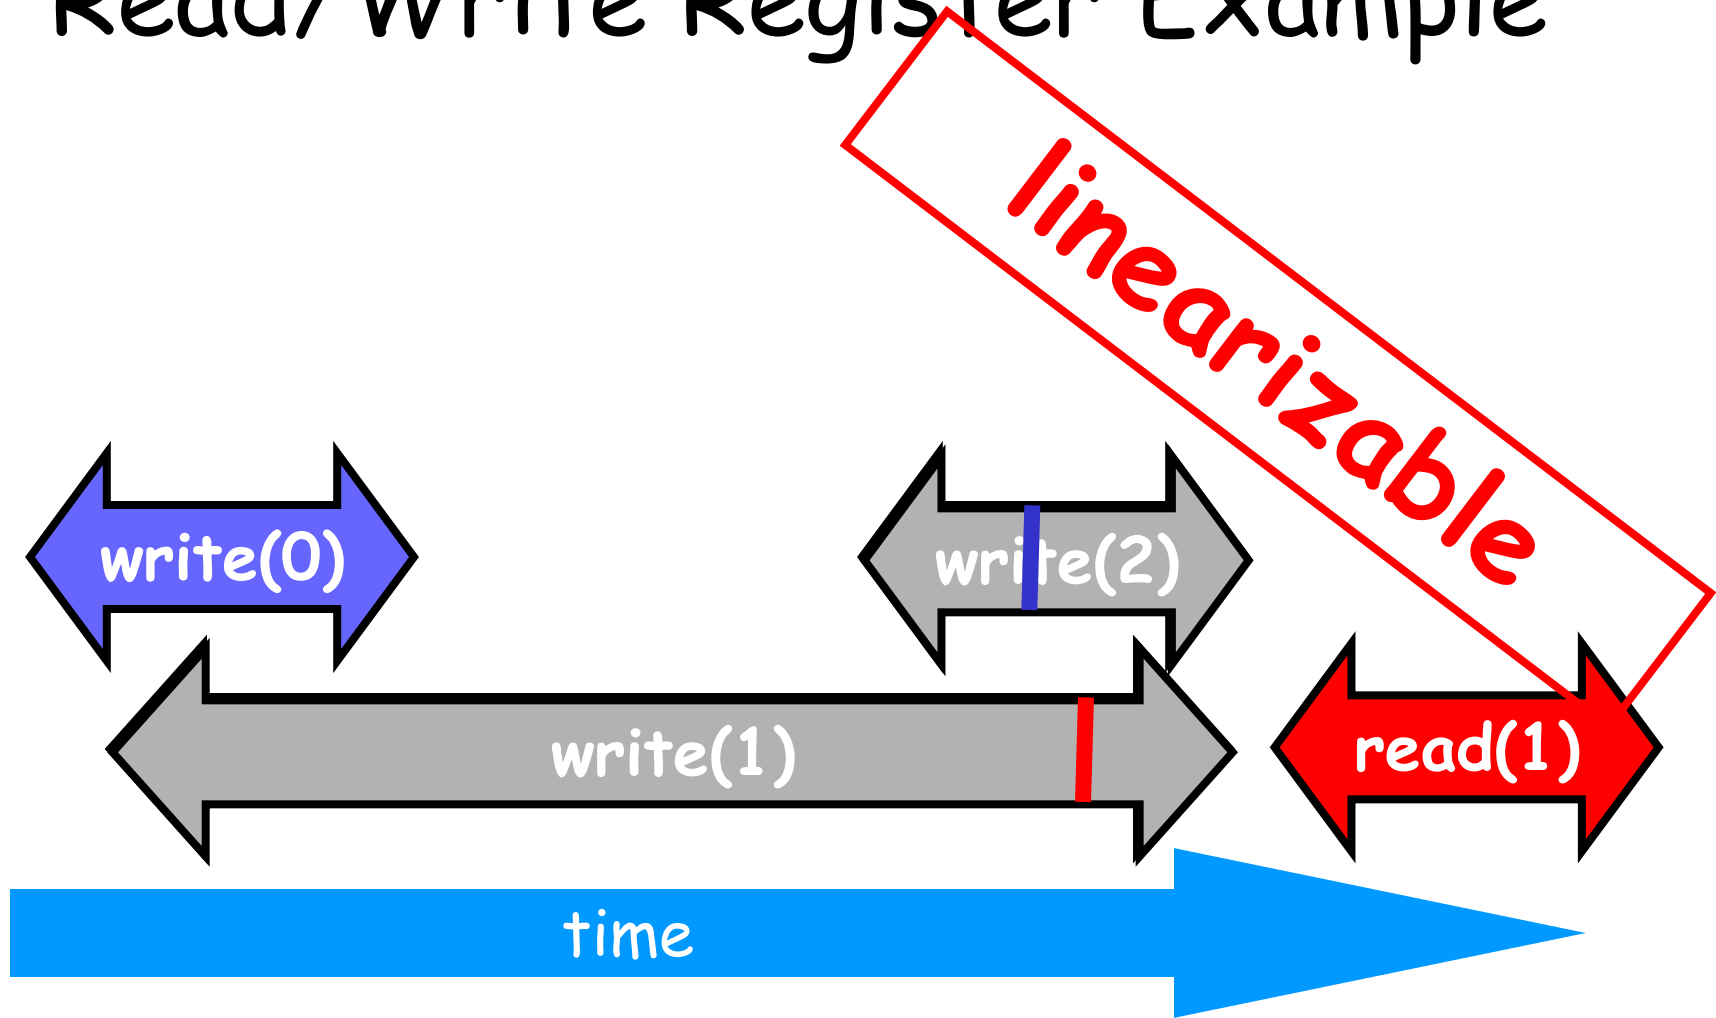
# Read/Write Register Example



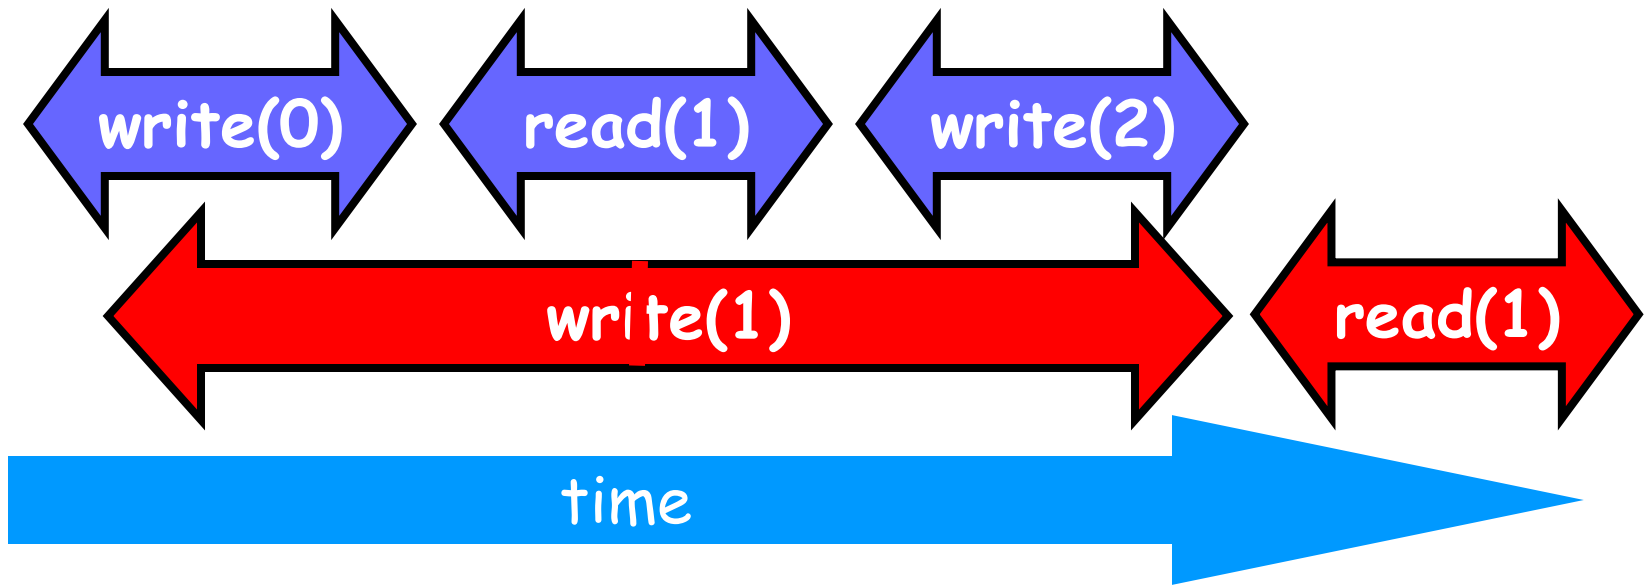
# Read/Write Register Example



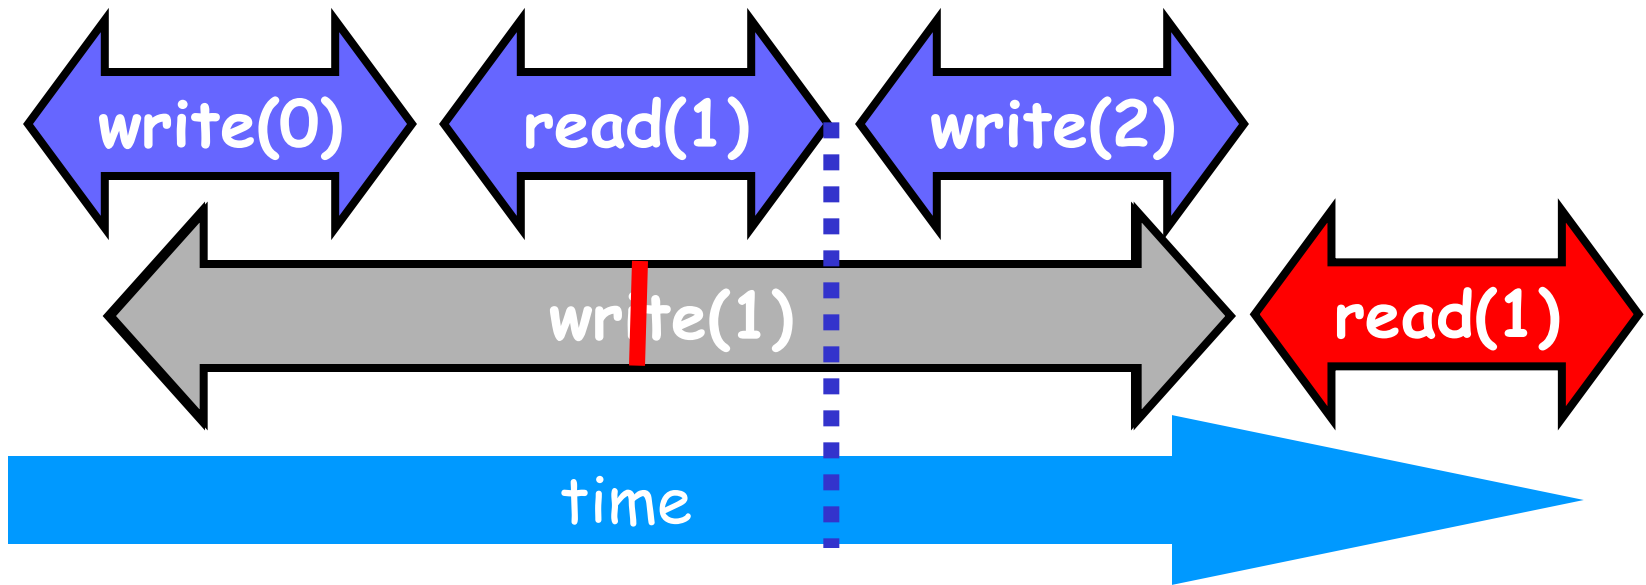
# Read/Write Register Example



# Read/Write Register Example

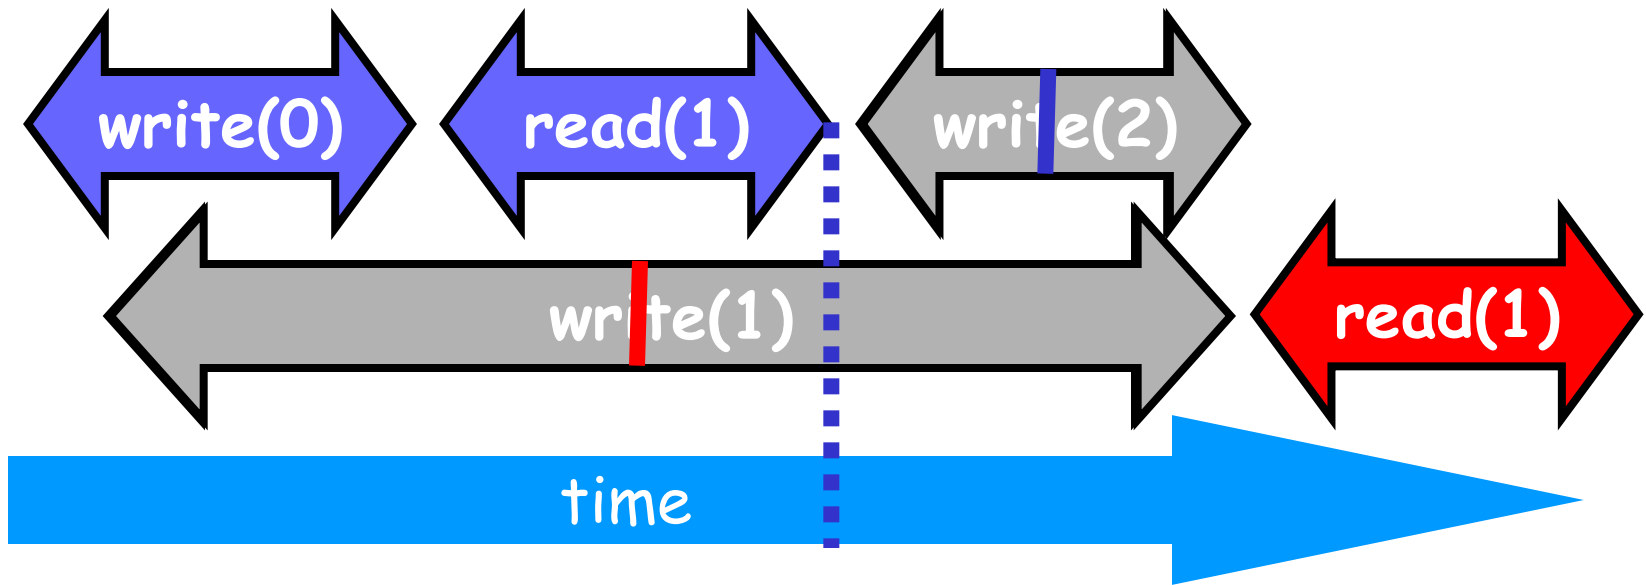


# Read/Write Register Example

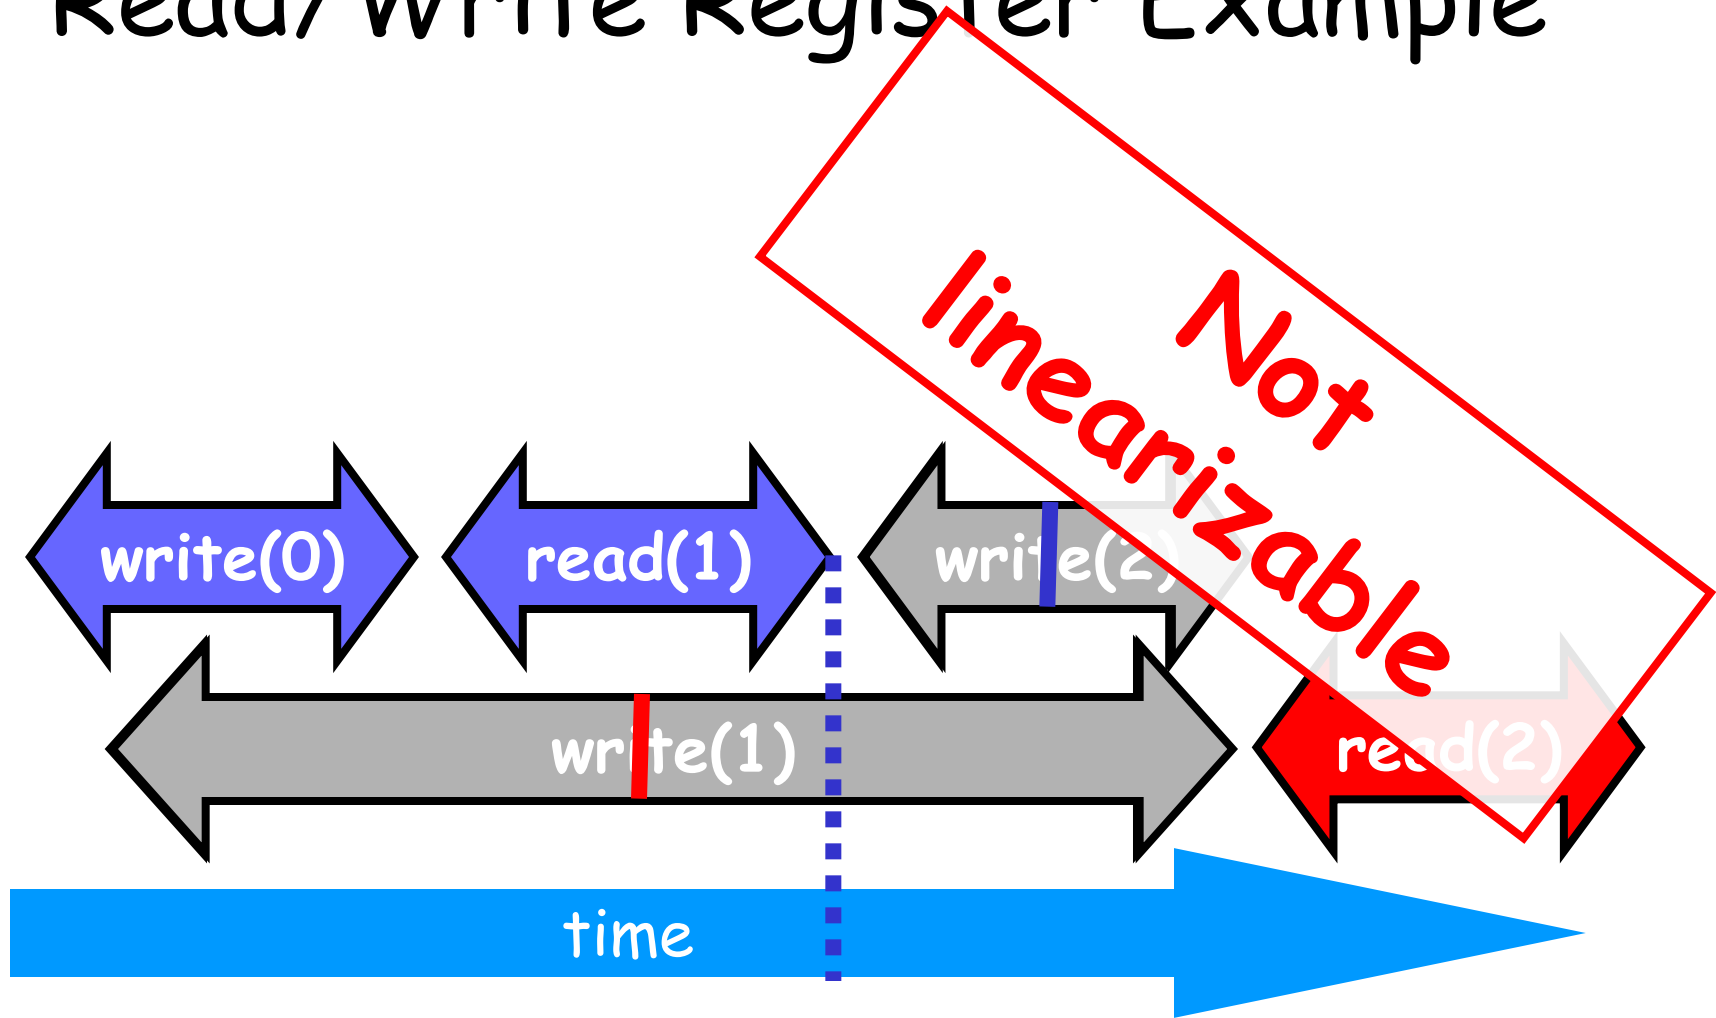




# Read/Write Register Example



# Read/Write Register Example



# Talking About Executions

- Why?
  - Can't we specify the linearization point of each operation without describing an execution?
- Not Always
  - In some cases, linearization point depends on the execution

# Formal Model of Executions

- Define precisely what we mean
  - Ambiguity is bad when intuition is weak
- Allow reasoning
  - Formal
  - But mostly informal
    - In the long run, actually more important
    - Ask me why!

# Split Method Calls into Two Events

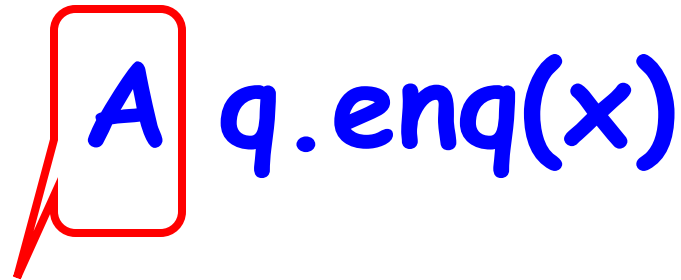
- Invocation
  - method name & args
  - `q.enq(x)`
- Response
  - result or exception
  - `q.enq(x)` returns void
  - `q.deq()` returns x
  - `q.deq()` throws empty

# Invocation Notation

$A \ q.enq(x)$

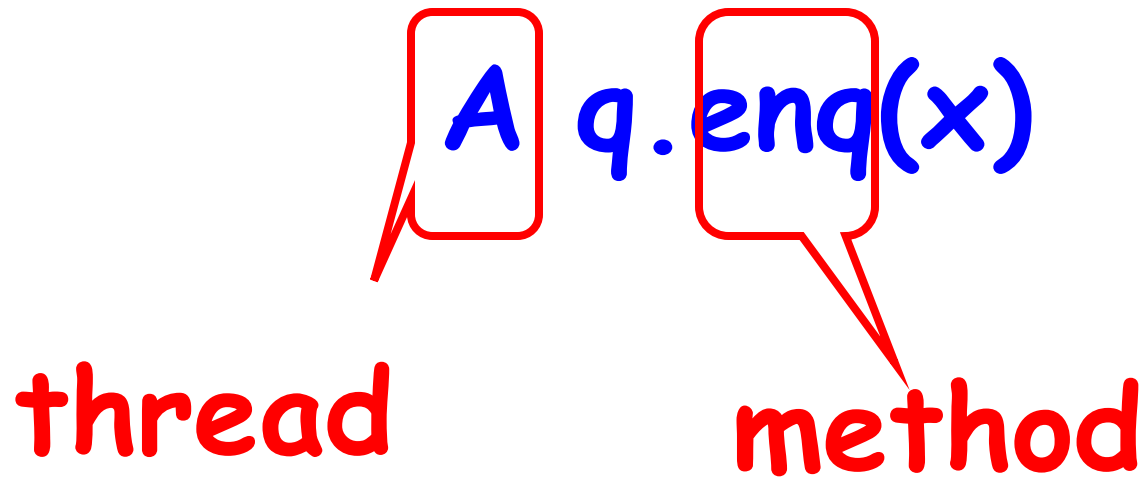


# Invocation Notation

 `A q.enq(x)`


**thread**

# Invocation Notation

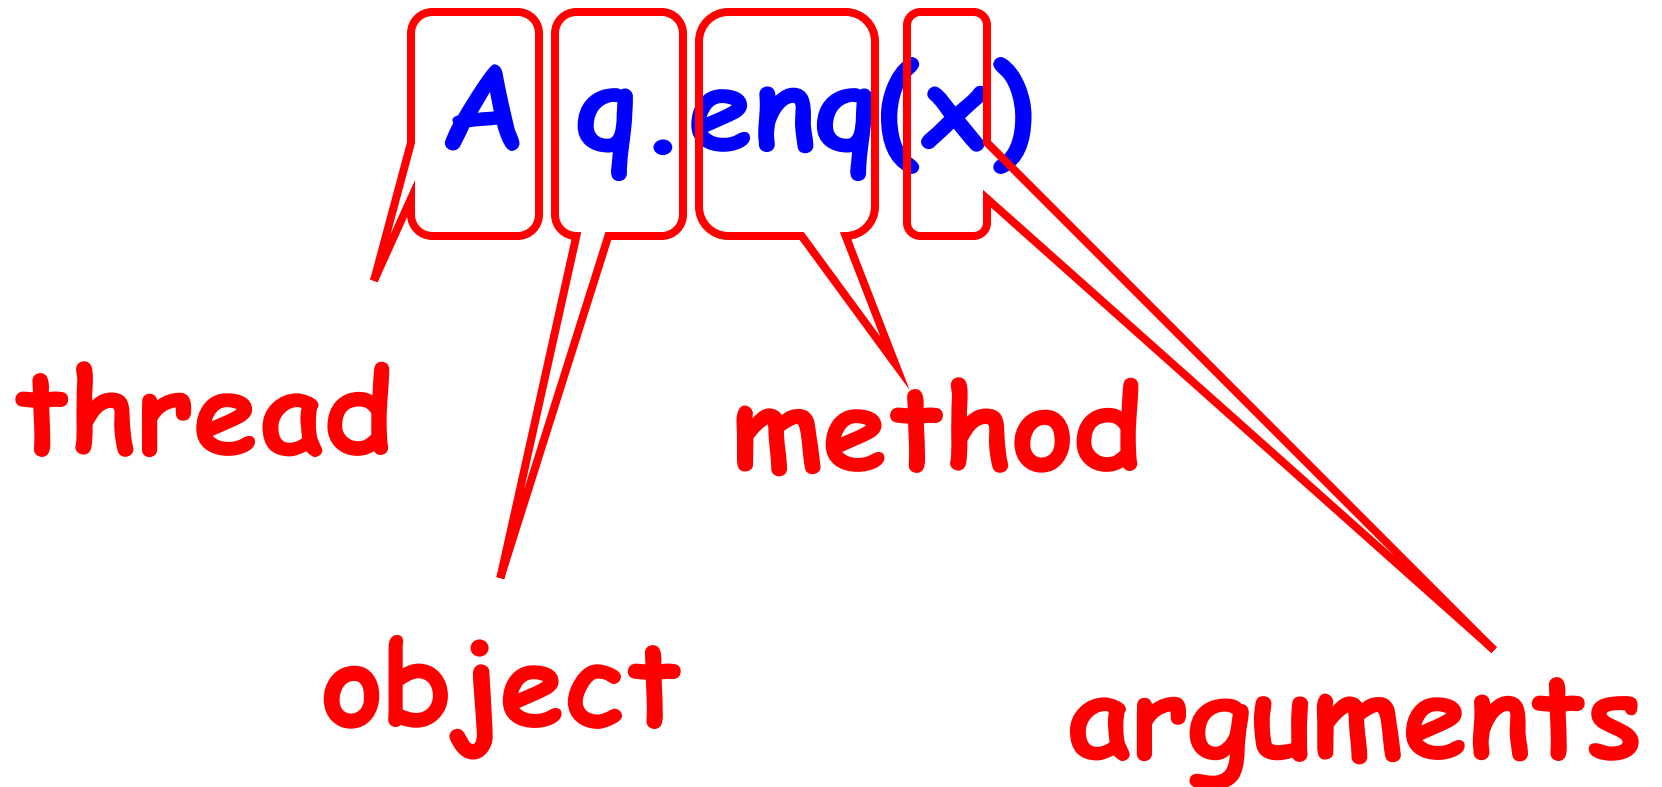




# Invocation Notation



# Invocation Notation



# Response Notation

**A q: void**

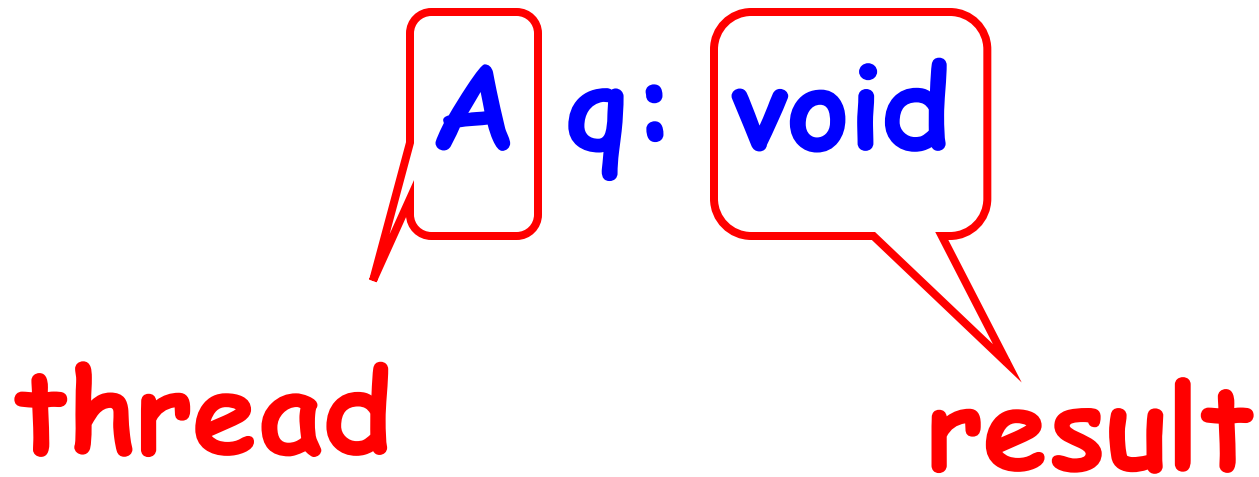


# Response Notation


**A** q: void

**thread**

# Response Notation



# Response Notation



# Response Notation

Method is implicit

A q: void

thread

result

object



# Response Notation

Method is implicit

A q: empty()

thread

exception

object





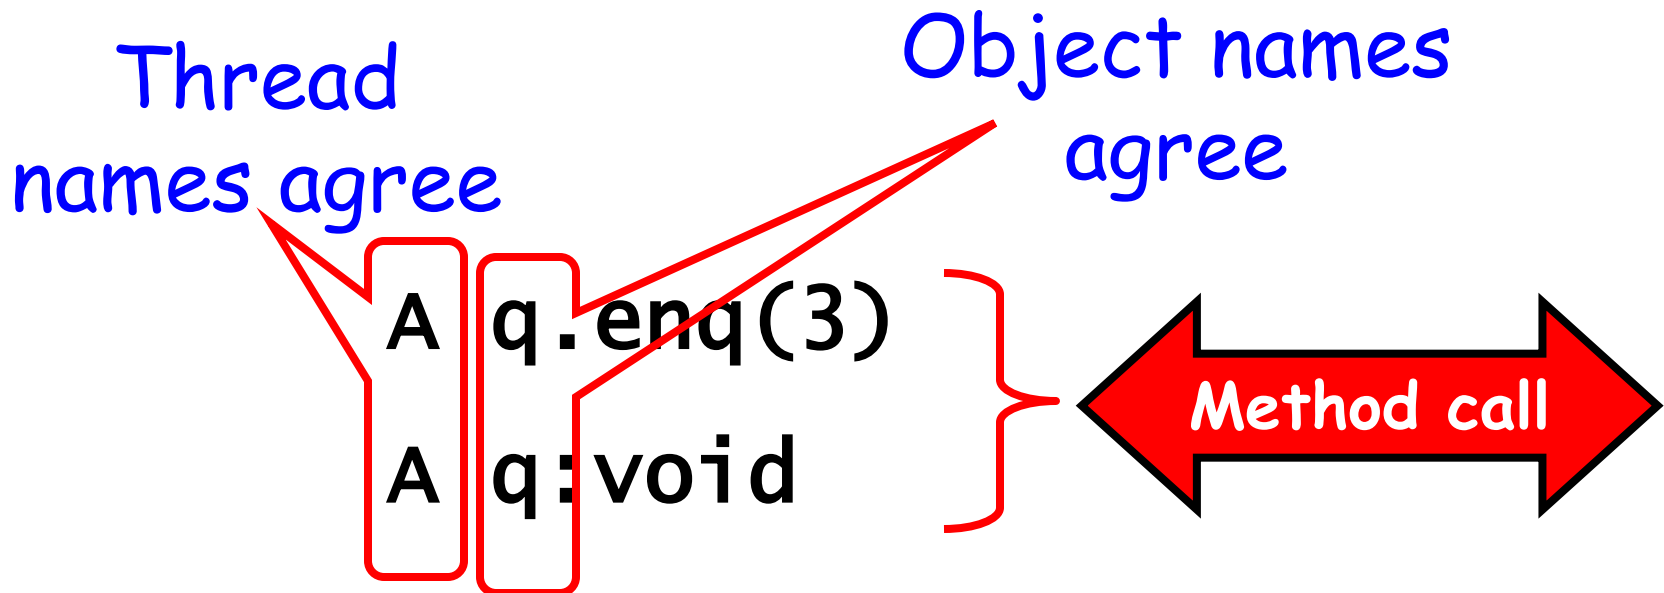
# History - Describing an Execution

$H =$  {  
A q.enq(3)  
A q:void  
A q.enq(5)  
B p.enq(4)  
B p:void  
B q.deq()  
B q:3

**Sequence of  
invocations and  
responses**

# Definition

- Invocation & response *match* if



# Object Projections

$H =$

- A `q.enq(3)`
- A `q:void`
- B `p.enq(4)`
- B `p:void`
- B `q.deq()`
- B `q:3`

# Object Projections

A q.enq(3)

A q:void

H|q =

B q.deq()

B q:3

# Thread Projections

$H =$

- A q.enq(3)
- A q:void
- B p.enq(4)
- B p:void
- B q.deq()
- B q:3

# Thread Projections

$H|B =$

- $B \quad p.enq(4)$
- $B \quad p:void$
- $B \quad q.deq()$
- $B \quad q:3$

# Complete Subhistory

$H =$

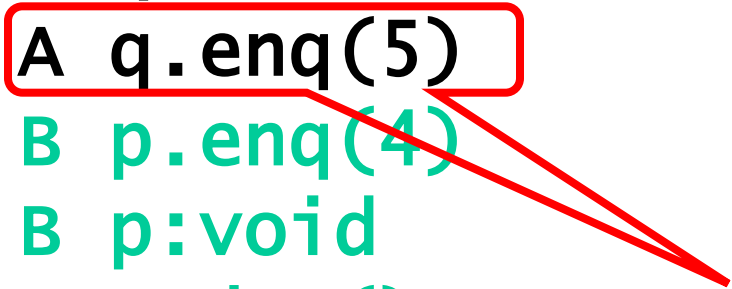
```
A q.enq(3)
A q:void
A q.enq(5)
B p.enq(4)
B p:void
B q.deq()
B q:3
```

**An invocation is pending if it has no matching response**

# Complete Subhistory

A q.enq(3)  
A q:void  
**A q.enq(5)**  
H = B p.enq(4)  
B p:void  
B q.deq()  
B q:3

**May or may not  
have taken effect**



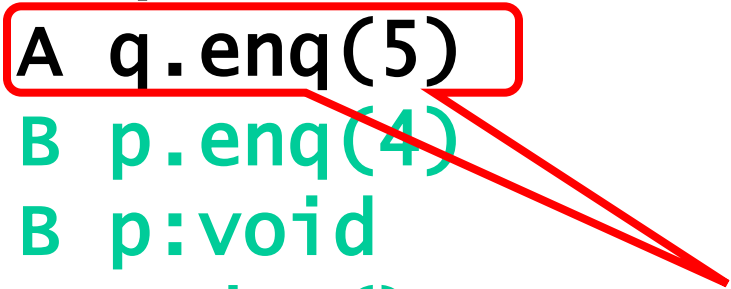


# Complete Subhistory

**H =**

- A q.enq(3)
- A q:void
- A q.enq(5)**
- B p.enq(4)
- B p:void
- B q.deq()
- B q:3

**discard pending invocations**



# Complete Subhistory

A q.enq(3)  
A q:void

**Complete(H) =** B p.enq(4)  
B p:void  
B q.deq()  
B q:3

# Sequential Histories

A q.enq(3)

A q:void

B p.enq(4)

B p:void

B q.deq()

B q:3

A q:enq(5)



# Sequential Histories

A q.enq(3)

A q:void

B p.enq(4)

B p:void

B q.deq()

B q:3

A q:enq(5)

match



# Sequential Histories

A q.enq(3)

A q:void

match

B p.enq(4)

B p:void

match

B q.deq()

B q:3

A q:enq(5)



# Sequential Histories

A q.enq(3)

A q:void

match

B p.enq(4)

B p:void

match

B q.deq()

B q:3

match

A q:enq(5)



# Sequential Histories

A q.enq(3)

A q:void

match

B p.enq(4)

B p:void

match

B q.deq()

B q:3

match

A q:enq(5)

Final pending  
invocation OK



# Sequential Histories

A q.enq(3)  
A q:void

B p.enq(4)  
B p:void

B q.deq()  
B q:3

A q:enq(5)

Method calls of  
do different threads  
do not interleave

match  
match  
match  
Final pending  
invocation OK





# Well-Formed Histories

**H=**

- A q.enq(3)
- B p.enq(4)
- B p:void
- B q.deq()
- A q:void
- B q:3

# Well-Formed Histories

Per-thread  
projections sequential

H=

A q.enq(3)  
B p.enq(4)  
B p:void  
B q.deq()  
A q:void  
B q:3

H | B=

B p.enq(4)  
B p:void  
B q.deq()  
B q:3

# Well-Formed Histories

Per-thread  
projections sequential

$H =$   
A q.enq(3)  
B p.enq(4)  
B p:void  
B q.deq()  
A q:void  
B q:3

$H | B =$   
B p.enq(4)  
B p:void  
B q.deq()  
B q:3

$H | A =$   
A q.enq(3)  
A q:void

# Equivalent Histories

Threads see the same thing in both

$$\left\{ \begin{array}{l} H|A = G|A \\ H|B = G|B \end{array} \right.$$

H=

```
A q.enq(3)
B p.enq(4)
B p:void
B q.deq()
A q:void
B q:3
```

G=

```
A q.enq(3)
A q:void
B p.enq(4)
B p:void
B q.deq()
B q:3
```

# Sequential Specifications


- A sequential specification is some way of telling whether a
  - Single-thread, single-object history
  - Is legal
- For example:
  - Pre and post-conditions
  - But plenty of other techniques exist ...

# Legal Histories


- A sequential (multi-object) history  $H$  is legal if
  - For every object  $x$
  - $H|x$  is in the sequential spec for  $x$

# Precedence

A q.enq(3)  
B p.enq(4)  
B p.void  
A q:void  
B q.deq()  
B q:3




A method call **precedes**  
another if response  
event precedes  
invocation event




# Non-Precedence

A q.enq(3)  
B p.enq(4)  
B p.void  
B q.deq()  
A q:void  
B q:3




Some method calls  
**overlap** one another





# Notation

- Given
  - History  $H$
  - method executions  $m_0$  and  $m_1$  in  $H$
- We say  $m_0 \rightarrow_H m_1$ , if
  - $m_0$  precedes  $m_1$
- Relation  $m_0 \rightarrow_H m_1$  is a
  - Partial order
  - Total order if  $H$  is sequential




# Linearizability

- History  $H$  is *linearizable* if it can be extended to  $G$  by
  - Appending zero or more responses to pending invocations
  - Discarding other pending invocations
- So that  $G$  is equivalent to
  - Legal sequential history  $S$
  - where  $\rightarrow_G \subset \rightarrow_S$

# What is $\rightarrow_G \subset \rightarrow_S$

$$\rightarrow_G = \{a \rightarrow c, b \rightarrow c\}$$

$$\rightarrow_S = \{a \rightarrow b, a \rightarrow c, b \rightarrow c\}$$

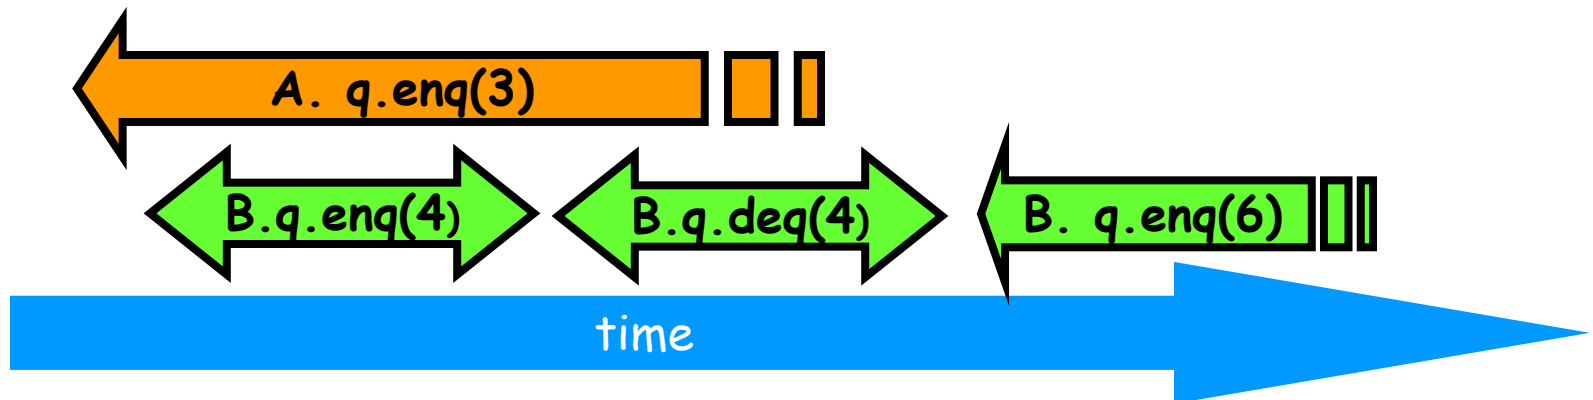


# Remarks

- Some pending invocations
  - Took effect, so keep them
  - Discard the rest
- Condition  $\rightarrow_G \subset \rightarrow_S$ 
  - Means that **S** respects “real-time order” of **G**

# Example

A q.enq(3)  
B q.enq(4)  
B q:void  
B q.deq()  
B q:4  
B q:enq(6)



# Example

A q.enq(3)

B q.enq(4)

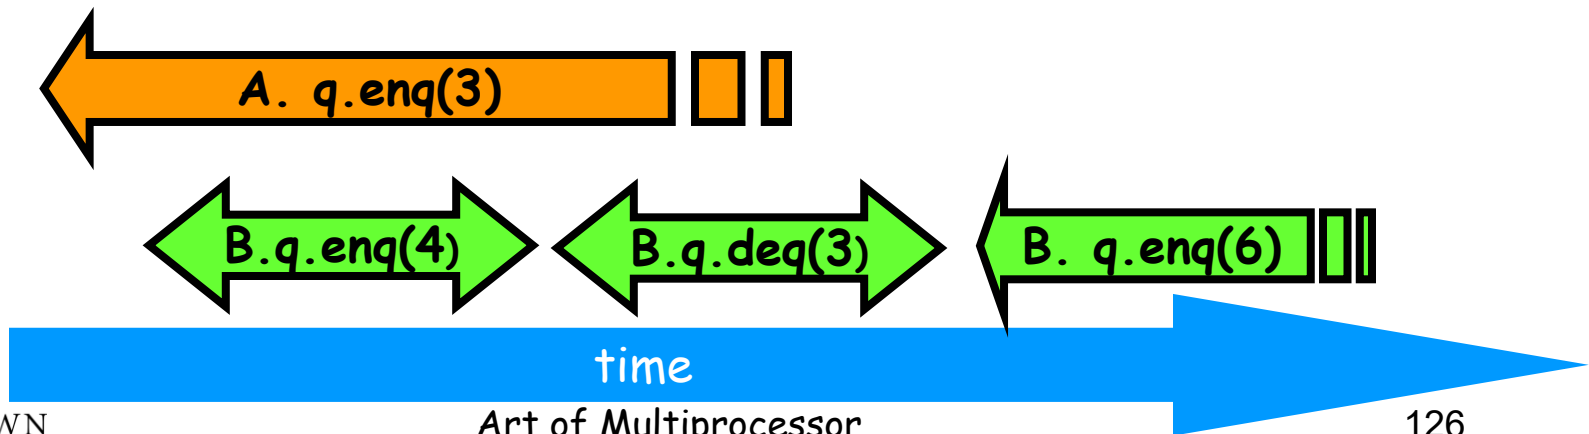
B q:void

B q.deq()

B q:4

B q:enq(6)

Complete this  
pending  
invocation



# Example

A q.enq(3)

B q.enq(4)

B q:void

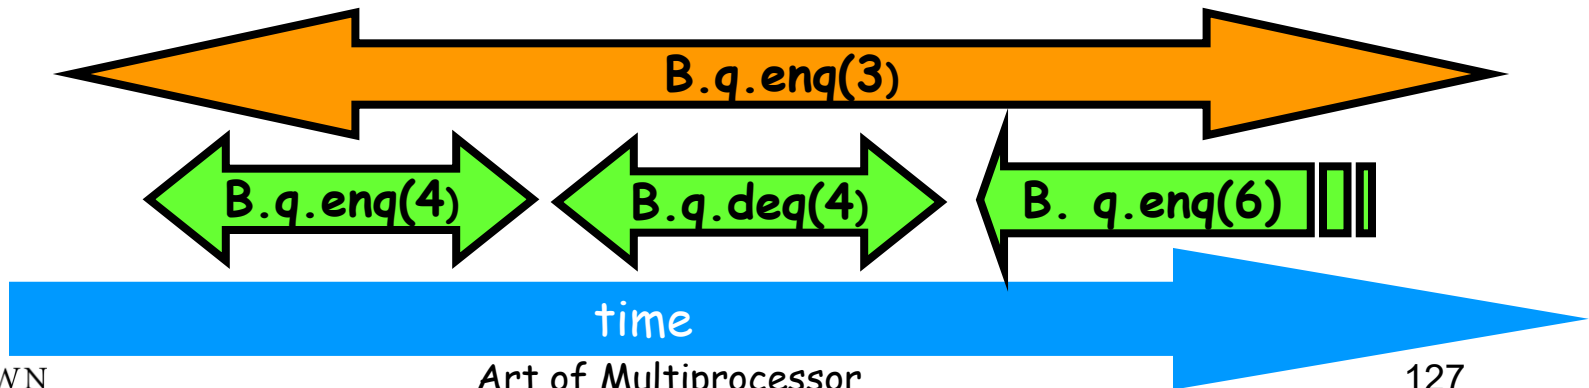
B q.deq()

B q:4

B q:enq(6)

A q:void

Complete this  
pending  
invocation



# Example

discard this one

A q.enq(3)

B q.enq(4)

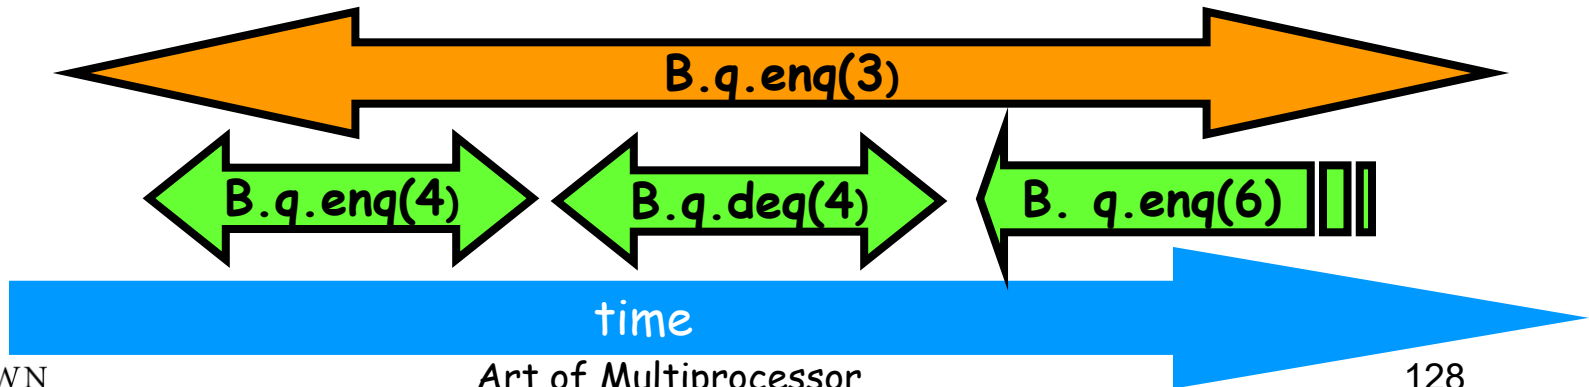
B q:void

B q.deq()

B q:4

B q:enq(6)

A q:void





# Example

A q.enq(3)

B q.enq(4)

B q:void

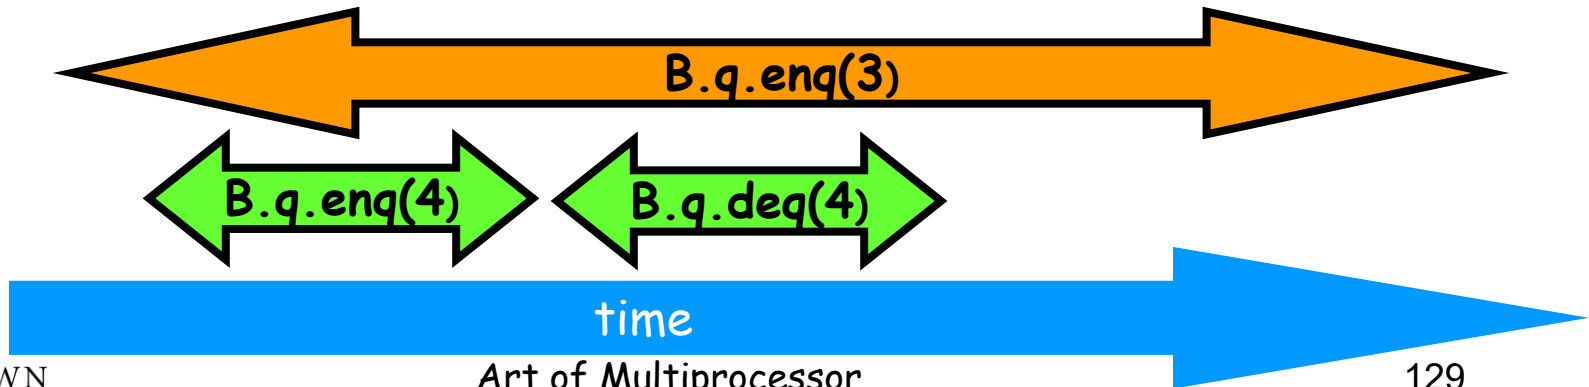
B q.deq()

B q:4

discard this one

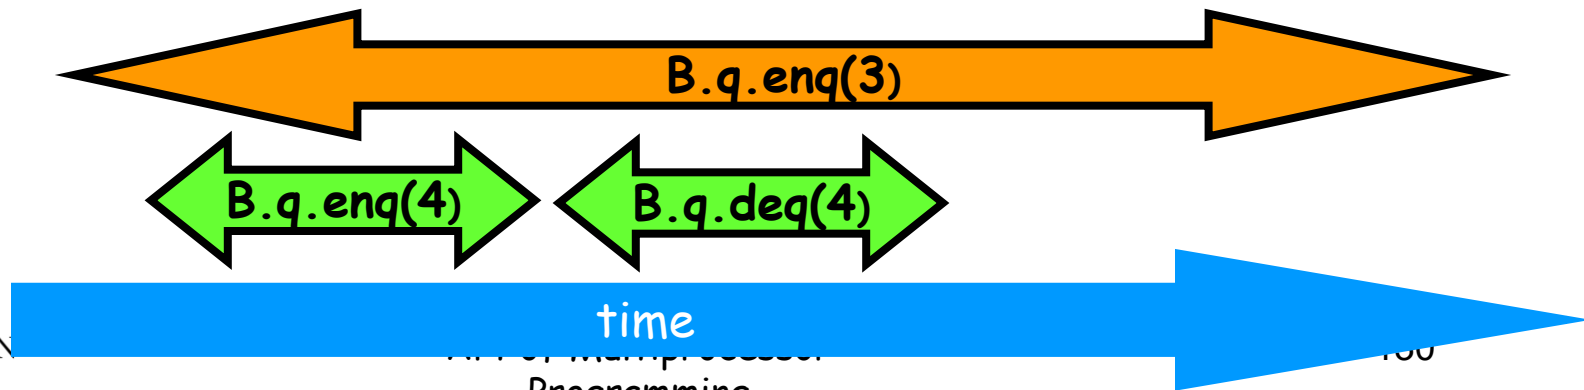


A q:void



# Example

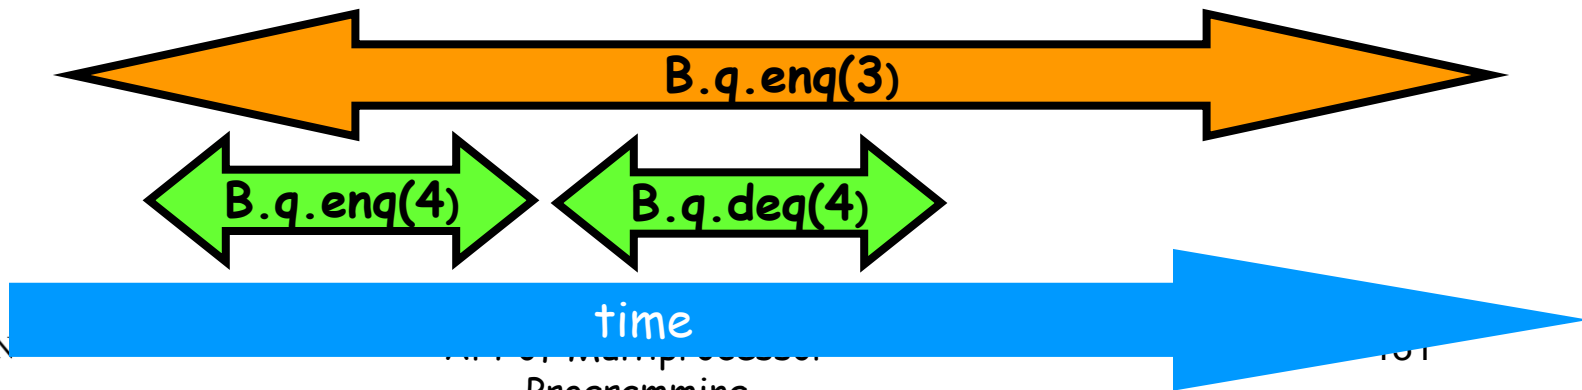
A q.enq(3)  
B q.enq(4)  
B q:void  
B q.deq()  
B q:4  
A q:void



# Example

A q.enq(3)  
B q.enq(4)  
B q:void  
B q.deq()  
B q:4  
A q:void

B q.enq(4)  
B q:void  
A q.enq(3)  
A q:void  
B q.deq()  
B q:4

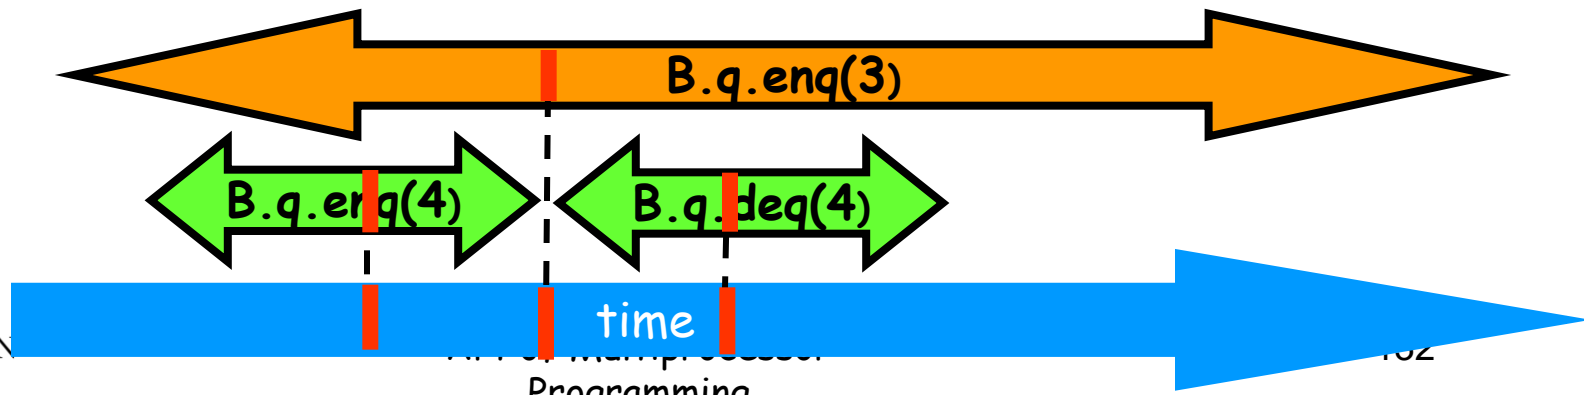


# Example

Equivalent sequential history

A q.enq(3)  
B q.enq(4)  
B q:void  
B q.deq()  
B q:4  
A q:void

B q.enq(4)  
B q:void  
A q.enq(3)  
A q:void  
B q.deq()  
B q:4



# Composability Theorem

- History  $H$  is linearizable if and only if
  - For every object  $x$
  - $H|x$  is linearizable
- We care about objects only!
  - (Materialism?)

# Why Does Composability Matter?

- Modularity
- Can prove linearizability of objects in isolation
- Can compose independently-implemented objects


# Reasoning About Linearity: Locking

```
public class Queue<T> {
```

```
    int head = 0, tail = 0;  
    T[QSIZE] items;
```

```
    public synchronized void enq(T x) {  
        while (tail - head == QSIZE)  
            this.wait();  
        items[tail % QSIZE] = x; tail++;  
        this.notifyAll();  
    }
```

```
    ...  
}}
```



# Implementation: Enq

```
public class Queue<T> {  
    int head = 0, tail = 0;  
    T[QSIZE] items;  
  
    public synchronized void enq(Tt x) {  
        while (tail - head == QSIZE)  
            this.wait();  
        items[tail % QSIZE] = x; tail++;  
        this.notifyAll();  
    }  
    ...  
}
```

Linearization order is order lock released





# More Reasoning: Lock-free


```
public class LockFreeQueue {
```

```
    int head = 0, tail = 0;  
    Item[QSIZE] items;
```

```
    public void enq(Item x) {  
        while (tail-head == QSIZE); // busy-wait  
        items[tail % QSIZE] = x; tail++;  
    }
```

```
    public Item deq() {  
        while (tail == head); // busy-wait  
        Item item = items[head % QSIZE]; head++;  
        return item;  
    }
```

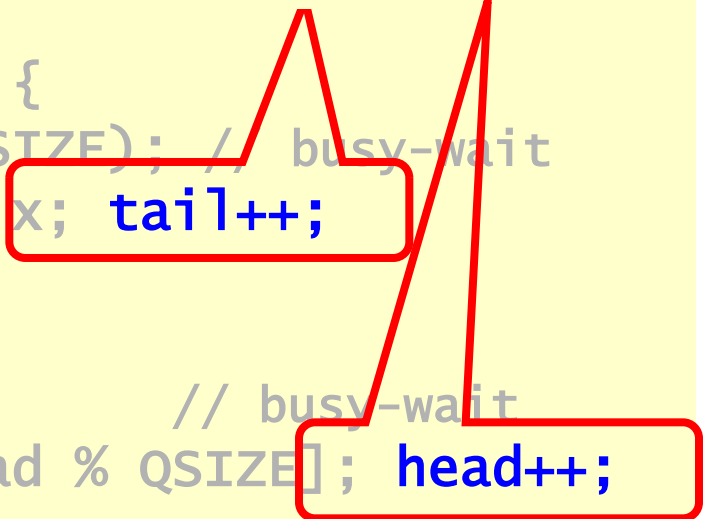
```
}}
```



# More Reasoning

```
public class LockFreeQ {  
    int head = 0, tail = 0;  
    Item[QSIZE] items;  
  
    public void enq(Item x) {  
        while (tail-head == QSIZE); // busy-wait  
        items[tail % QSIZE] = x; tail++;  
    }  
  
    public Item deq() {  
        while (tail == head); // busy-wait  
        Item item = items[head % QSIZE]; head++;  
        return item;  
    }  
}
```

Linearization order is order head and tail fields modified



# Strategy

- Identify one atomic step where method “happens”
  - Critical section
  - Machine instruction
- Doesn't always work
  - Might need to define several different steps for a given method

# Linearizability: Summary

- Powerful specification tool for shared objects
- Allows us to capture the notion of objects being “atomic”
- Don't leave home without it

# Alternative: Sequential Consistency

- History  $H$  is **Sequentially Consistent** if it can be extended to  $G$  by
  - Appending zero or more responses to pending invocations
  - Discarding other pending invocations
- So that  $G$  is equivalent to a
  - Legal sequential history  $S$


Differs from linearizability

~~Where  $\rightarrow G \subseteq \rightarrow S$~~



# Alternative: Sequential Consistency

- No **need to preserve** real-time order
  - Cannot re-order operations done by the same thread
  - Can re-order non-overlapping operations done by different threads
- Often used to describe multiprocessor memory architectures

# Example






# Example










# Example









# Example










# Example









# Example

**not linearizable**









If we can ignore  
real-time order





Sequentially Consistent




# Theorem

Sequential Consistency is not a  
local property


(and thus we lose compasability...)

# FIFO Queue Example






# FIFO Queue Example




# FIFO Queue Example




History H



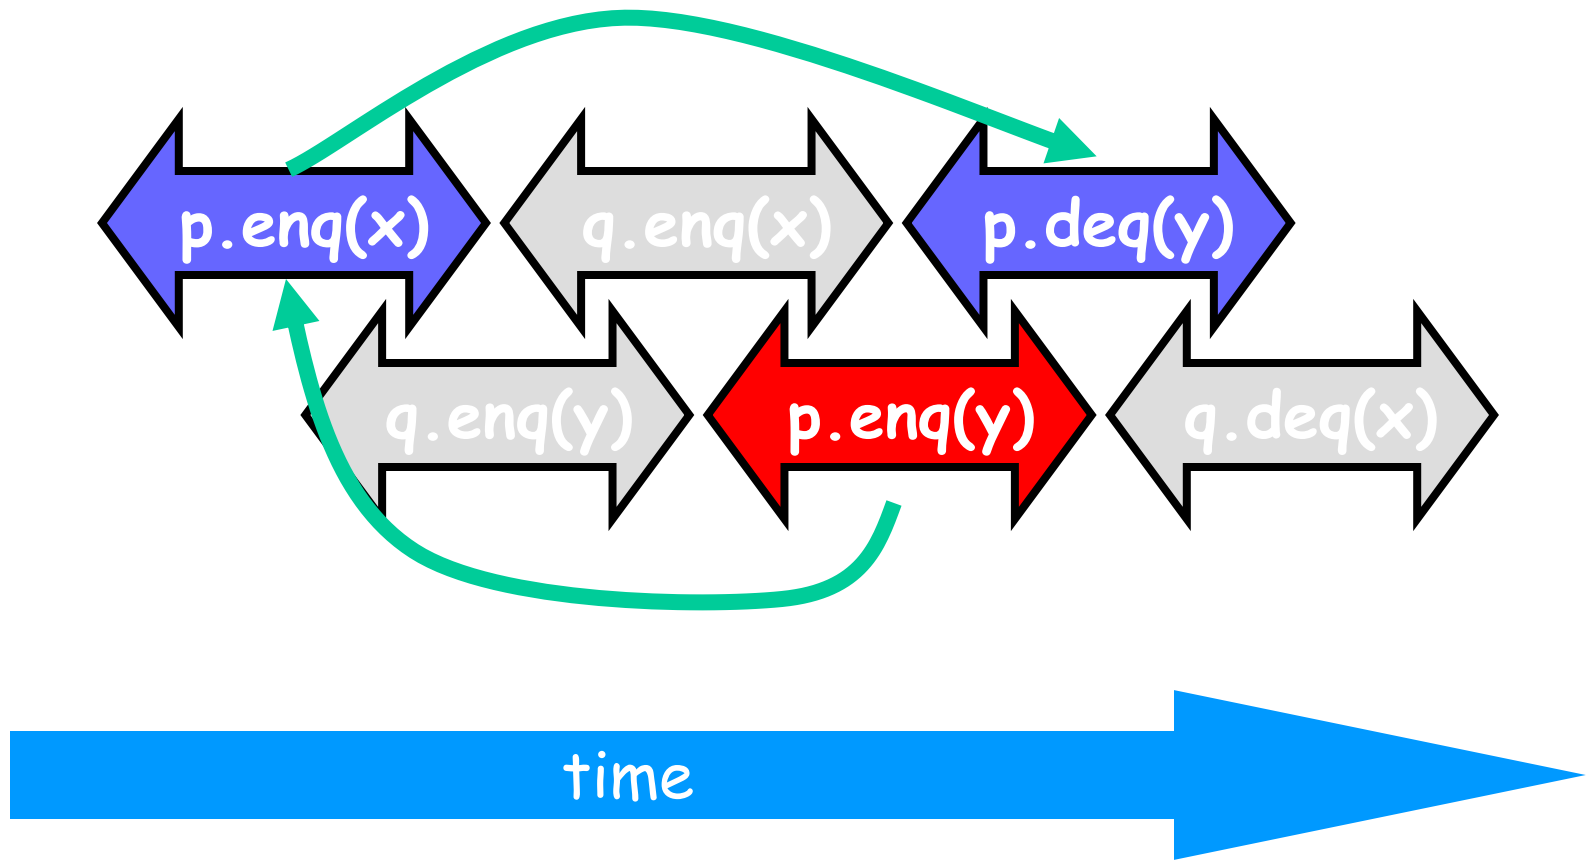
# Help Sequentially Consistent




# H|q Sequentially Consistent



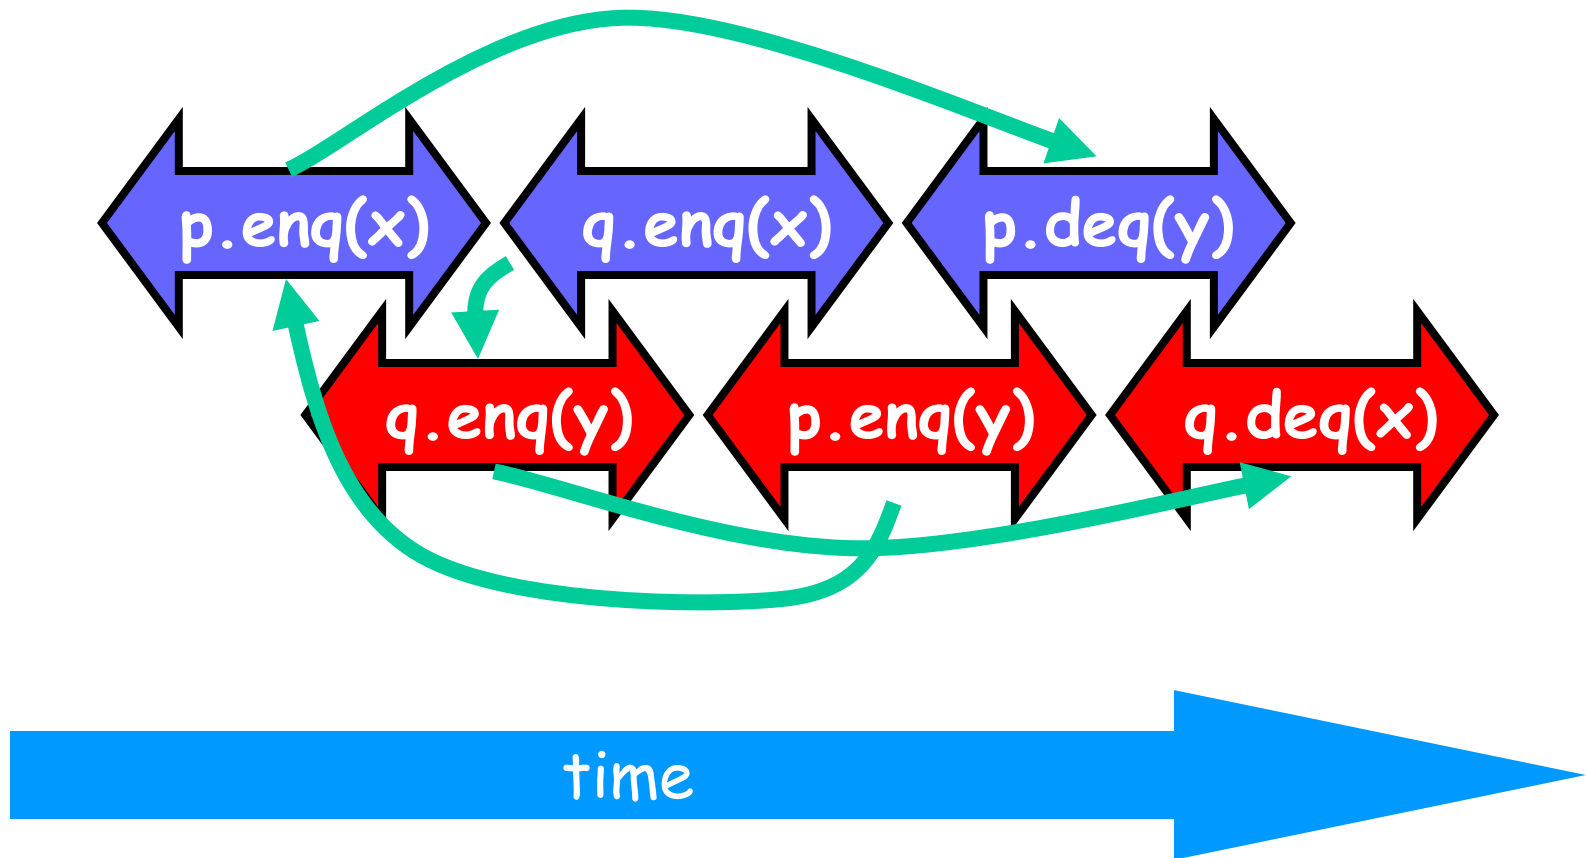
# Ordering imposed by p



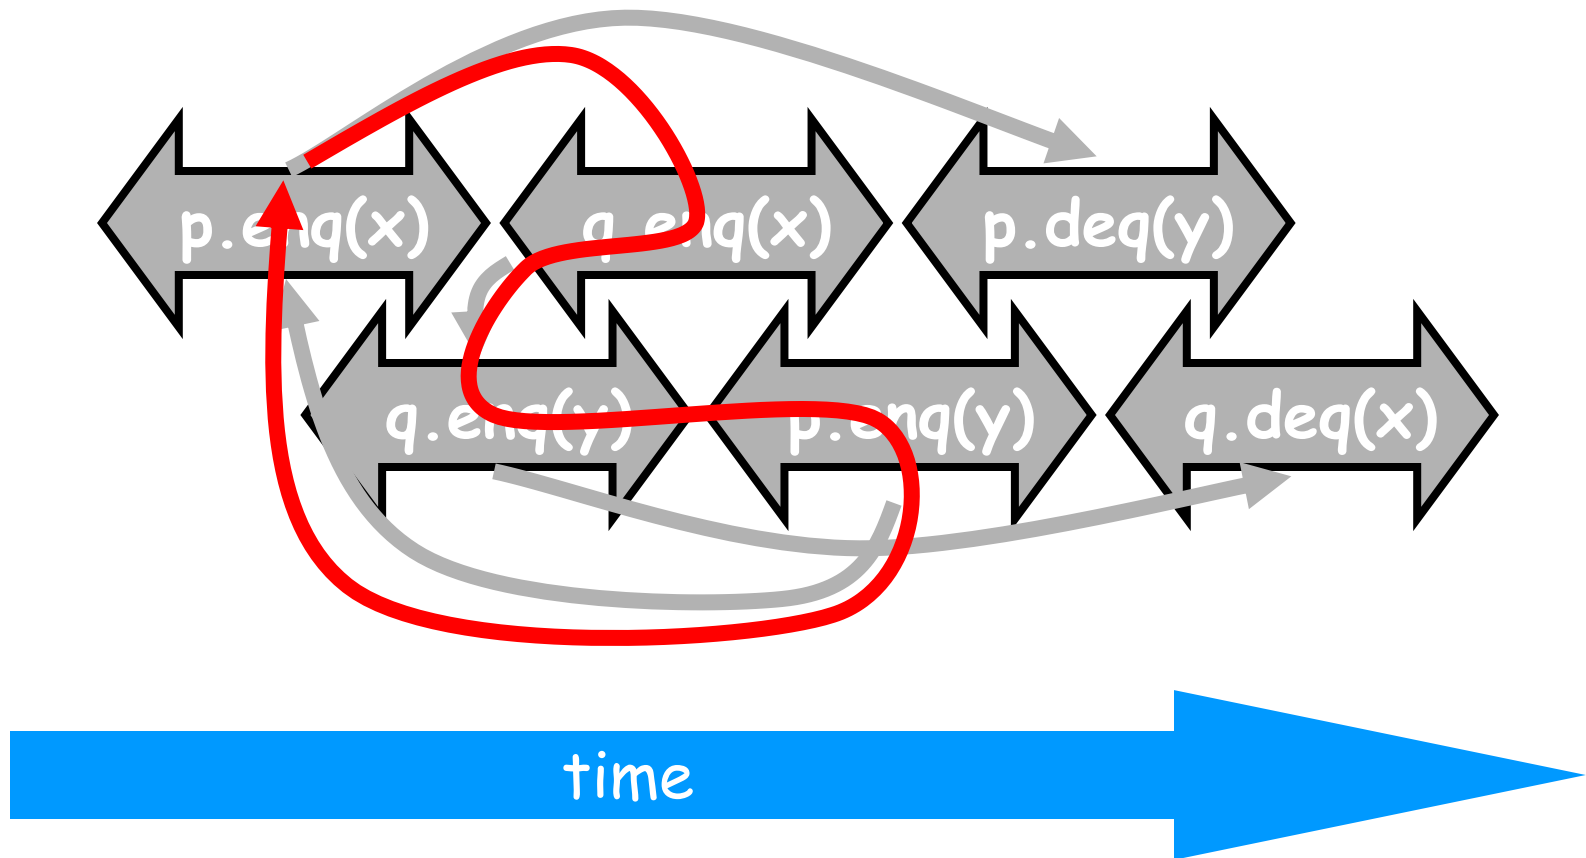
# Ordering imposed by q



# Ordering imposed by both



# Combining orders







# Fact

- Most hardware architectures don't support sequential consistency
- Because they think it's too strong
- Here's another story ...

# The Flag Example




# The Flag Example




- Each thread's view is sequentially consistent
  - It went first

# The Flag Example



- Entire history isn't sequentially consistent
  - Can't both go first

# The Flag Example



- Is this behavior really so wrong?
  - We can argue either way ...

# Opinion1: It's Wrong

- This pattern
  - Write mine, read yours
- Is exactly the flag principle
  - Beloved of Alice and Bob
  - Heart of mutual exclusion
    - Peterson
    - Bakery, etc.
- It's non-negotiable!

# Opinion2: But It Feels So Right

...

- Many hardware architects think that sequential consistency is too strong
- Too expensive to implement in modern hardware
- OK if flag principle
  - violated by default
  - Honored by explicit request

# Memory Hierarchy

- On modern multiprocessors, processors do not read and write directly to memory.
- Memory accesses are very slow compared to processor speeds,
- Instead, each processor reads and writes directly to a **cache**



# Memory Operations

- To read a memory location,
  - load data into cache.
- To write a memory location
  - update cached copy,
  - Lazily write cached data back to memory

# While Writing to Memory

- A processor can execute hundreds, or even thousands of instructions
- Why delay on every memory write?
- Instead, write back in parallel with rest of the program.

# Revisionist History

- Flag violation history is actually OK
  - processors delay writing to memory
  - Until after reads have been issued.
- Otherwise unacceptable delay between read and write instructions.
- Who knew you wanted to synchronize?

# Who knew you wanted to synchronize?

- Writing to memory = mailing a letter
- Vast majority of reads & writes
  - Not for synchronization
  - No need to idle waiting for post office
- If you want to synchronize
  - Announce it explicitly
  - Pay for it **only** when you need it

# Explicit Synchronization

- Memory barrier instruction
  - Flush unwritten caches
  - Bring caches up to date
- Compilers often do this for you
  - Entering and leaving critical sections
- Expensive

# Volatile

- In Java, can ask compiler to keep a variable up-to-date with `volatile` keyword
- Also inhibits reordering, removing from loops, & other “optimizations”

# Real-World Hardware Memory

- Weaker than sequential consistency
- But you can get sequential consistency at a price
- OK for expert, tricky stuff
  - assembly language, device drivers, etc.
- Linearizability more appropriate for high-level software

# Critical Sections

- Easy way to implement linearizability
  - Take sequential object
  - Make each method a critical section
- Like synchronized methods in Java™
- Problems
  - Blocking
  - No concurrency



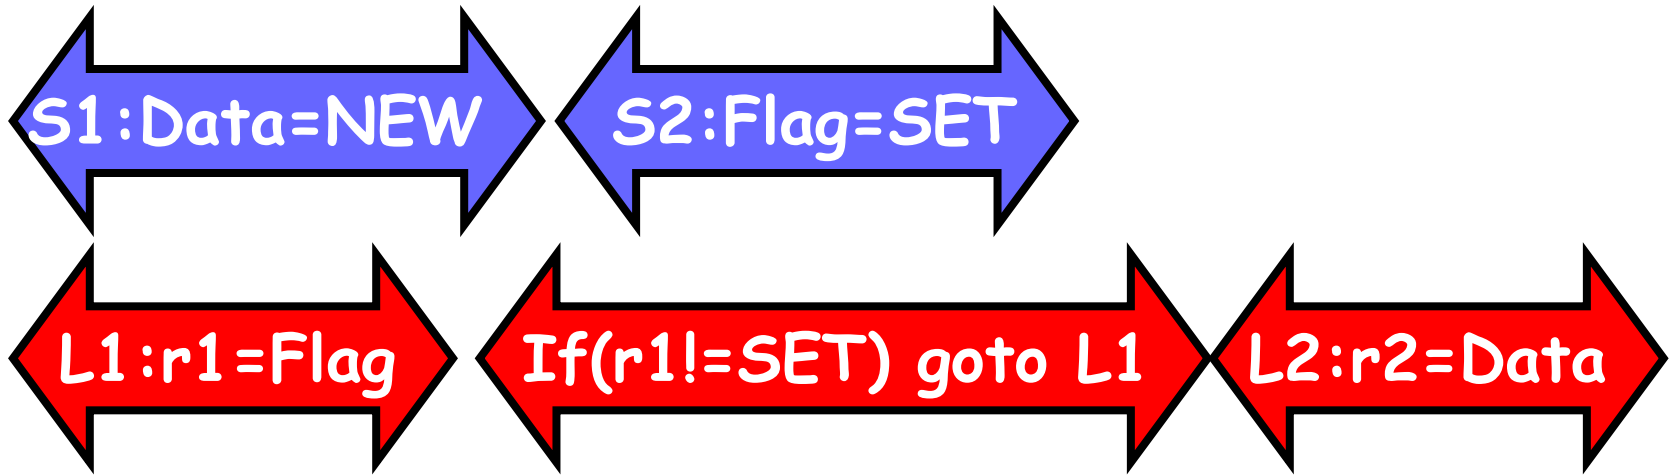
# Summary

- Linearizability
  - Operation takes effect instantaneously between invocation and response
  - Uses sequential specification, locality implies composability
  - Good for high level objects

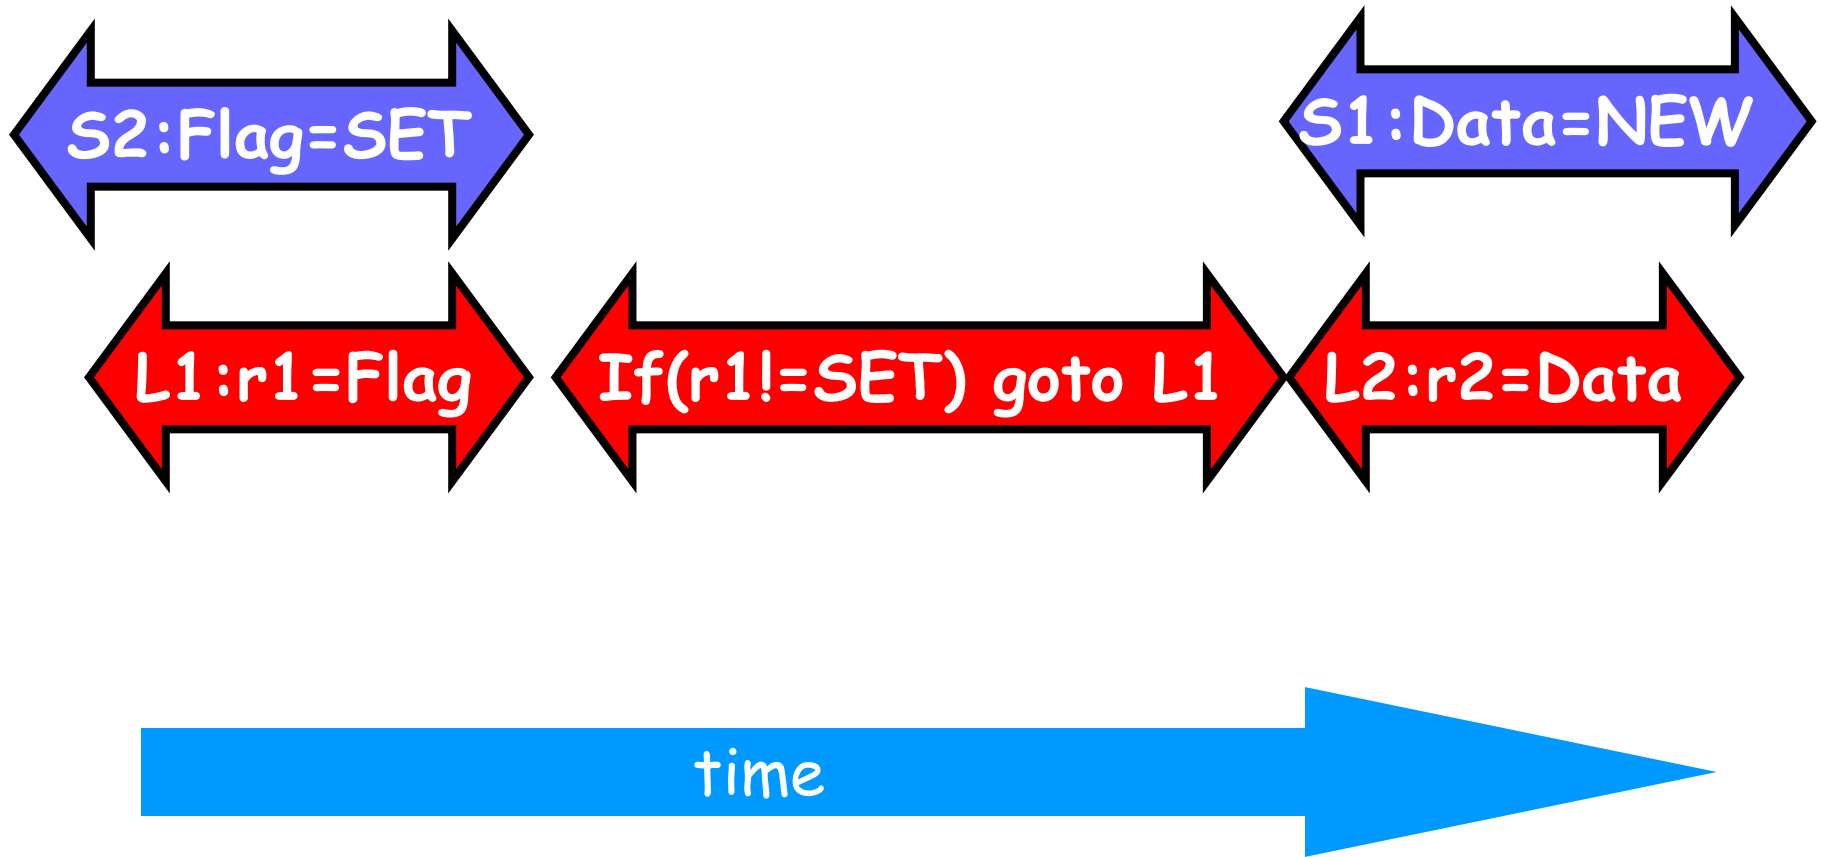
# Summary

- Sequential Consistency
  - Not composable
  - Harder to work with
  - Good way to think about hardware models
- We will use *linearizability* in the remainder of this course unless stated otherwise

# With multi core, What happens ?



# S1 and S2 can be reordered



# Dekker's algorithm

Initially,  $x=0$  &  $y=0$

Core C1

Core C2

S1:  $x=New$ ; S2:  $y=New$ ;


L1:  $r1=y$ ; L2:  $r2=x$ ;

Can both  $r1$  and  $r2$  be set to 0 ?


program order (<p) of Core C1

memory order (<m)


program order (<p) of Core C2





(a) SC Execution 1



time

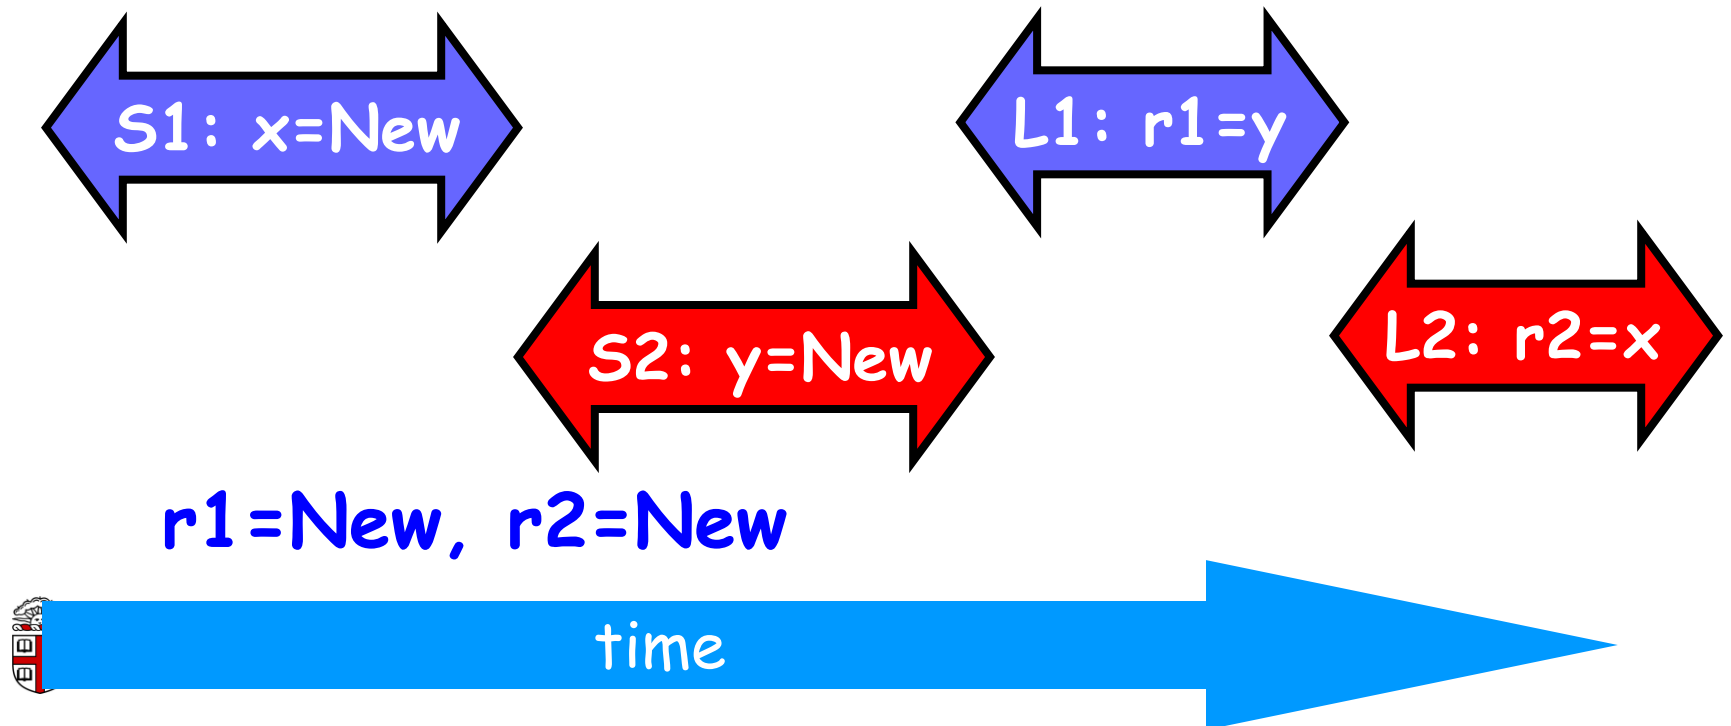
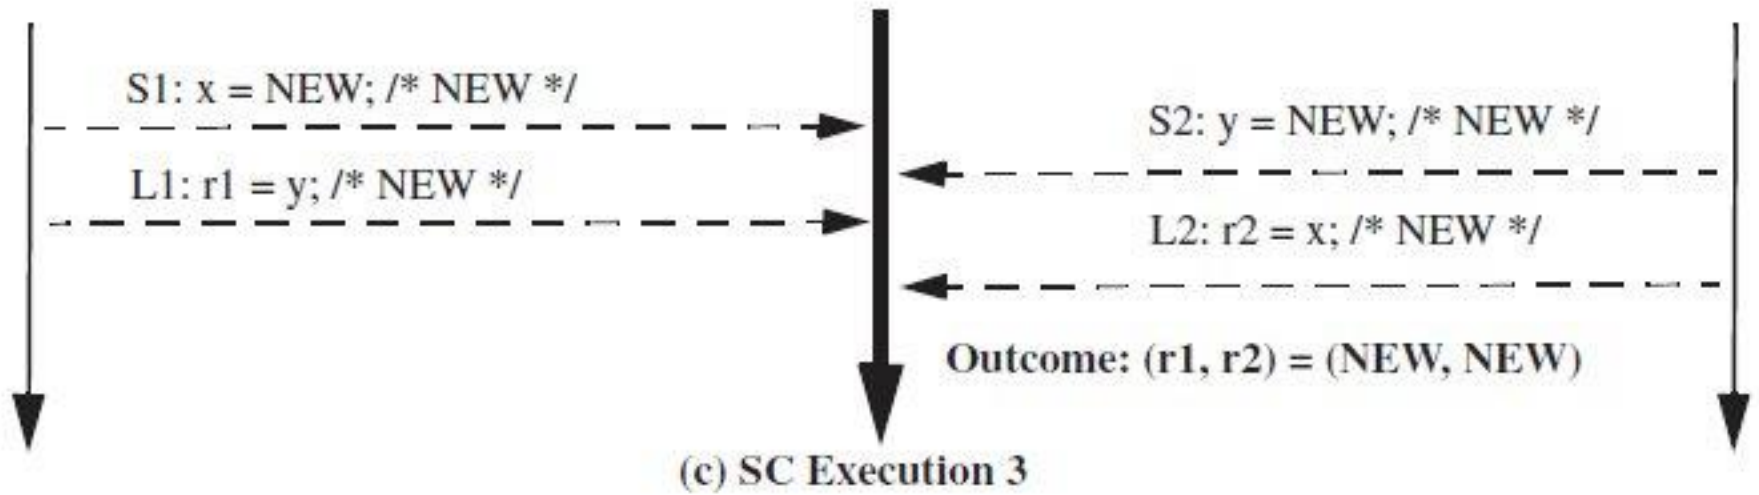


(b) SC Execution 2




`r1=New, r2=0`













(d) NOT an SC Execution

reorder



$r1 = 0, r2 = 0$



Possible for x86, AMD



# An SC execution requires

- All cores insert their loads and stores into the order  $\langle m \rangle$  respecting their program order, regardless of whether they are to the same or different addresses
- Every load gets its value from the last store before it to the same address

# SC ordering Rules

	Operation 2		
Operation 1	Load	Store	RMW
Load	X	X	X
Store	X	X	X
RMW	X	X	X

# Total Store Order

- SPARC
- X86
- Remnants of the write buffer
- Write takes longer
  - It sometimes does, too
  - Needs to get the write permission

# Total Store Order

- Load -> Load
  - Load -> Store
  - Store -> Store
  - Store -> Load (omitted for TSO)
- 
- Omitting the 4<sup>th</sup> constraint allows each core to use a write buffer

# TSO behavior

Initially,  $x=0$  &  $y=0$

Core C1

S1:  $x=New;$

L1:  $r1=x;$

L2:  $r2=y;$

Core C2


S2:  $y=New;$

L3:  $r3=y;$

L4:  $r4=x;$

If  $r2=0$  &  $r4=0$ , can  $r1$  or  $r3$  be set to 0 ?

# Bypassing



**r1=New & r3=New !!!**

# To make TSO sequentially consistent

- Use FENCE
- FENCE ensures the memory operations before the FENCE get placed before the memory operations after the FENCE on the core



# TSO ordering Rules

	Operation 2				
Operation 1	Load	Store	RMW	FENCE	
Load	X	X	X	X	
Store	B	X	X	X	
RMW	X	X	X	X	
FENCE	X	X	X	X	

# This work is licensed under a [Creative Commons Attribution-ShareAlike 2.5 License](https://creativecommons.org/licenses/by-sa/3.0/).

- **You are free:**
  - **to Share** — to copy, distribute and transmit the work
  - **to Remix** — to adapt the work
- **Under the following conditions:**
  - **Attribution.** You must attribute the work to “The Art of Multiprocessor Programming” (but not in any way that suggests that the authors endorse you or your use of the work).
  - **Share Alike.** If you alter, transform, or build upon this work, you may distribute the resulting work only under the same, similar or a compatible license.
- For any reuse or distribution, you must make clear to others the license terms of this work. The best way to do this is with a link to
  - <http://creativecommons.org/licenses/by-sa/3.0/>.
- Any of the above conditions can be waived if you get permission from the copyright holder.
- Nothing in this license impairs or restricts the author's moral rights.