

# Reagents:

Functional programming  
meets scalable concurrency

Aaron Turon

Northeastern University

# Concurrency $\neq$ Parallelism

**Concurrency** is overlapped execution of processes.

**Parallelism** is simultaneous execution of computations.

The trouble is that *essentially all the interesting applications of concurrency involve the deliberate and controlled mutation of shared state*, such as screen real estate, the file system, or the internal data structures of the program. The right solution, therefore, is to provide mechanisms which allow (though alas they cannot enforce) the safe mutation of shared state.

-- Peyton Jones, Gordon, and Finne  
in *Concurrent Haskell*

# Concurrency $\cap$ Parallelism = *Scalable* Concurrency

## **Use cases:**

- Concurrent programs on parallel hardware (e.g. OS kernels)
- *Implementing* parallel abstractions (e.g. work stealing for data parallelism)
- “Last mile” of parallel programming (where we must resort to concurrency)

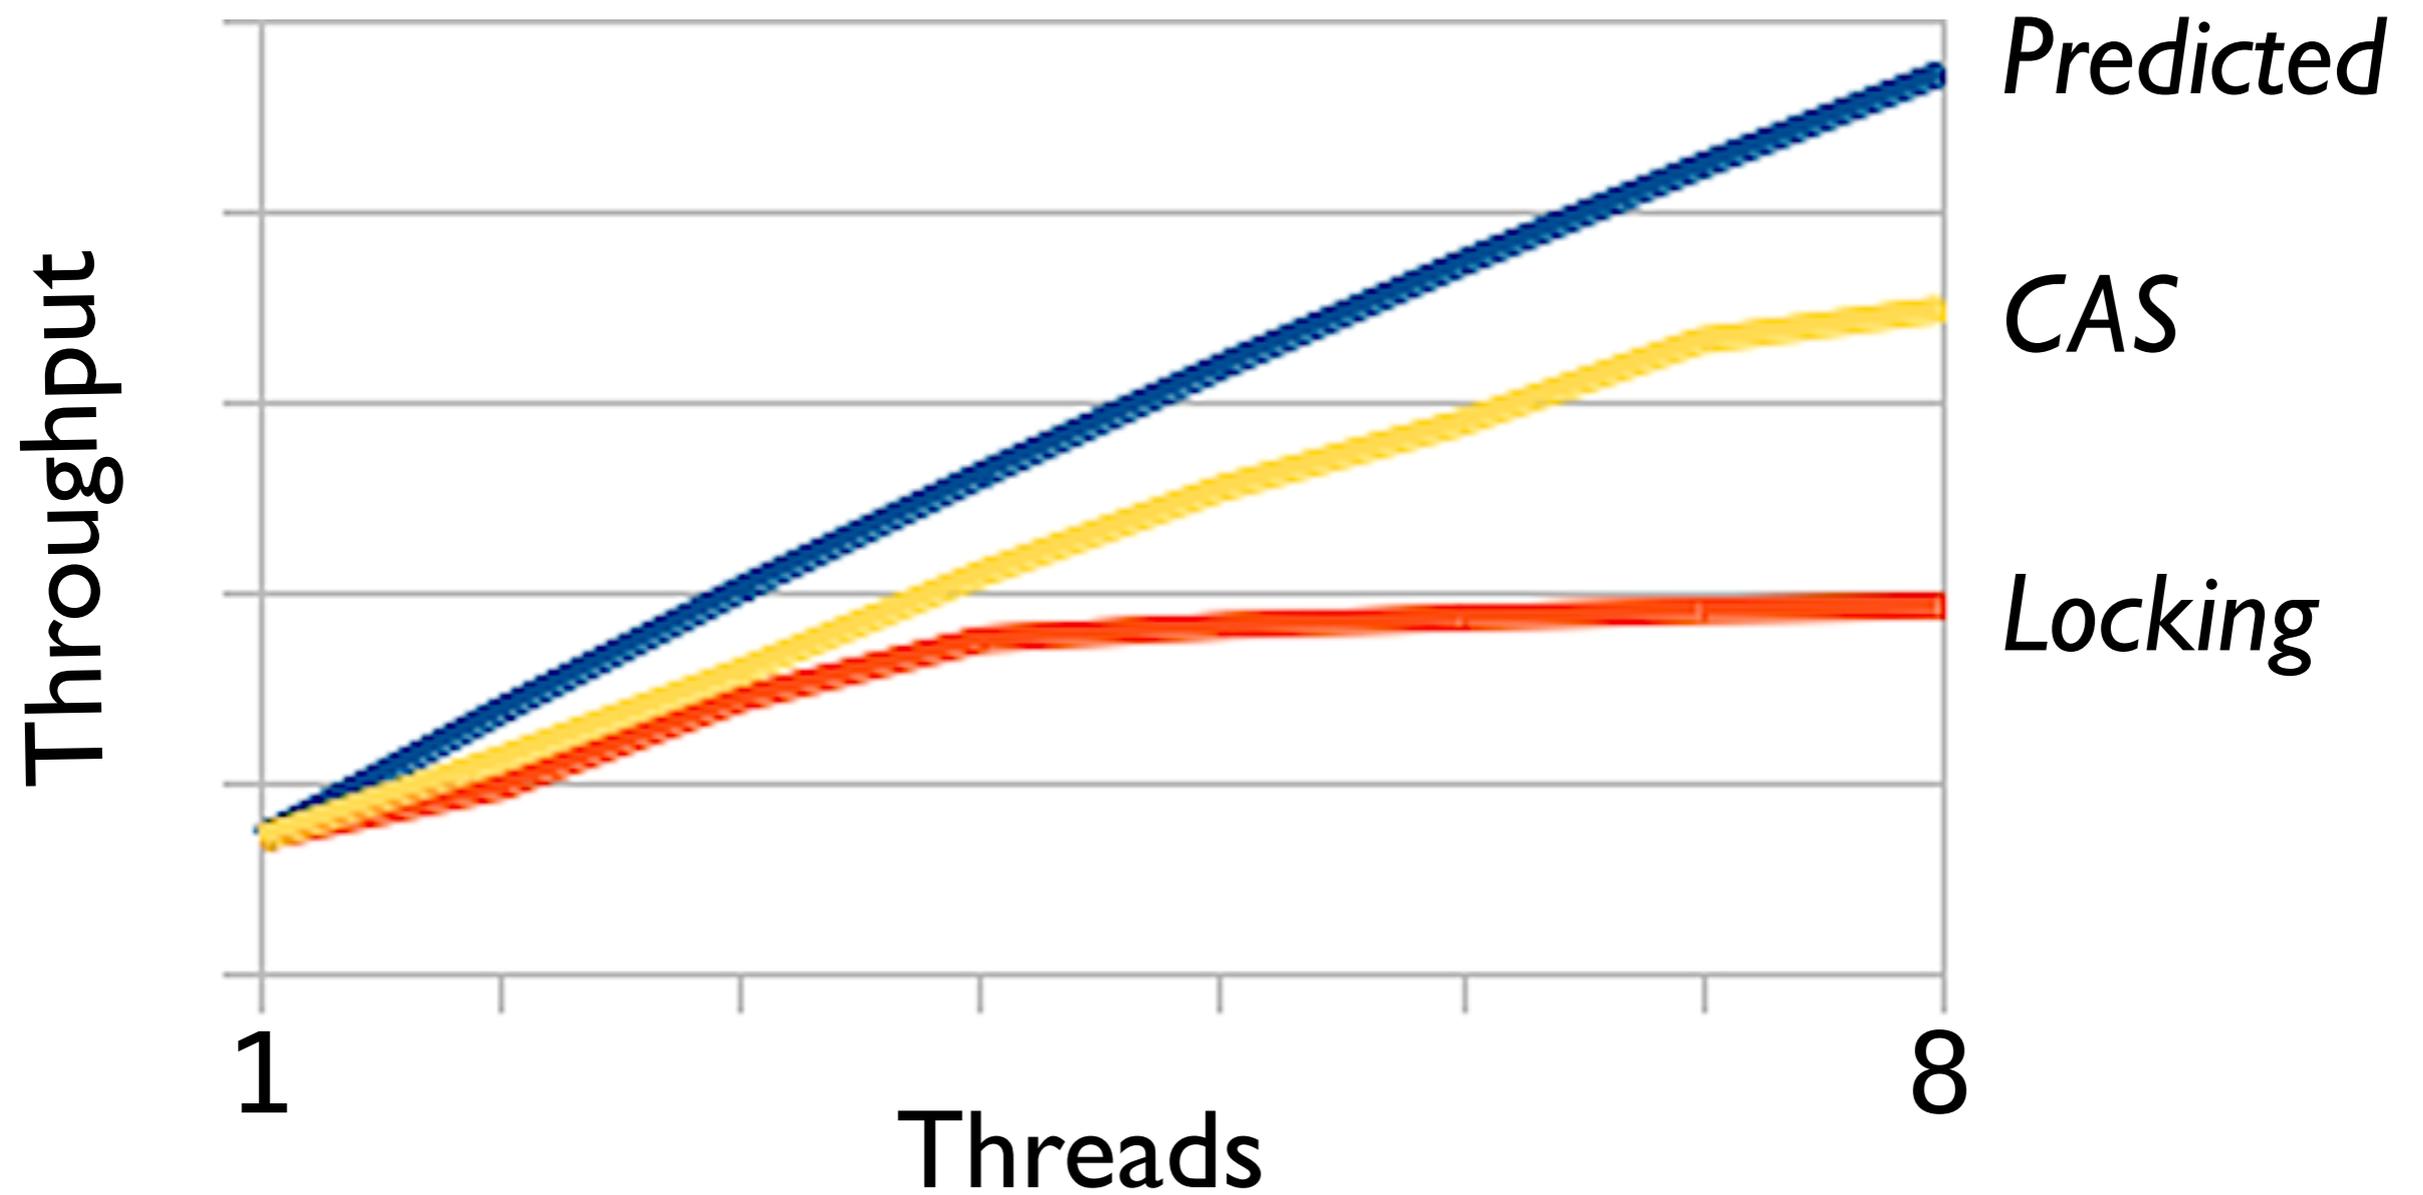
```
class LockCounter {  
  private var c: Int = 0  
  private var l = new Lock  
  def inc: Int = {  
    l.lock()  
    val old = c  
    c = old + 1  
    l.unlock()  
    old  
  }  
}
```

```
class CASCounter {  
  private var c = new AtomicRef[Int](0)  
  def inc: Int = {  
    while (true) {  
      val old = c  
      if (c.cas(old, old+1)) return old  
    }  
  }  
}
```

# A simple test

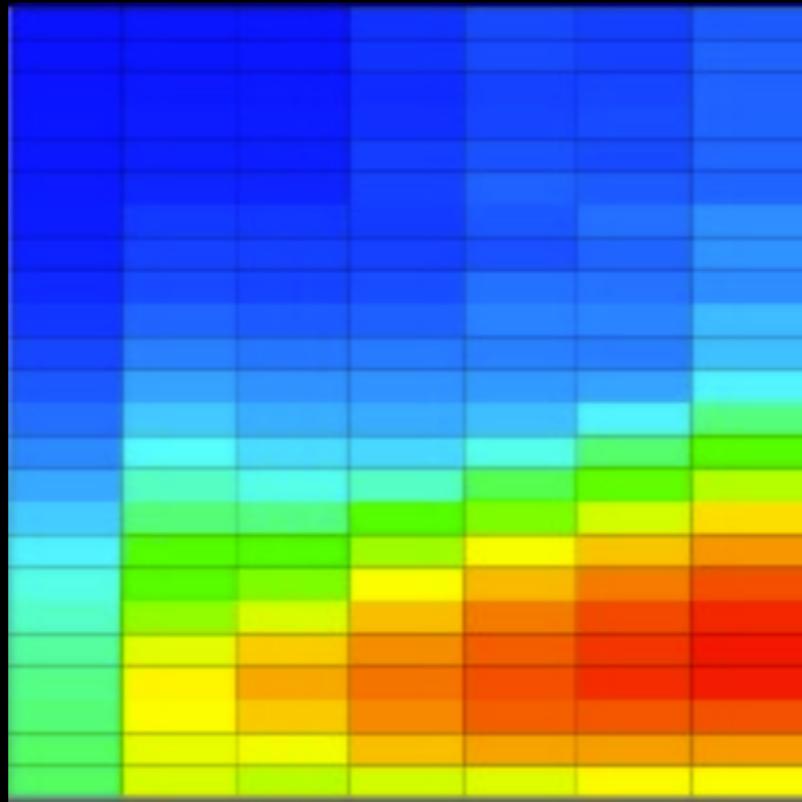
- Increment counter
- Busywait for  $t$  cycles (no cache interaction)
- Repeat

# Results for 98% parallelism



# Lock-based

Parallelism (log-scale)



2 4 6 8

Threads

# CAS-based

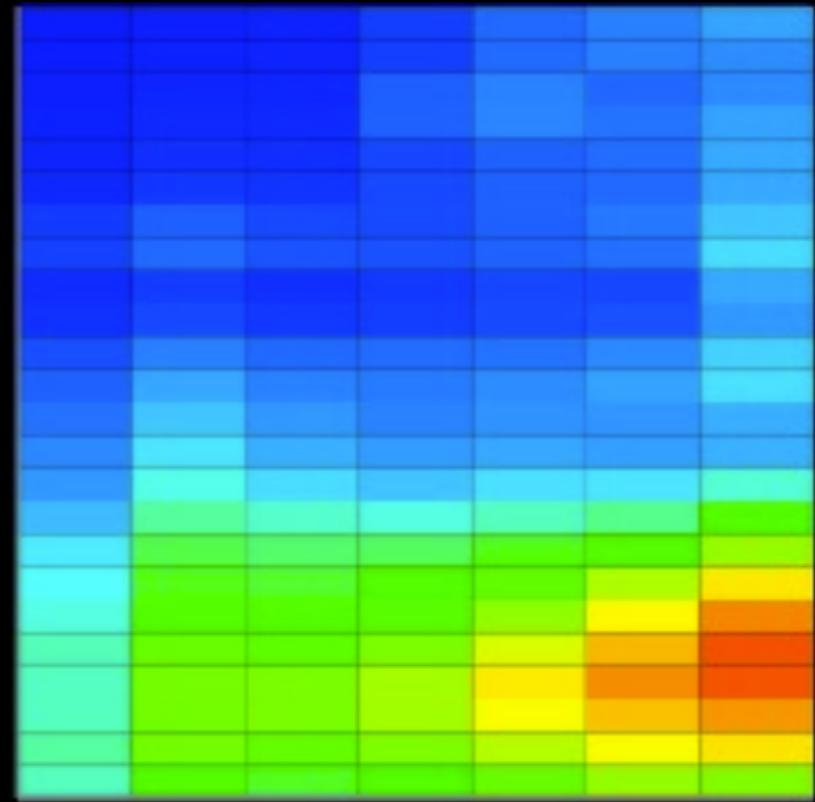
99.9%

99.7%

98%

88%

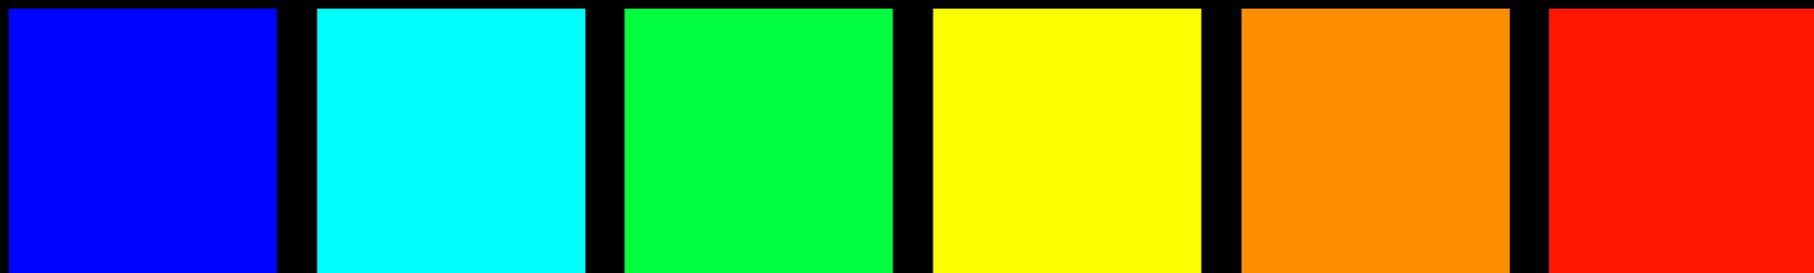
63%



2 4 6 8

Threads

Throughput  
Optimal



1.0

0.87

0.74

0.61

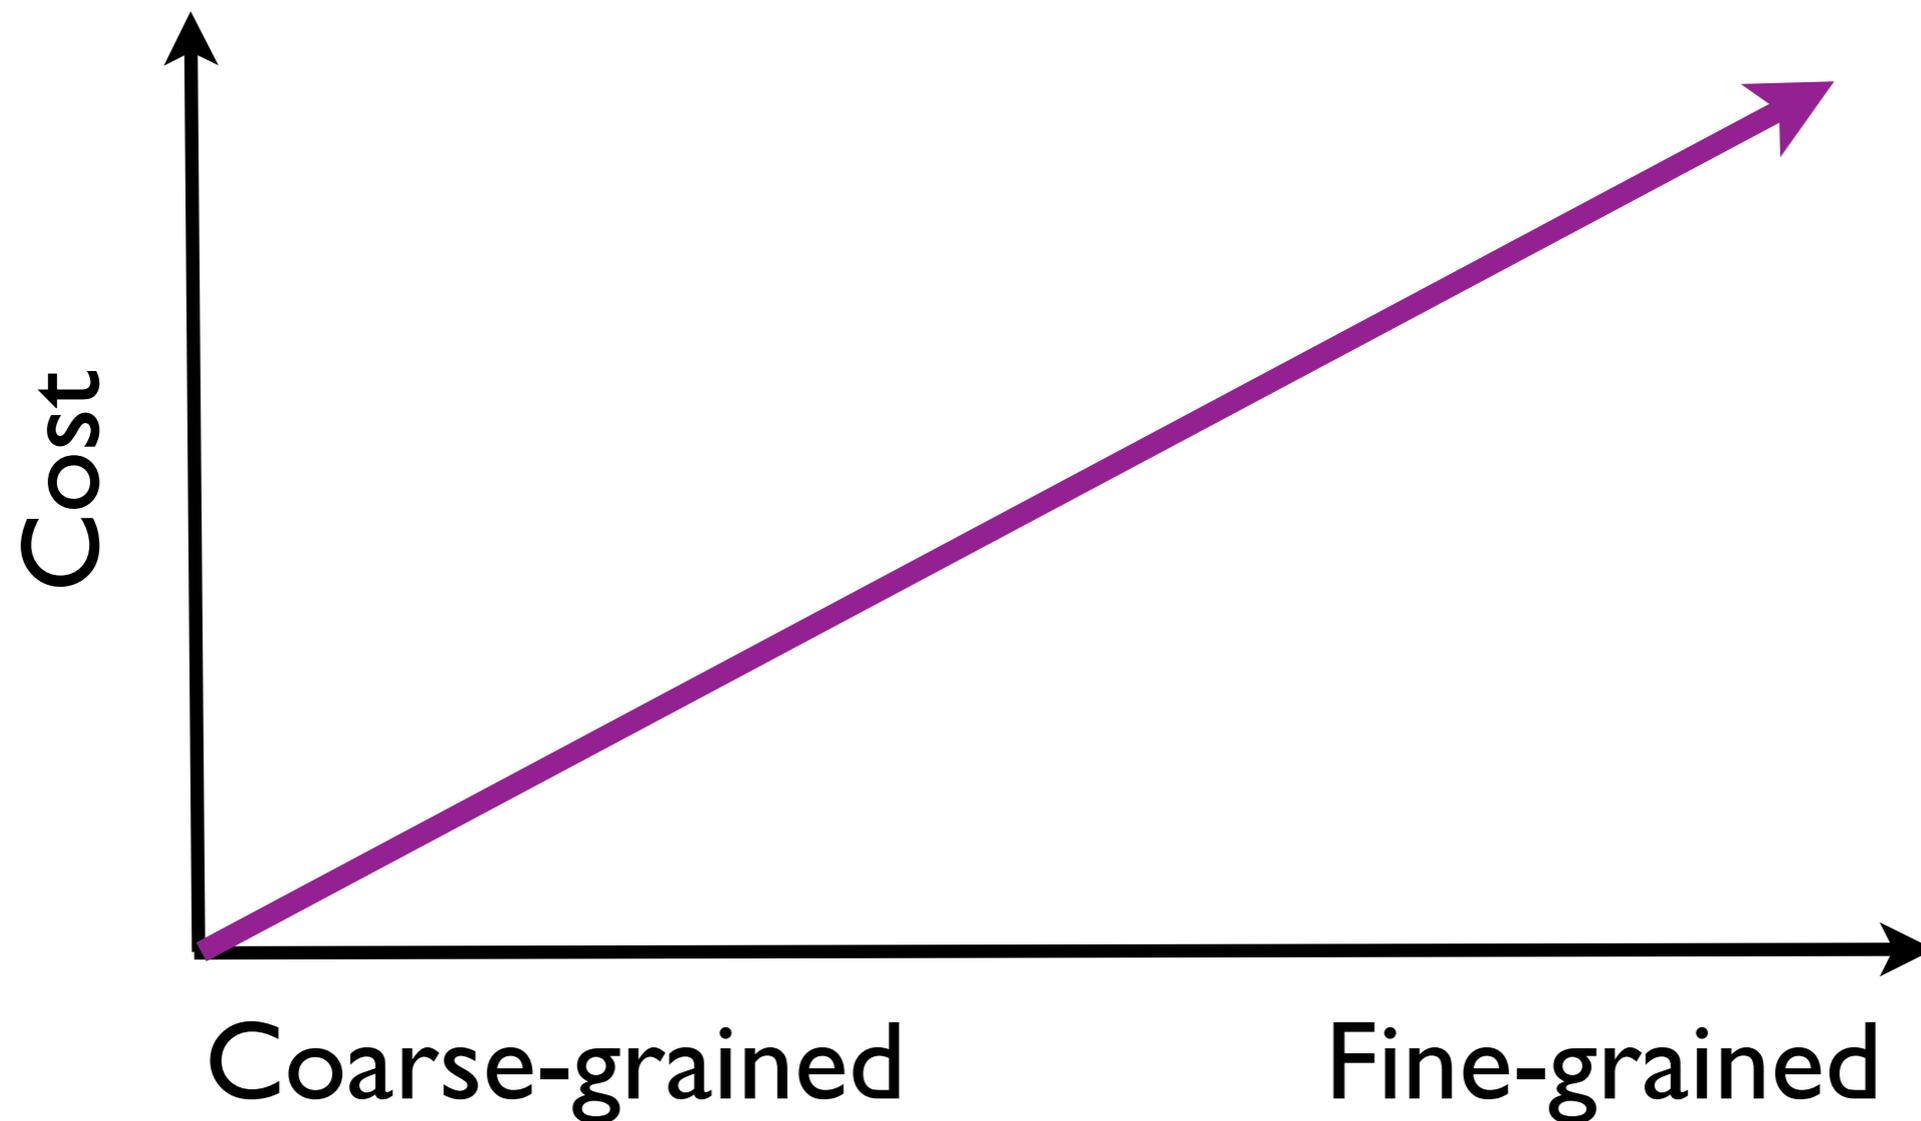
0.48

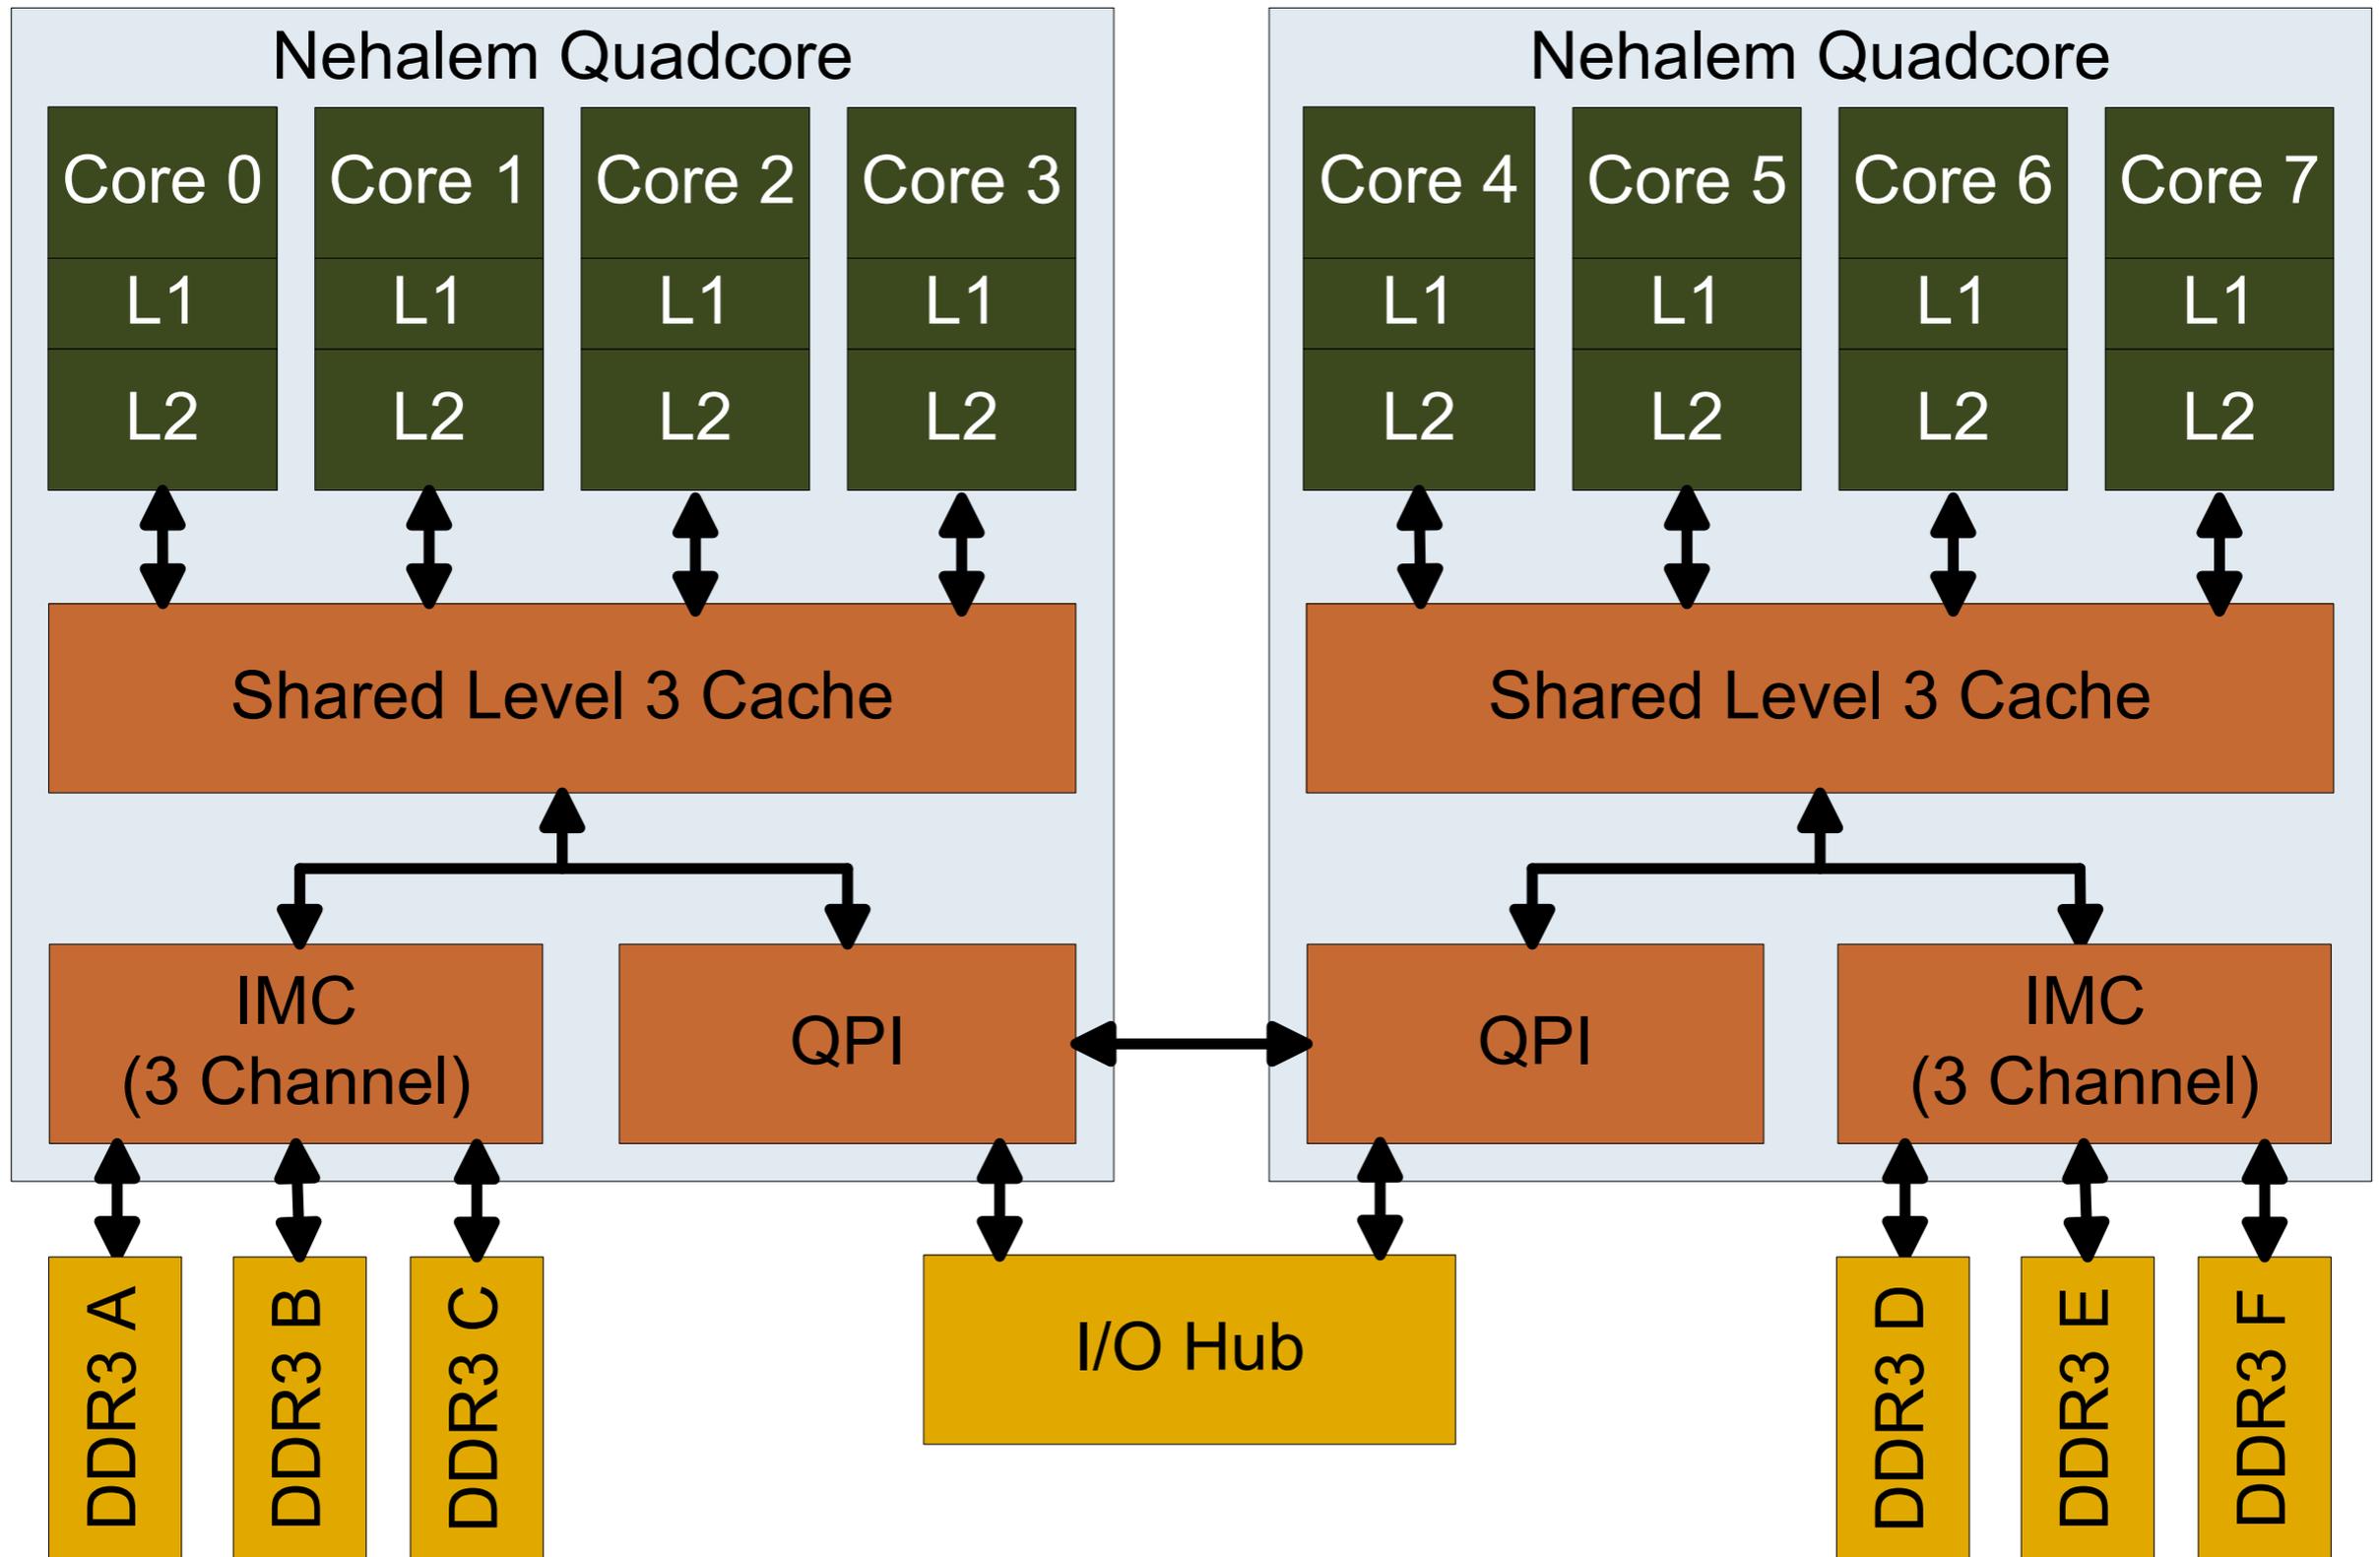
0.35

What's going on here?

# What's going on here?

## **Communication**







# java.util.concurrent

## **Synchronization**

Reentrant locks

Semaphores

R/W locks

Reentrant R/W locks

Condition variables

Countdown latches

Cyclic barriers

Phasers

Exchangers

## **Data structures**

Queues

Nonblocking

Blocking (array & list)

Synchronous

Priority, nonblocking

Priority, blocking

Dequeues

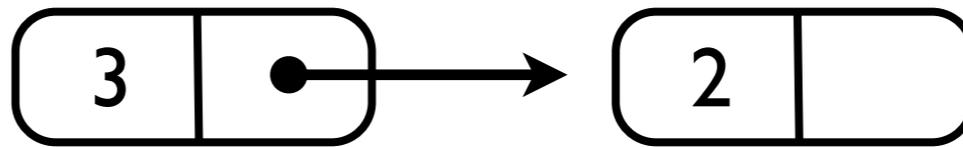
Sets

Maps (hash & skiplist)

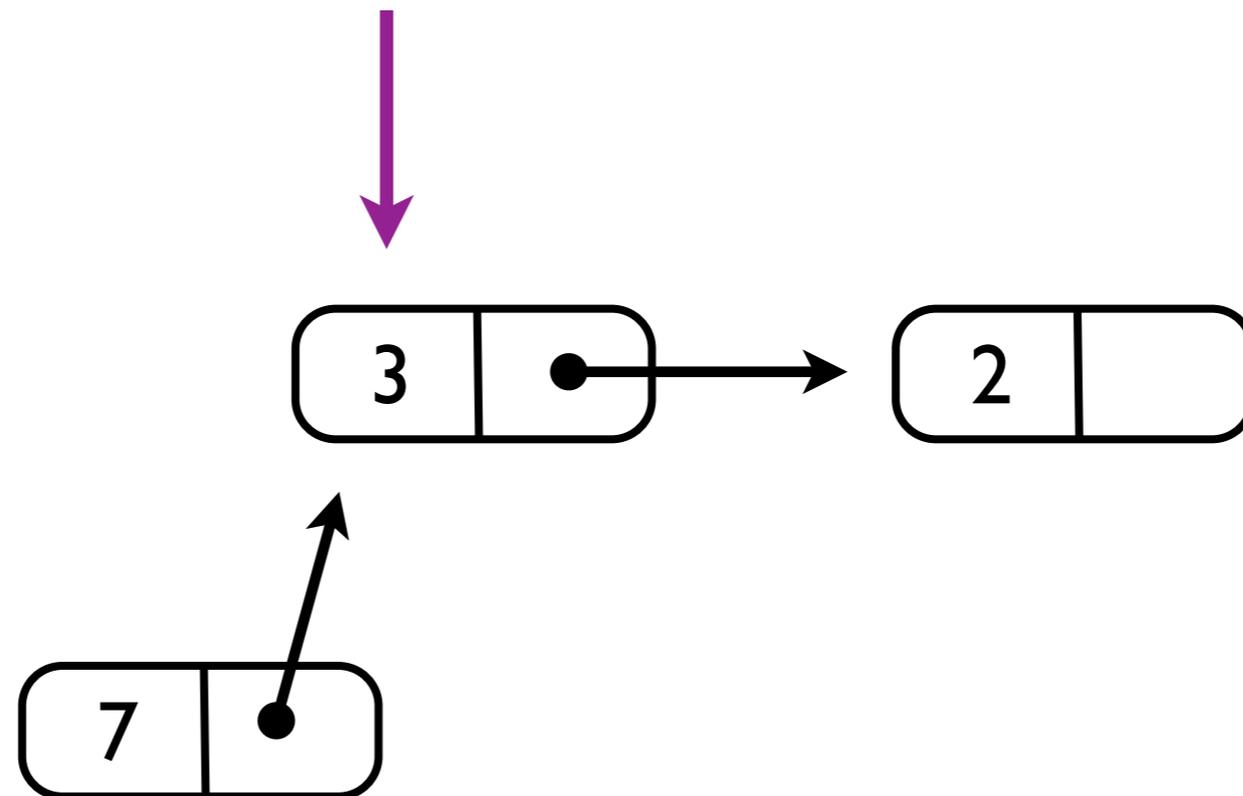
```
class TreiberStack[A] {  
  private val head =  
    new AtomicRef[List[A]](Nil)  
  
  def push(a: A) {  
    val backoff = new Backoff  
    while (true) {  
      val cur = head.get()  
      if (head.cas(cur, a :: cur)) return  
      backoff.once()  
    }  
  }  
}
```

...

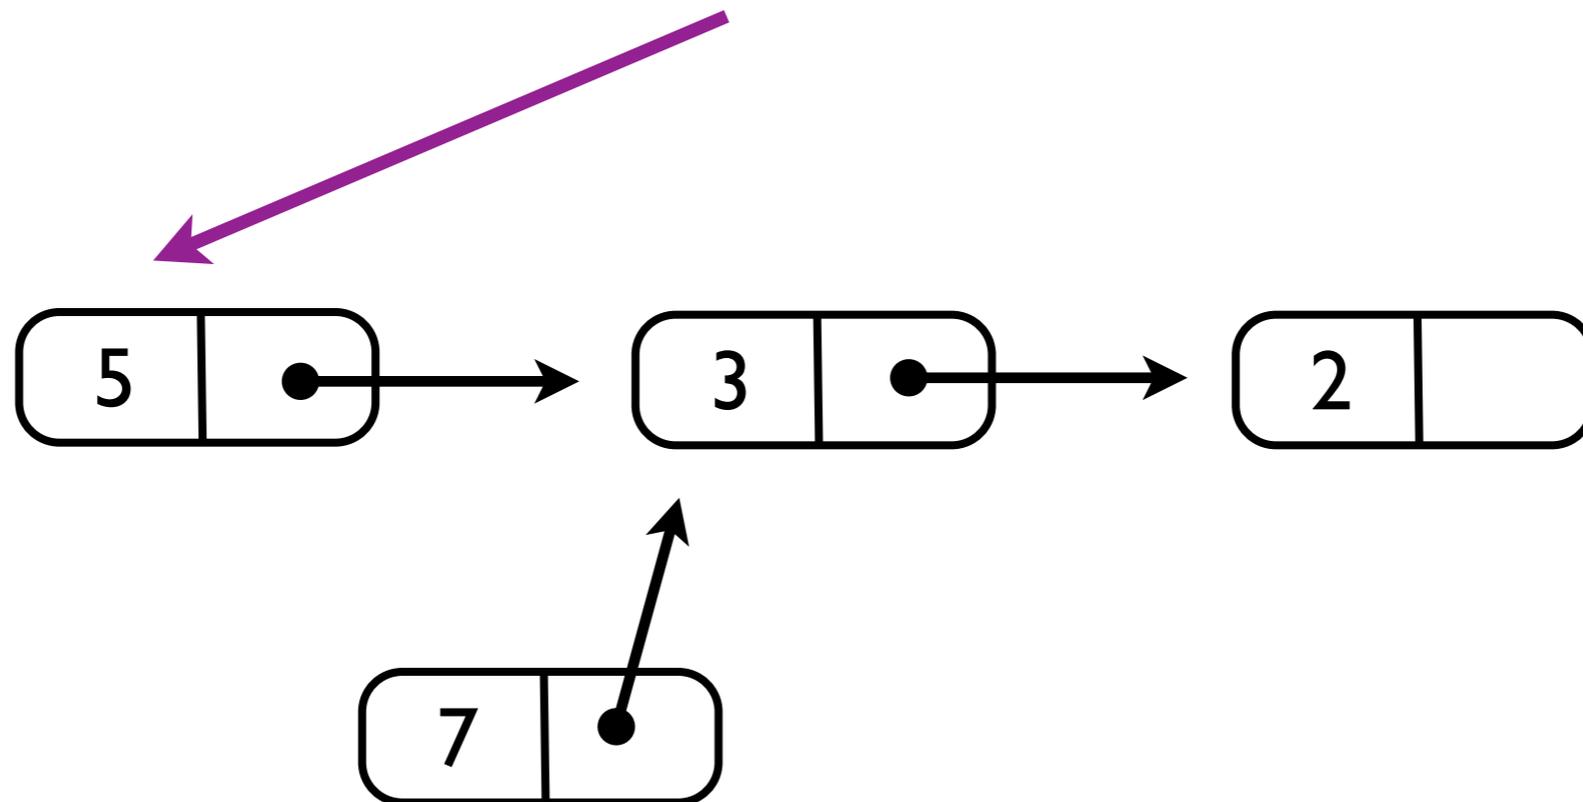
Head

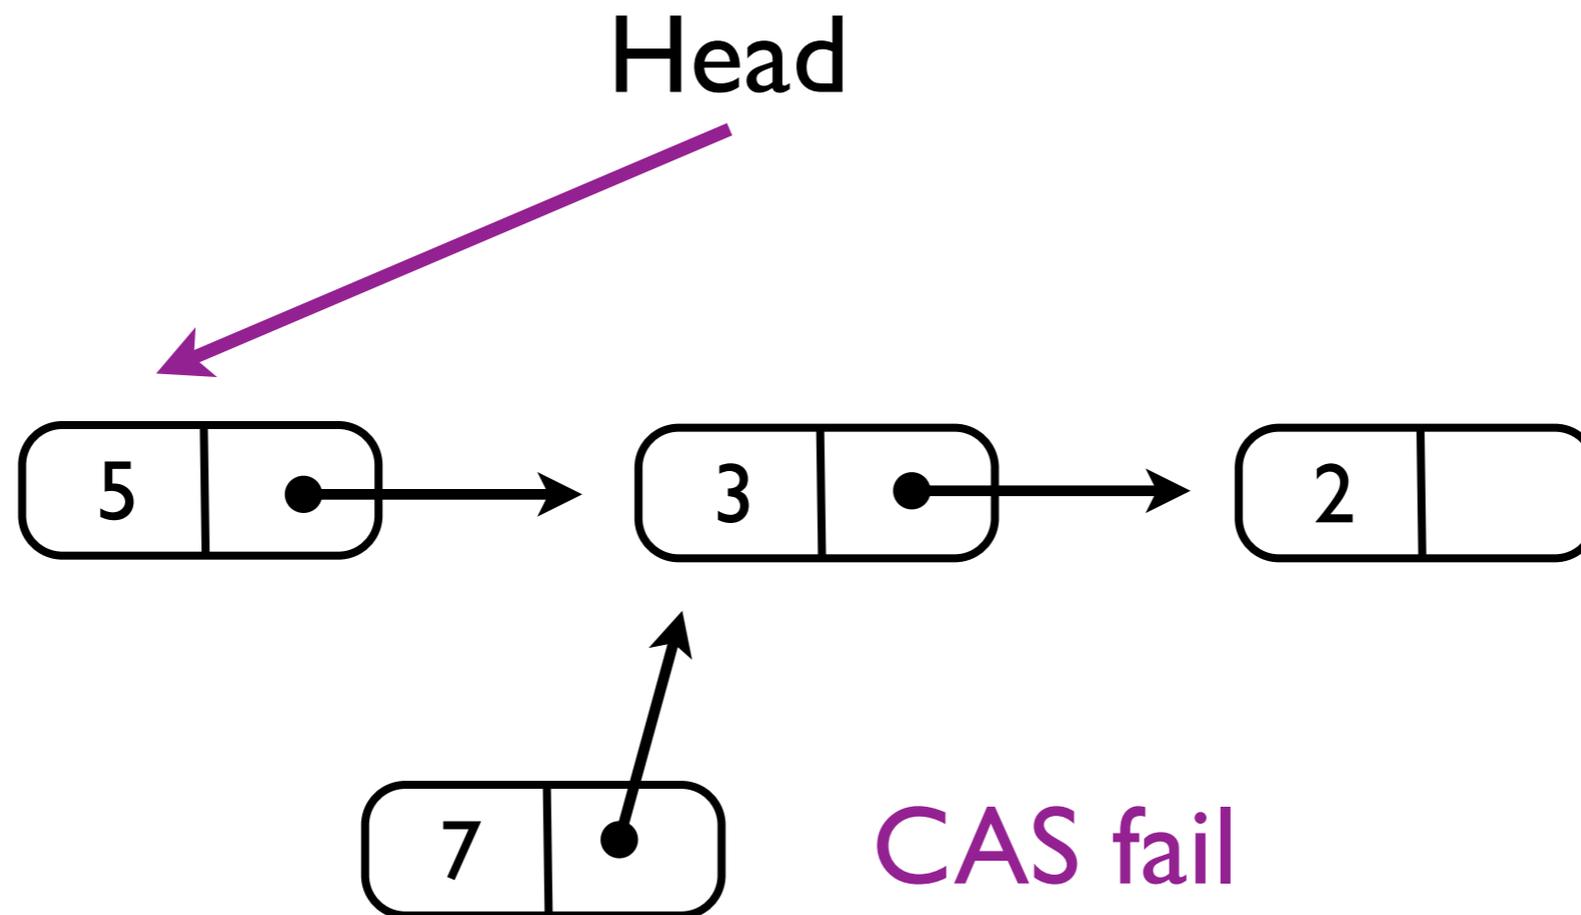


Head

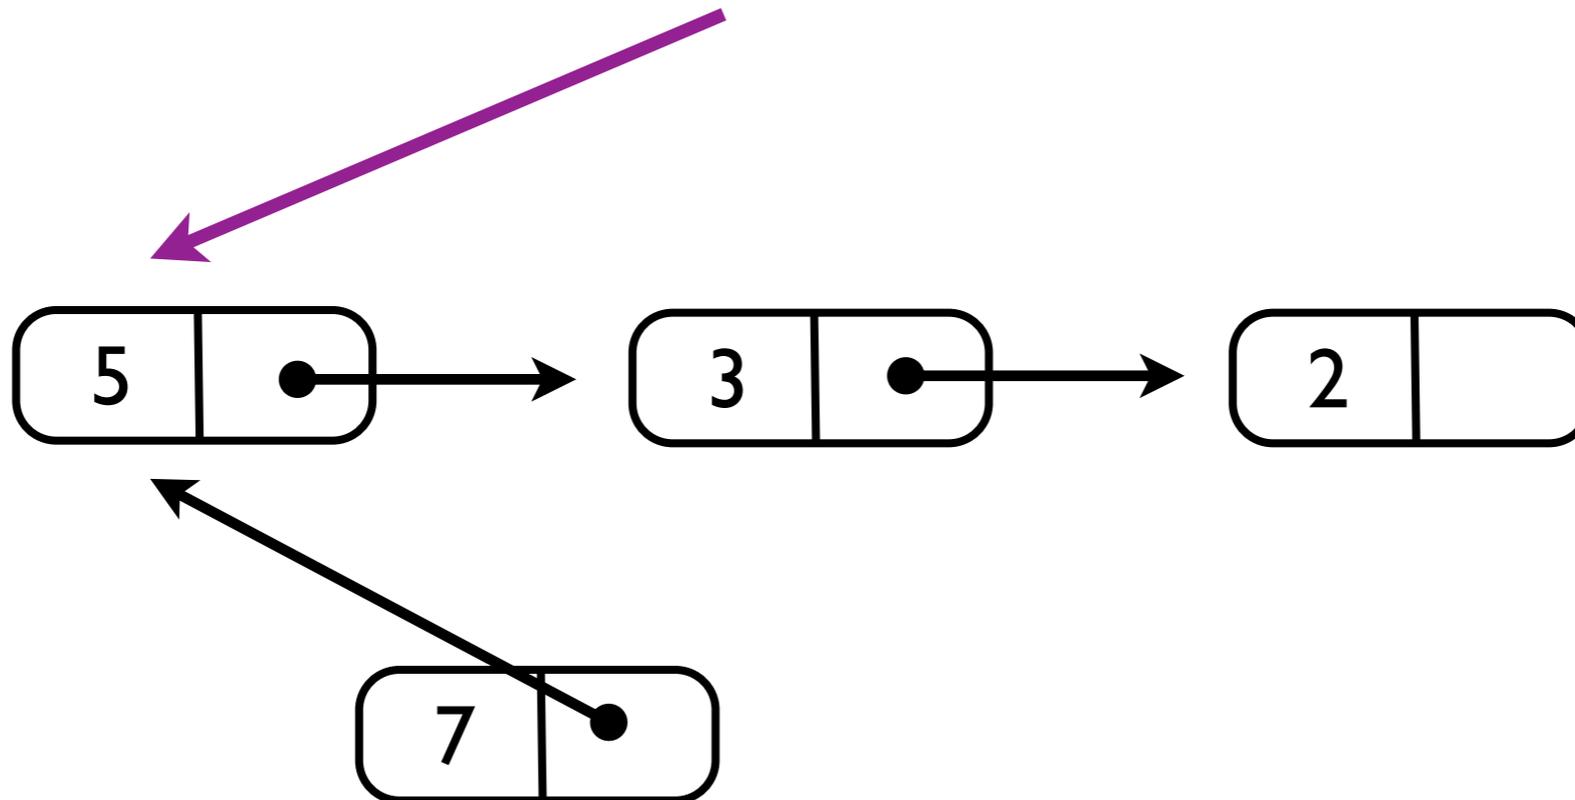


Head

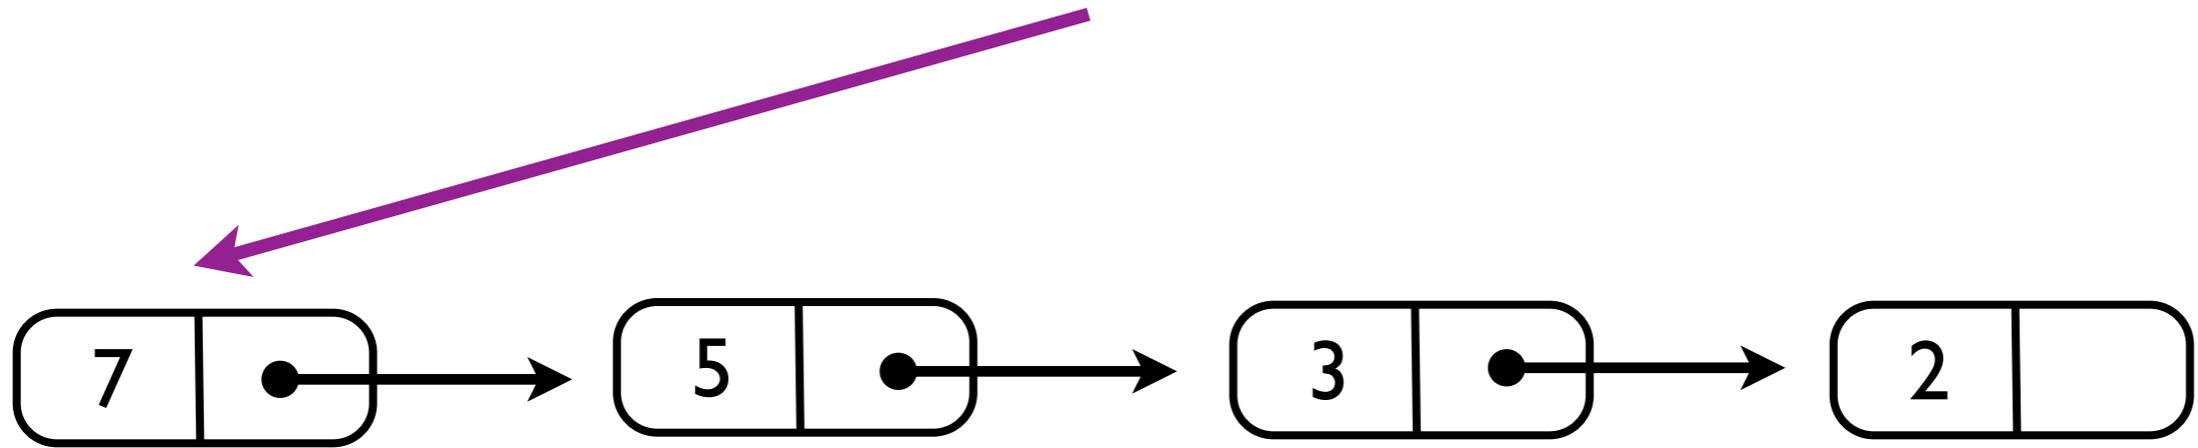




Head



Head



```
def tryPop(): Option[A] = {  
  val backoff = new Backoff  
  while (true) {  
    val cur = head.get()  
    cur match {  
      case Nil => return None  
      case a::tail =>  
        if (head.cas(cur, tail))  
          return Some(a)  
    }  
    backoff.once()  
  }  
}
```

# The Problem:

Concurrency libraries are  
**indispensable**, but hard to  
**build** and **extend**

# The Proposal:

Build and extend scalable  
concurrent algorithms using  
a monad with shared-state and  
message-passing operations

# Design

# Reagents are (first) arrows

Lambda abstraction:

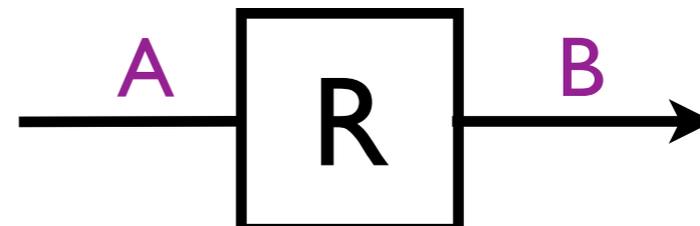


# Reagents are (first) arrows

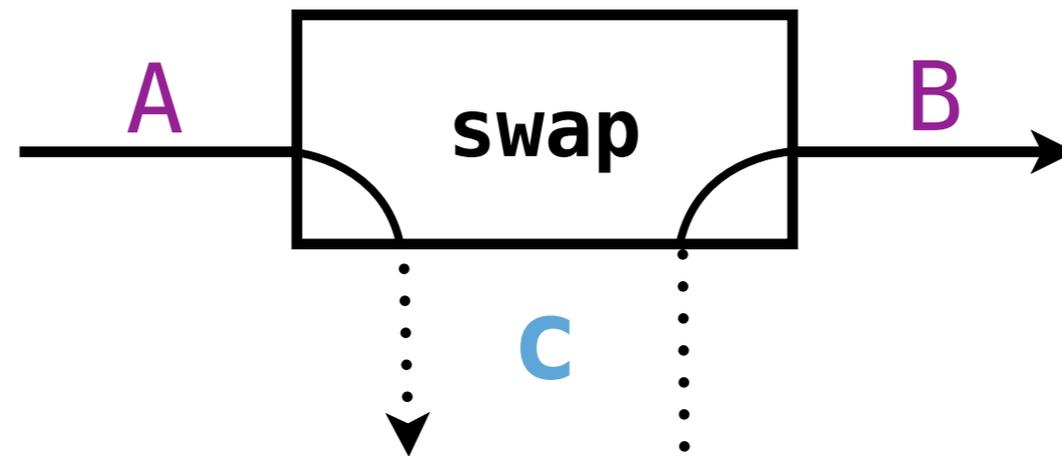
Lambda abstraction:



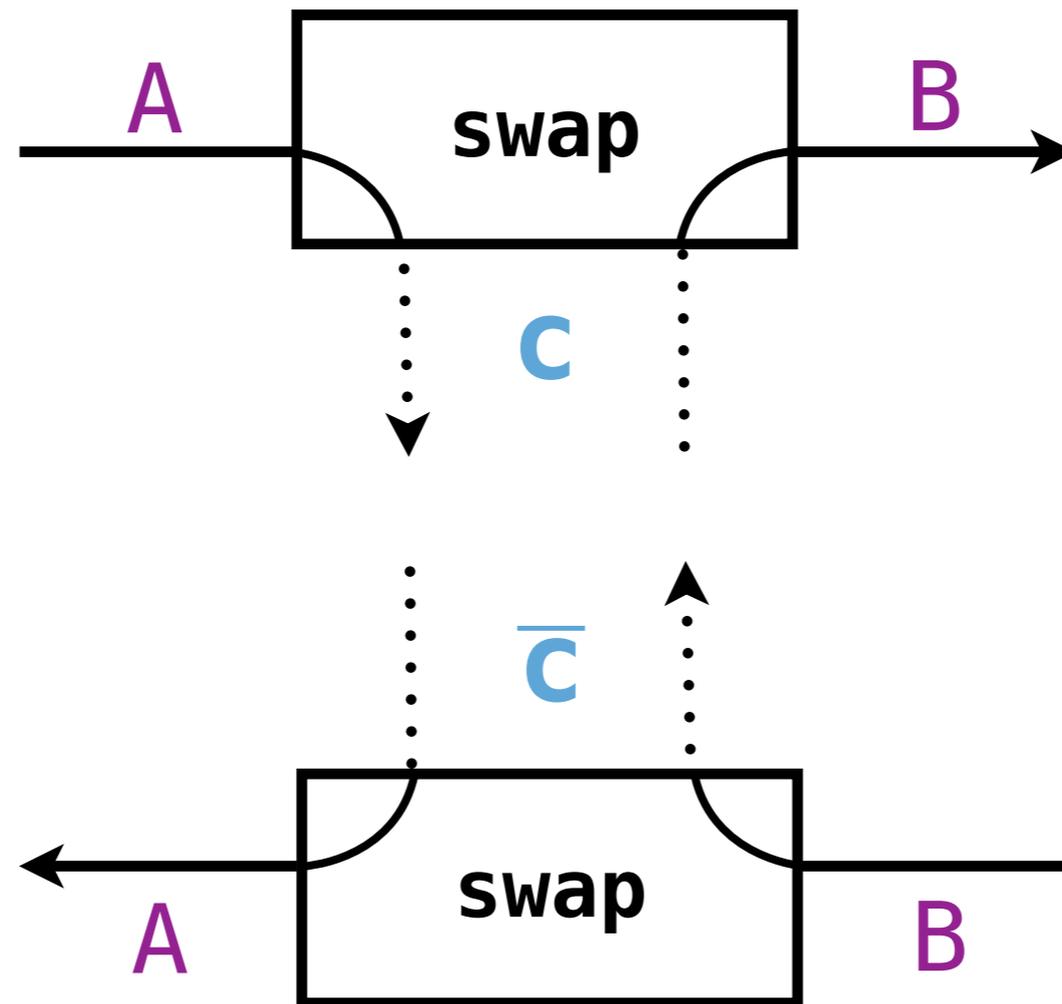
Reagent abstraction:



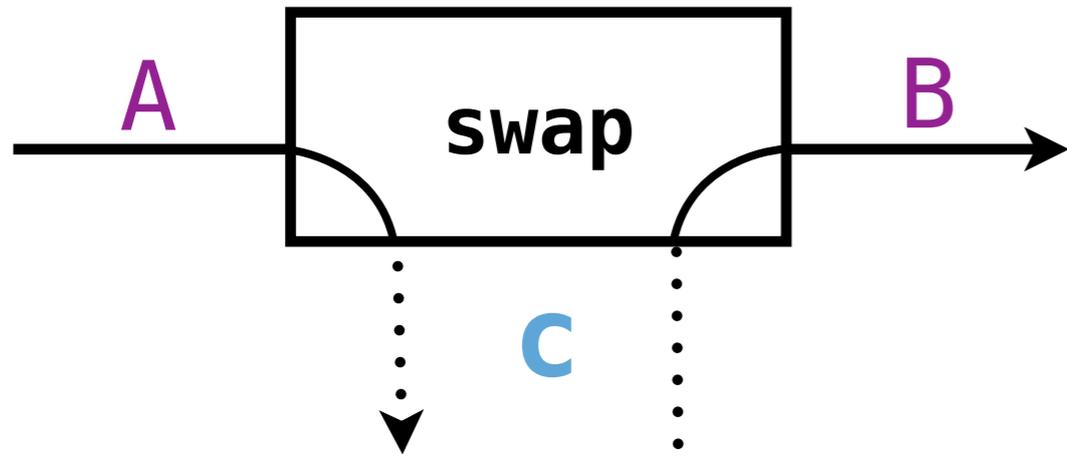
**c:** Chan [A, B]



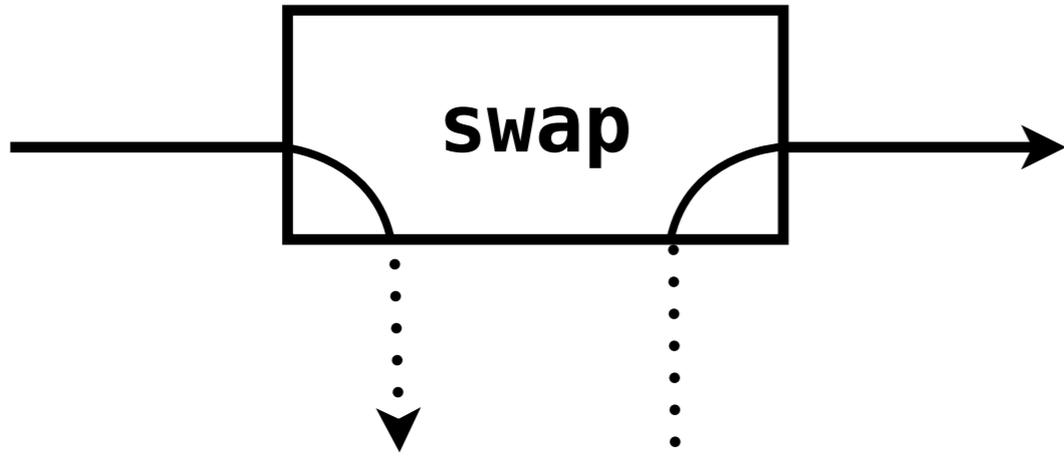
**c**: Chan [A, B]



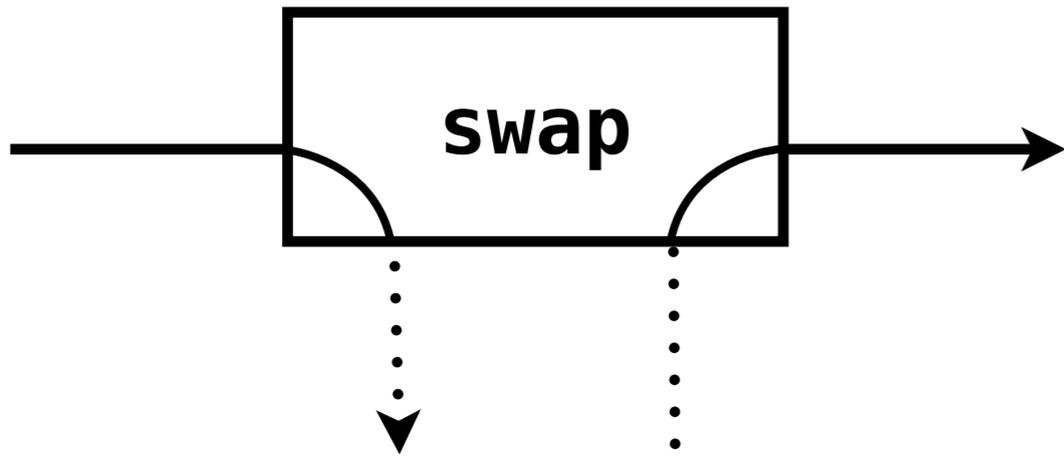
**c:** Chan [A, B]



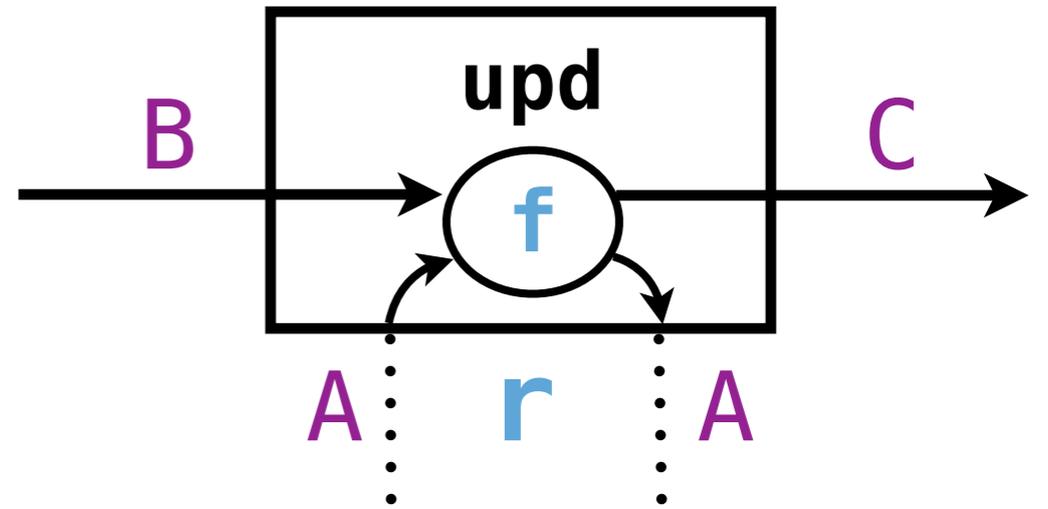
# Message passing



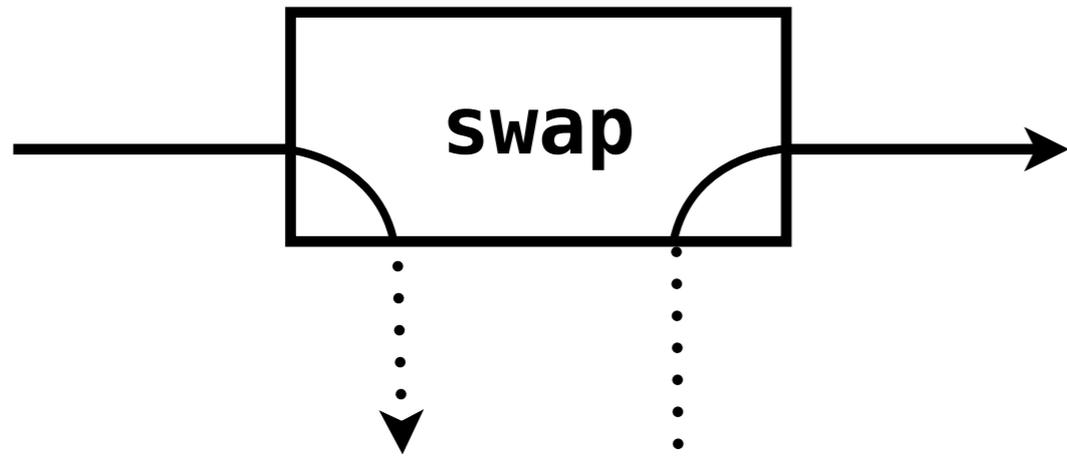
# Message passing



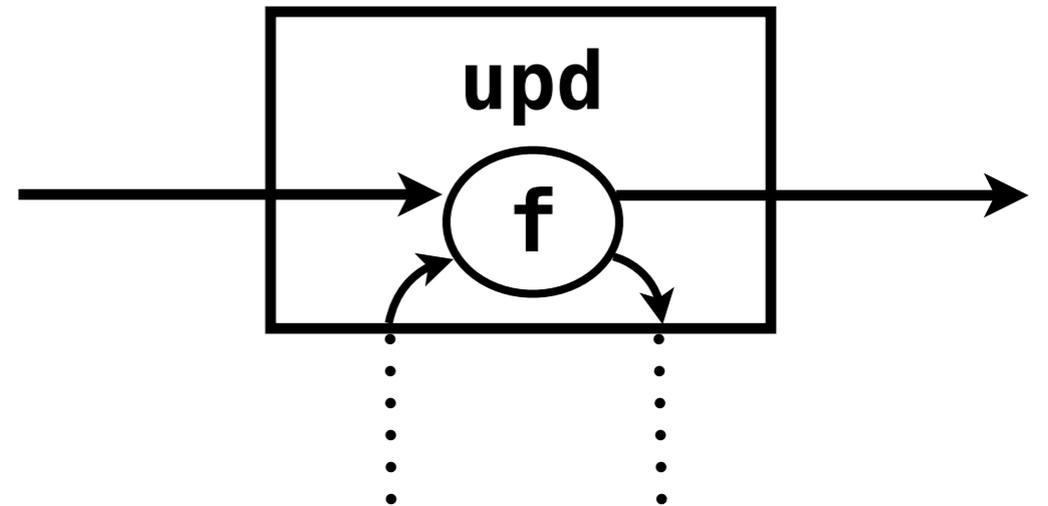
$r$ : Ref[A]  
 $f$ : (A, B)  $\rightarrow$  (A, C)



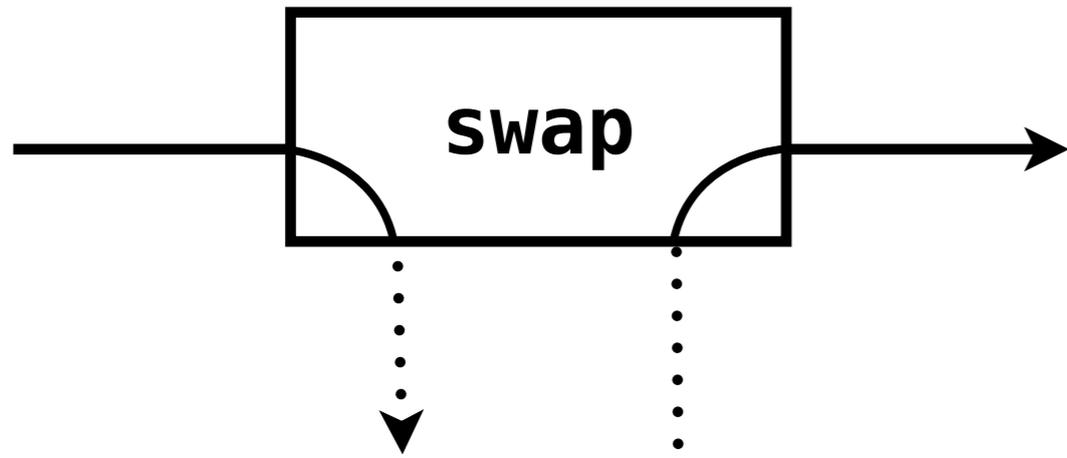
# Message passing



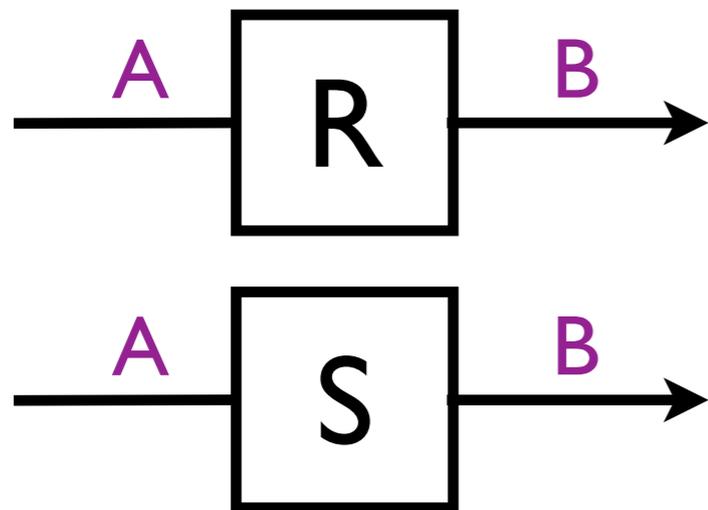
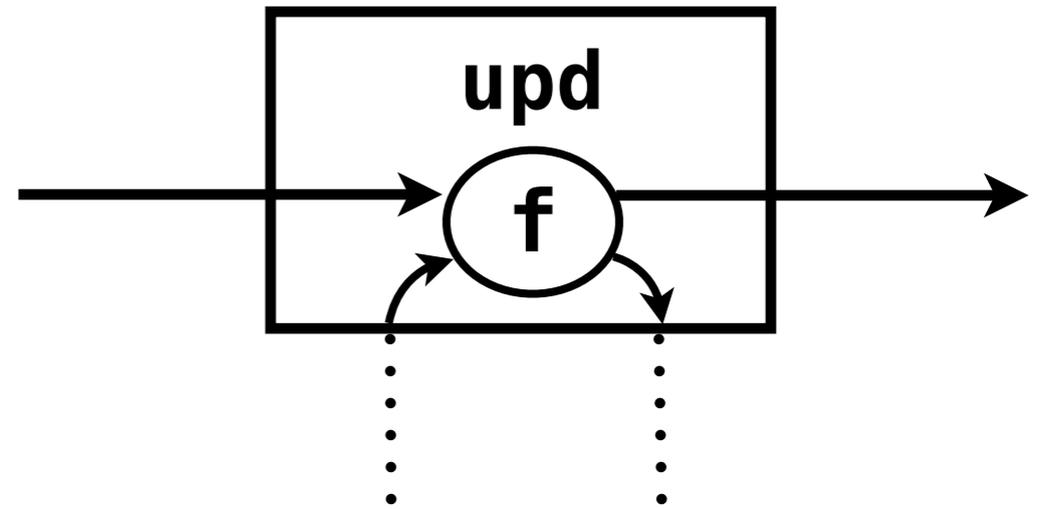
# Shared state



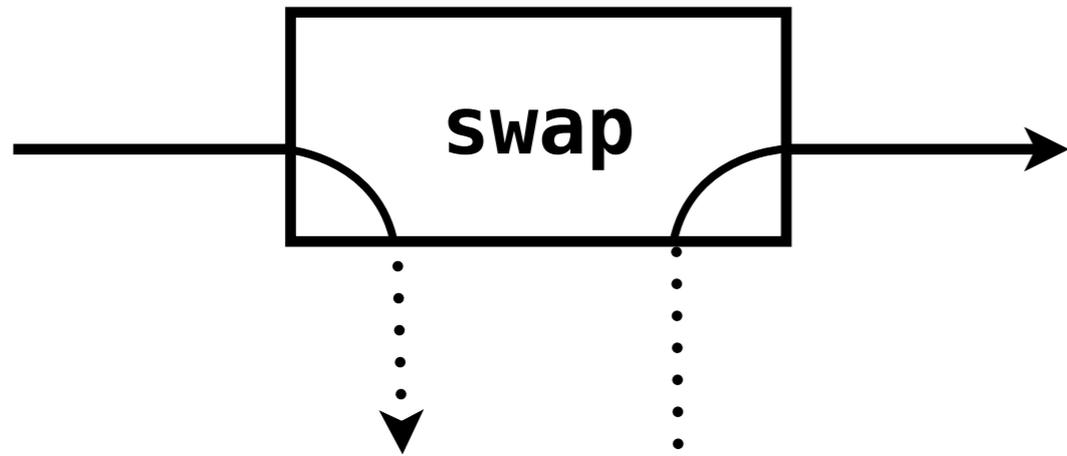
# Message passing



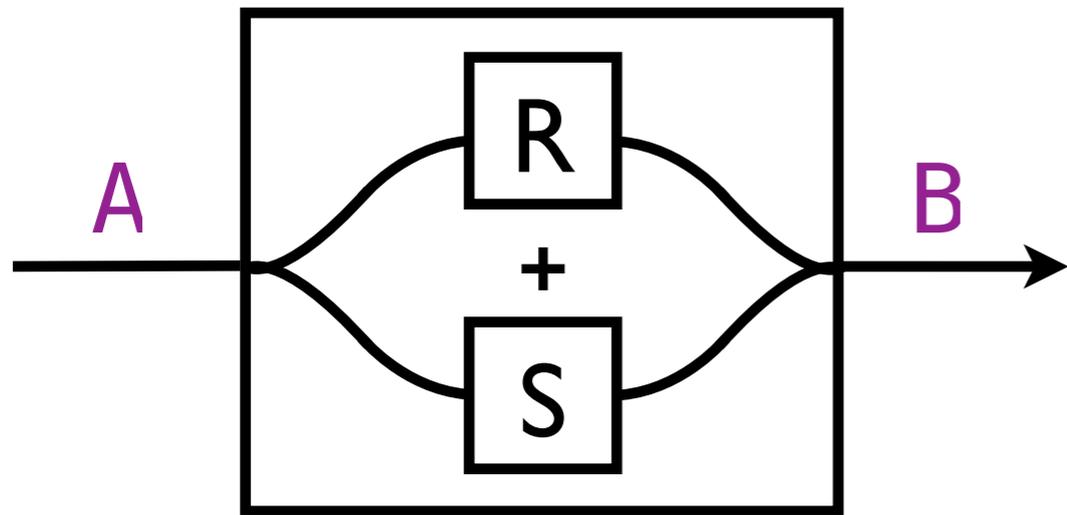
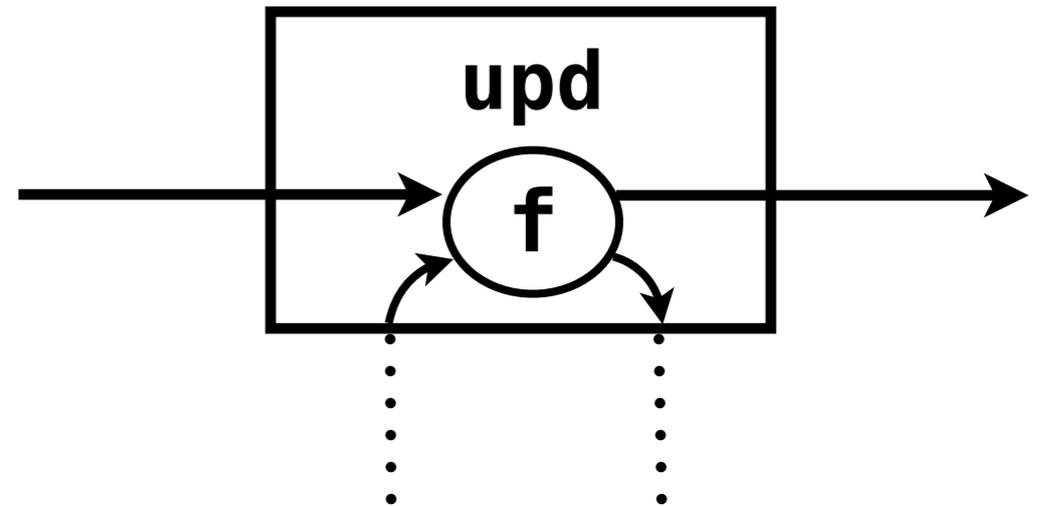
# Shared state



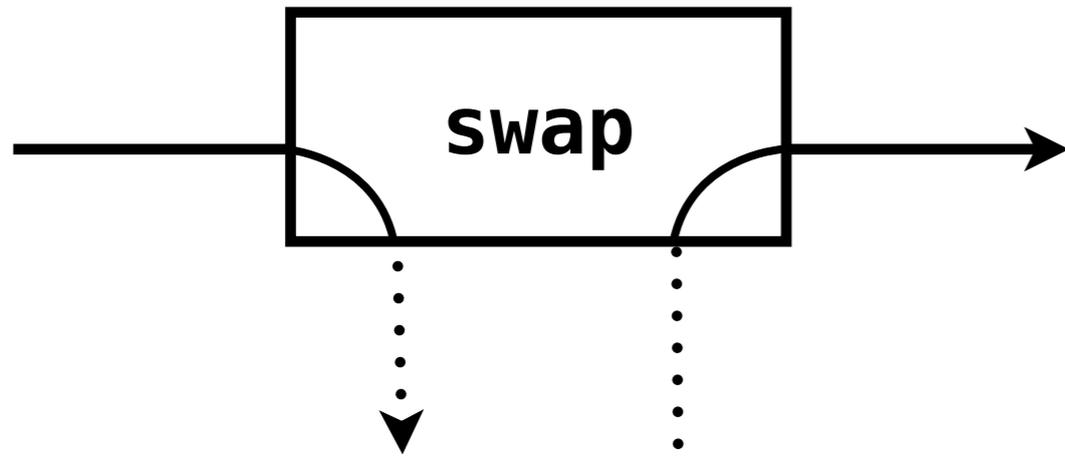
# Message passing



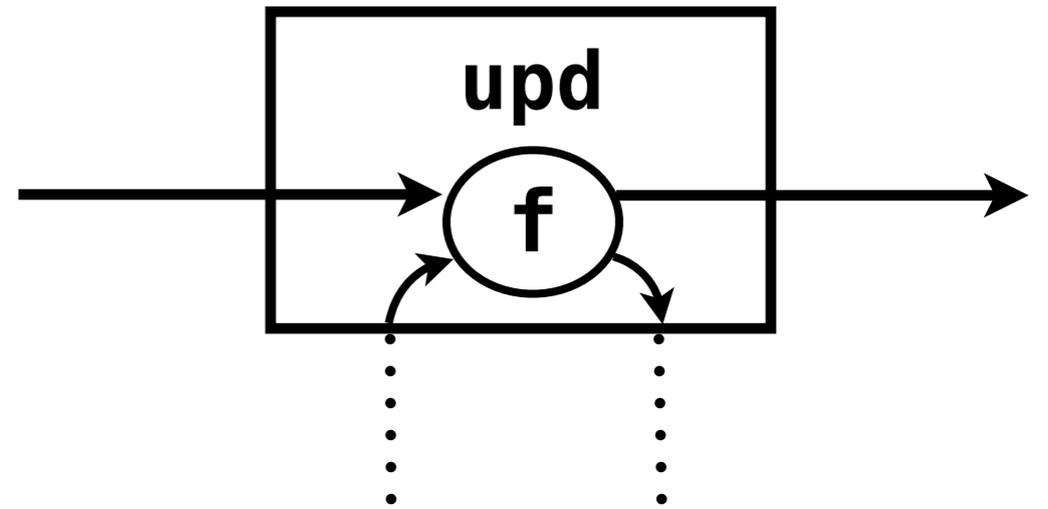
# Shared state



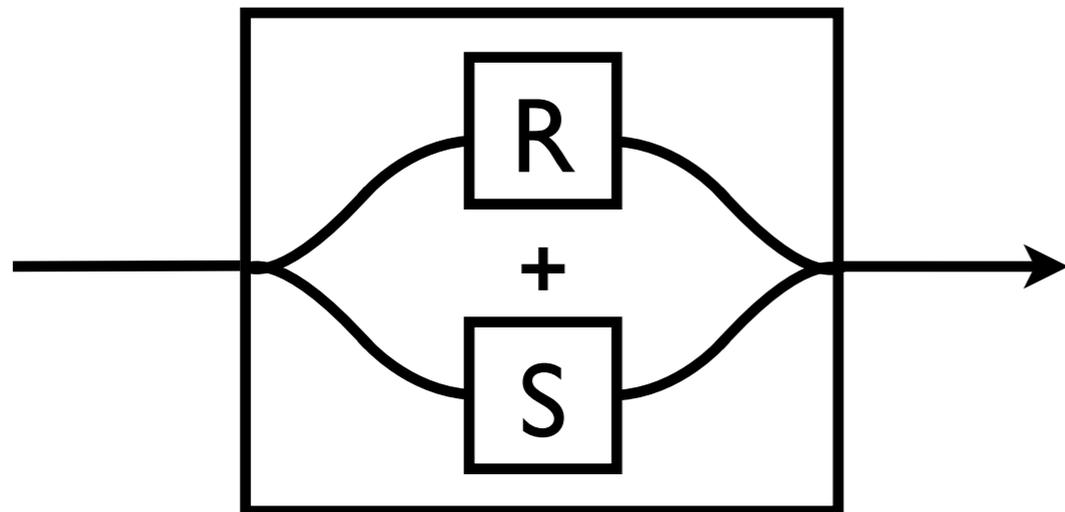
# Message passing



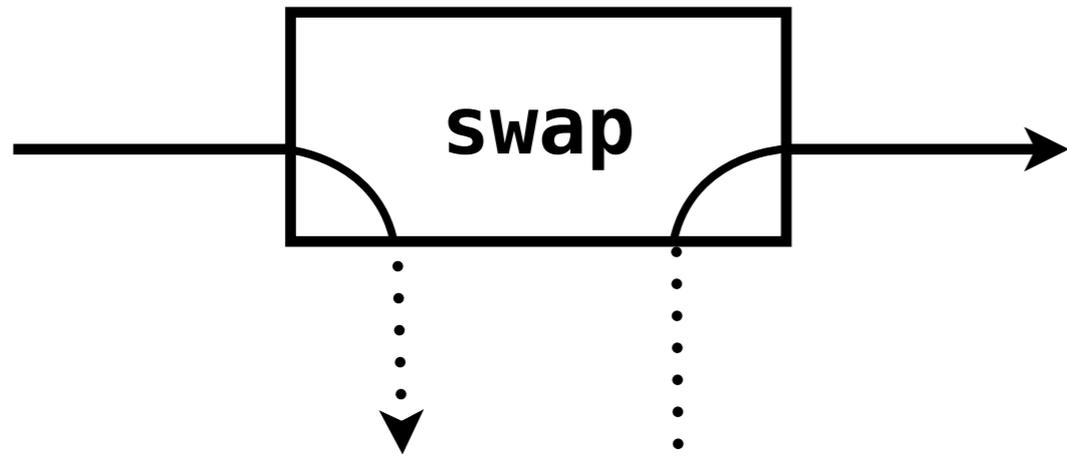
# Shared state



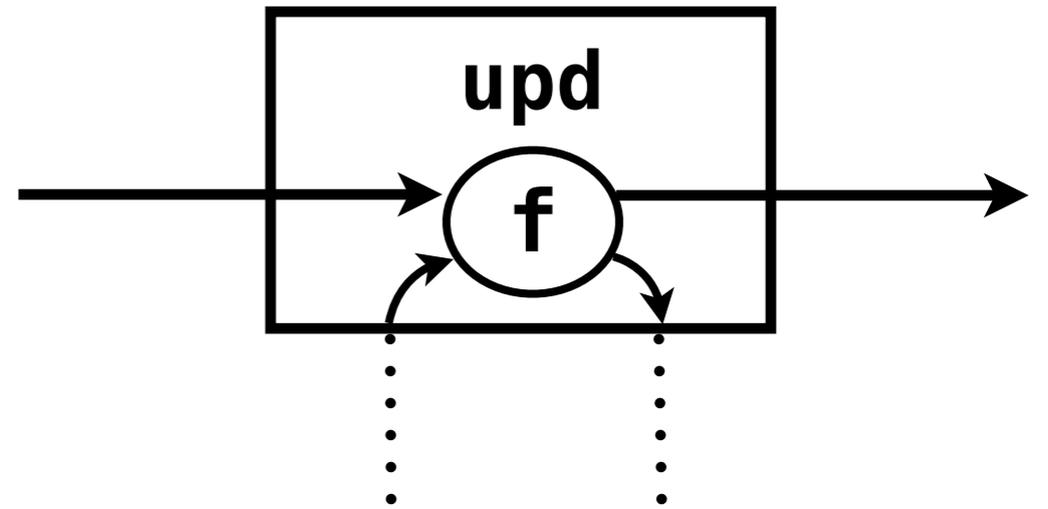
# Disjunction



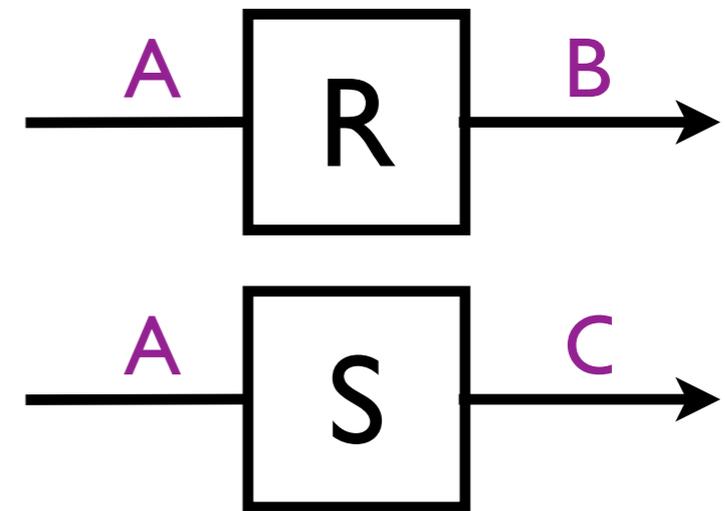
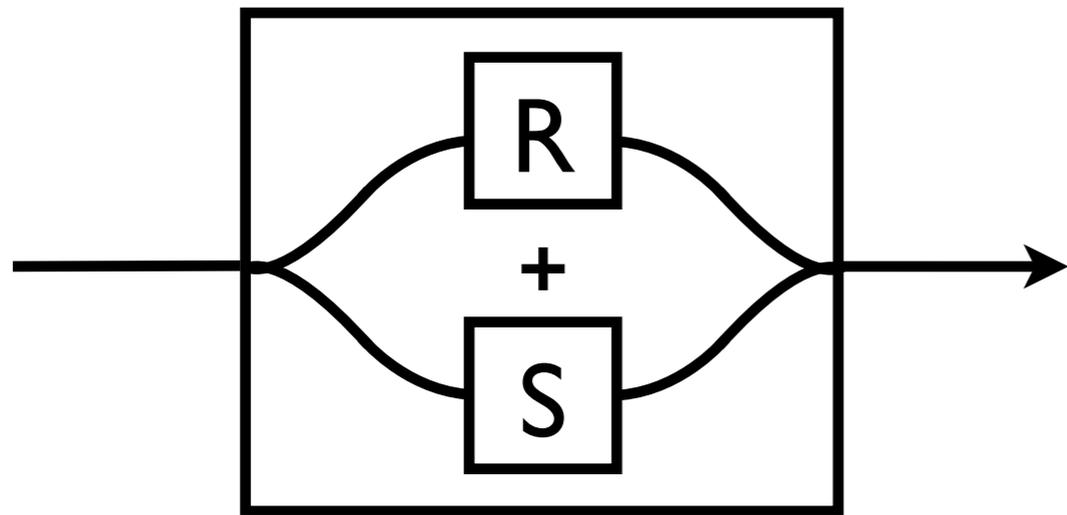
# Message passing



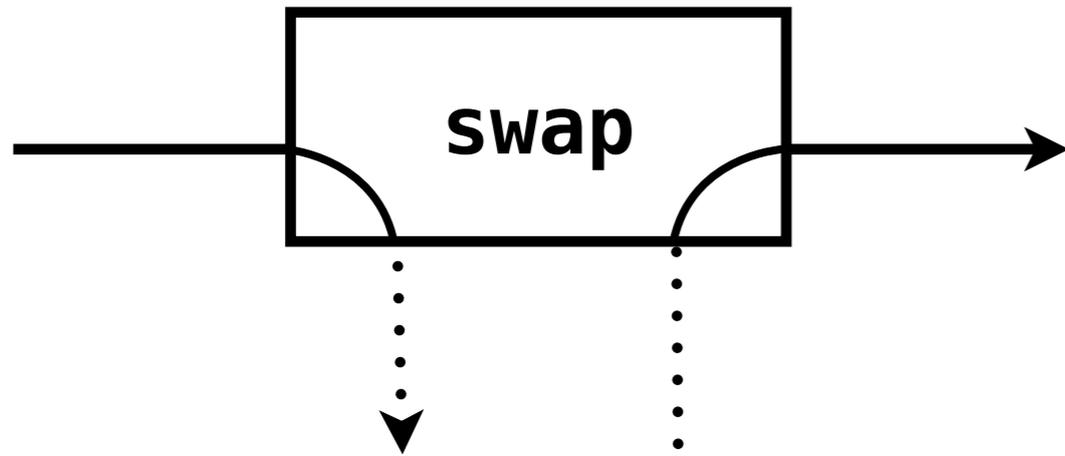
# Shared state



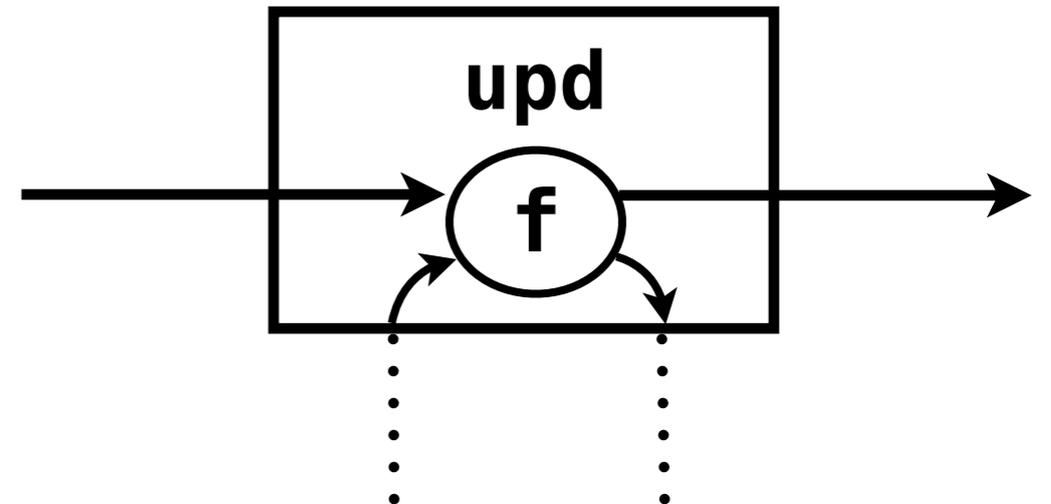
# Disjunction



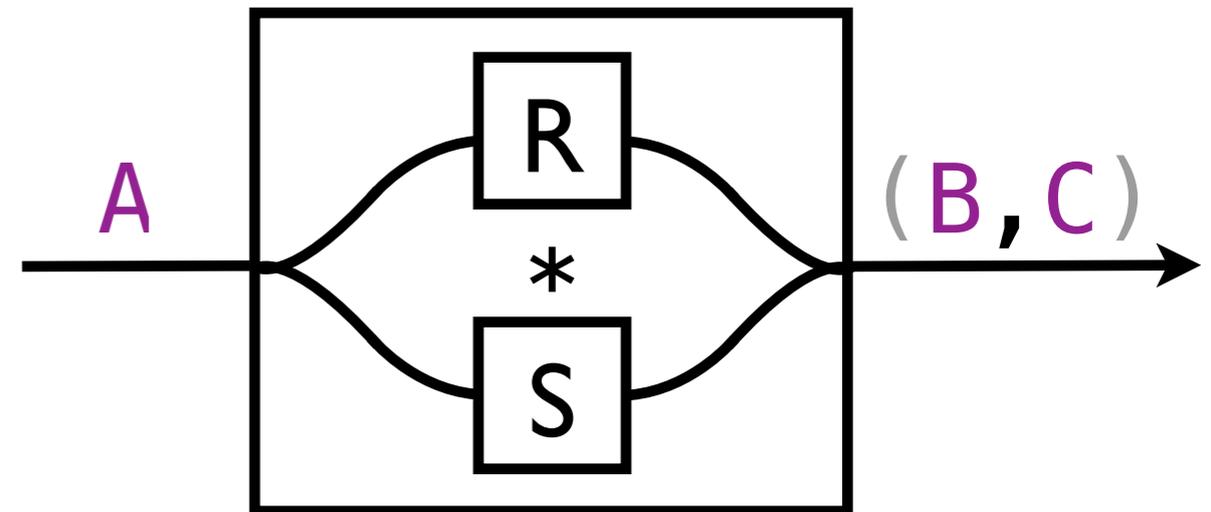
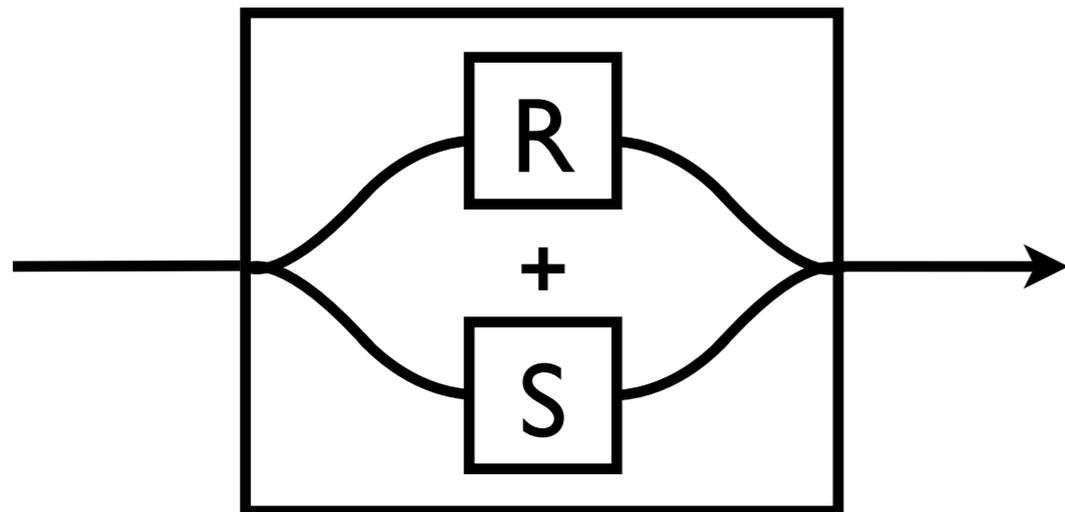
# Message passing



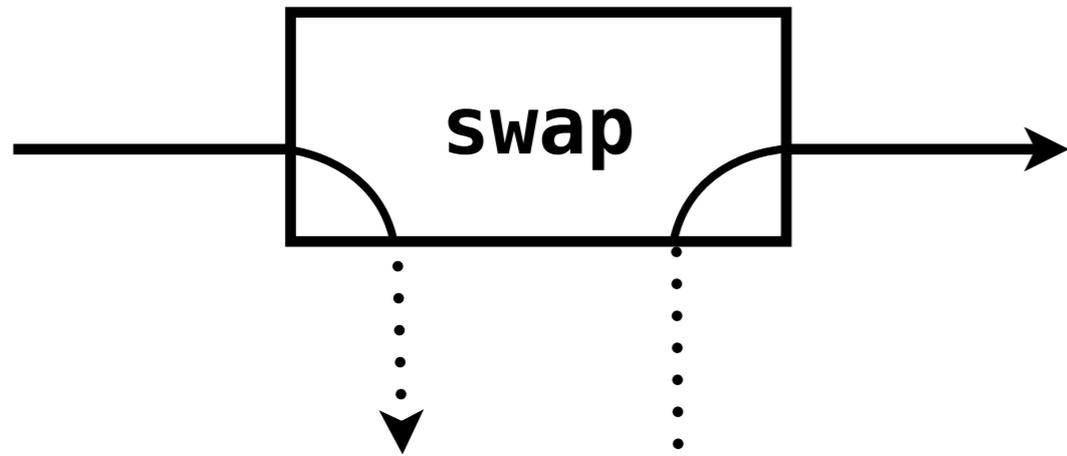
# Shared state



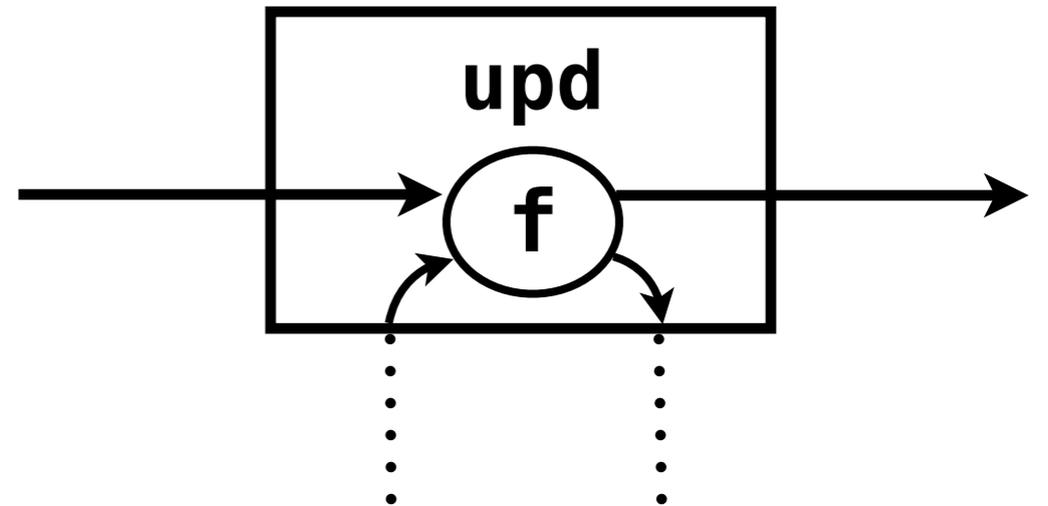
# Disjunction



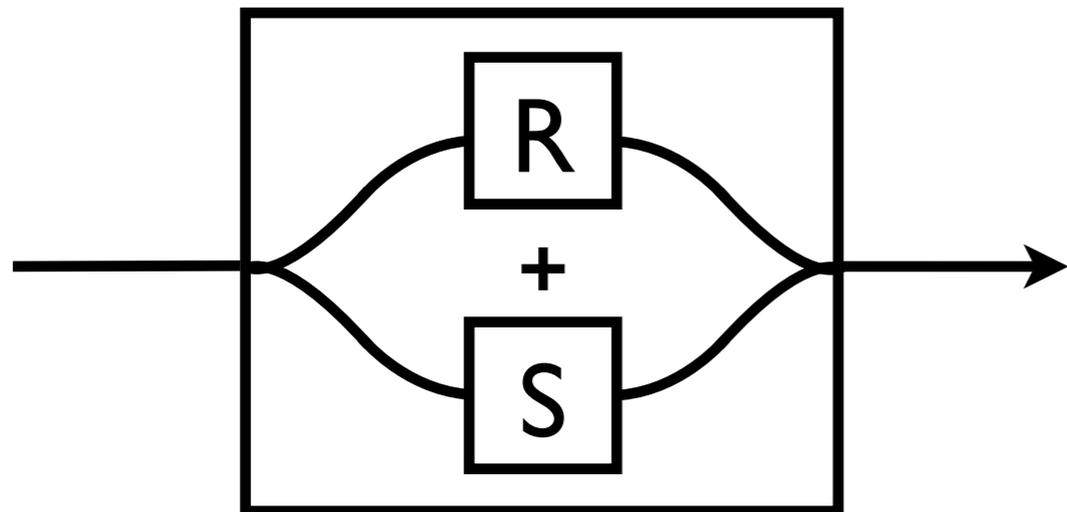
## Message passing



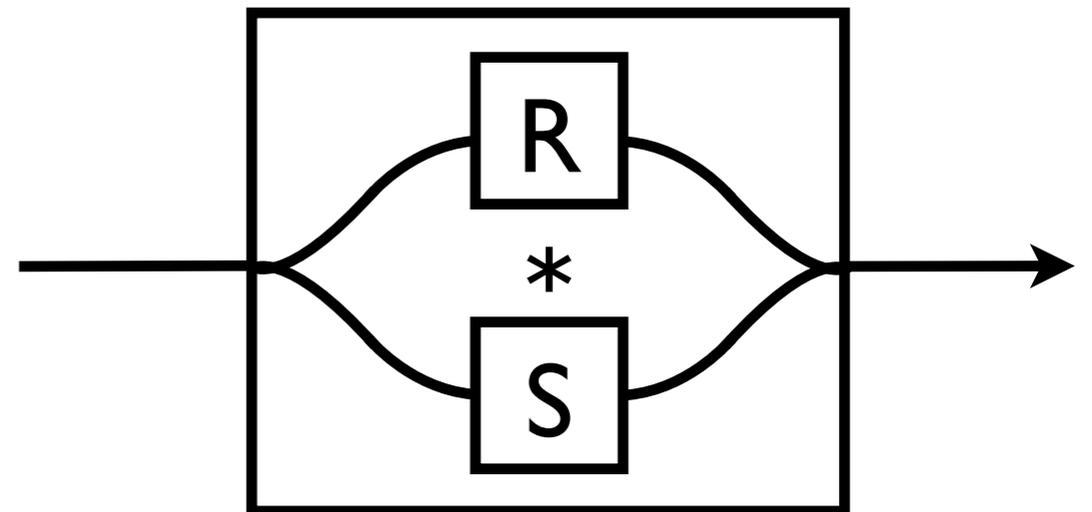
## Shared state



## Disjunction



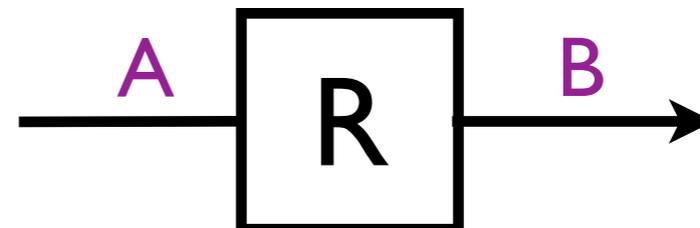
## Conjunction



Lambda abstraction:



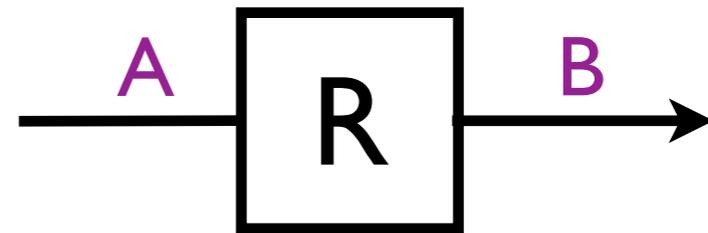
Reagent abstraction:



Lambda abstraction:



Reagent abstraction:



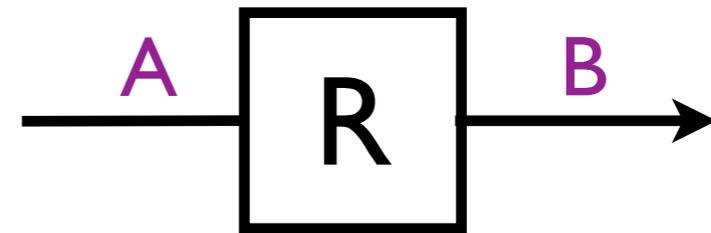
Lambda abstraction:

application:



$$f(a) = b$$

Reagent abstraction:



Lambda abstraction:

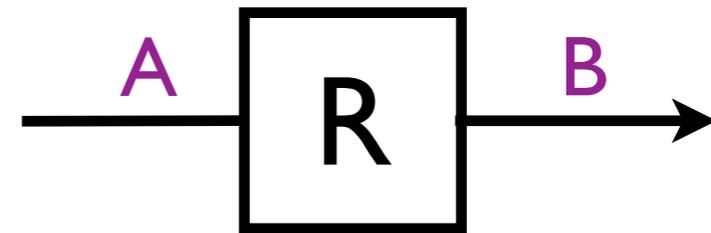
application:



$$f(a) = b$$

Reagent abstraction:

apply as reactant:



$$R ! a = b$$

Lambda abstraction:

application:

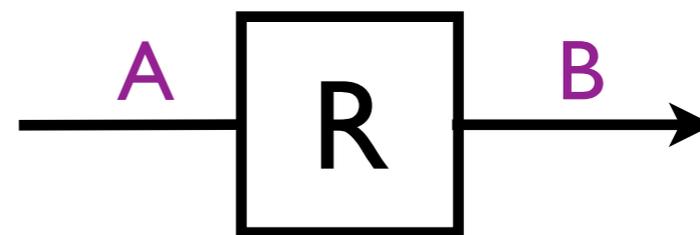


$$f(a) = b$$

Reagent abstraction:

apply as reactant:

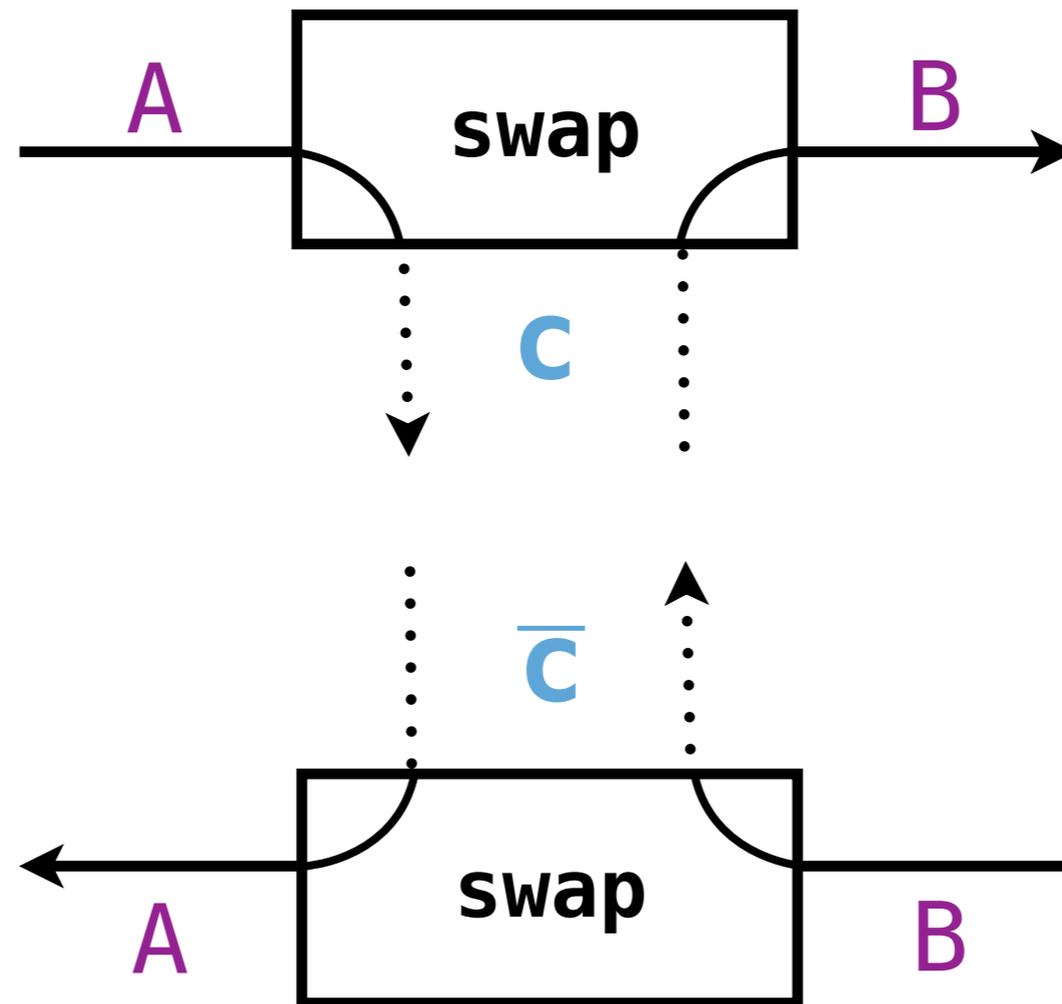
apply as catalyst:



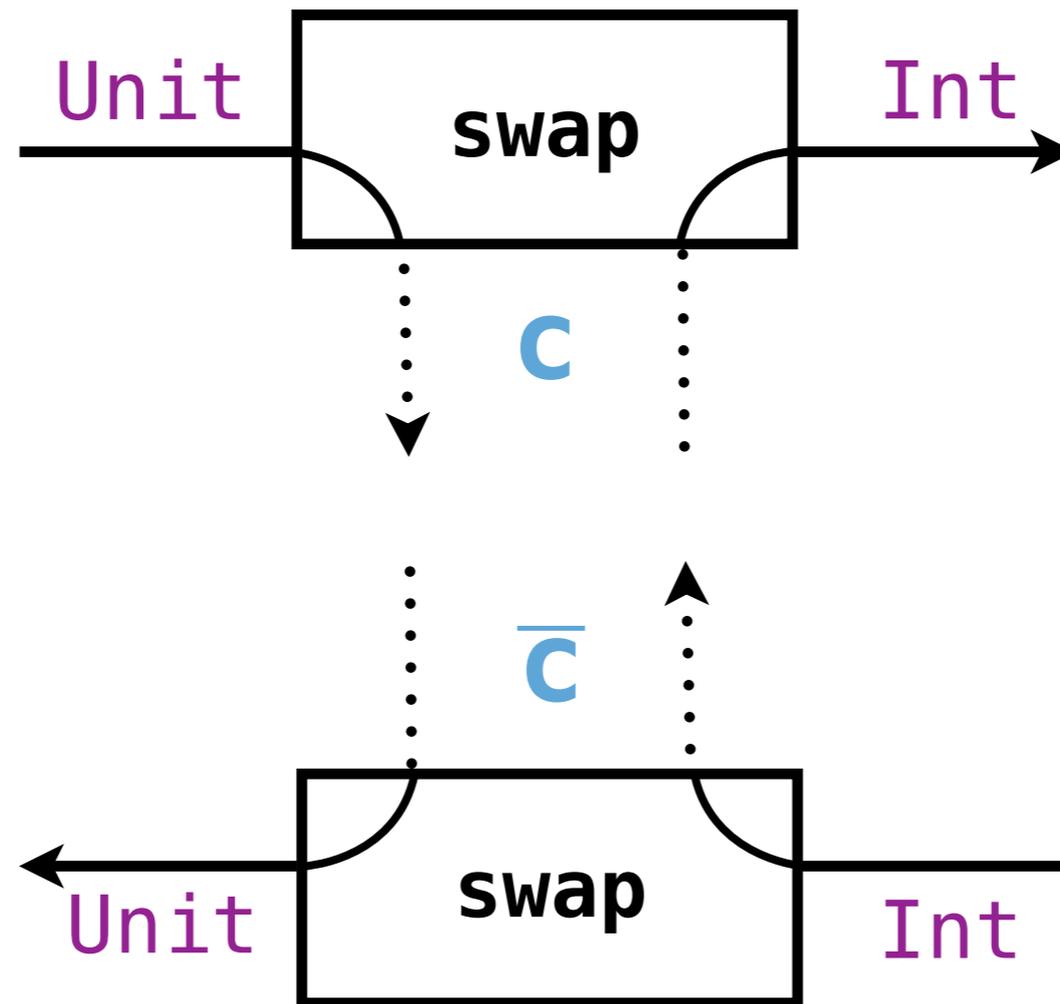
$$R ! a = b$$

dissolve(R)

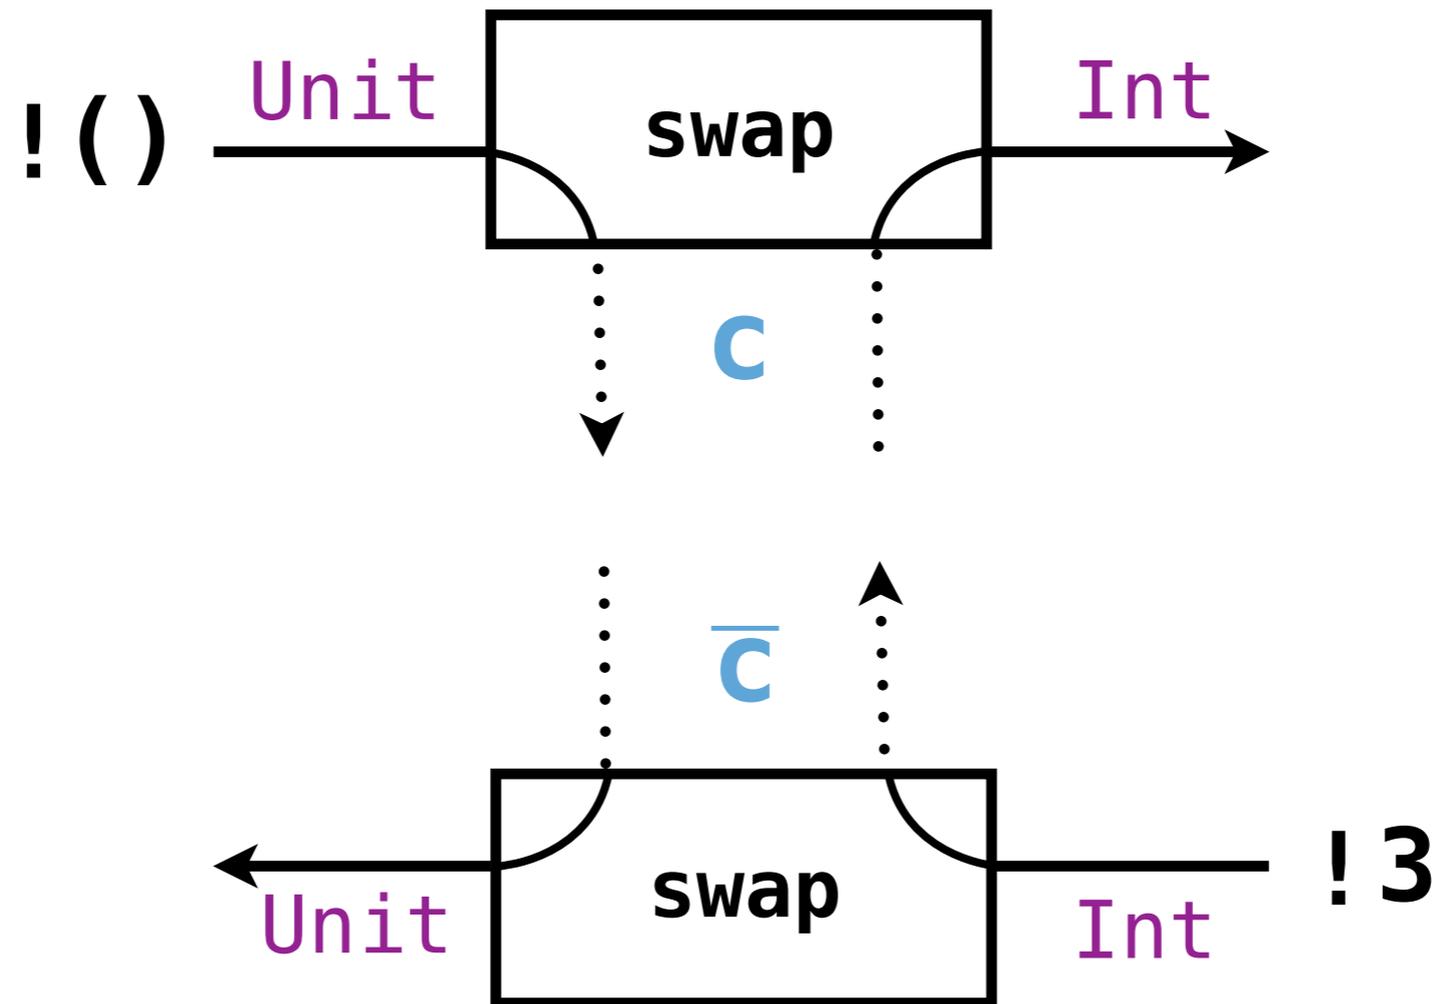
**c**: Chan [A, B]



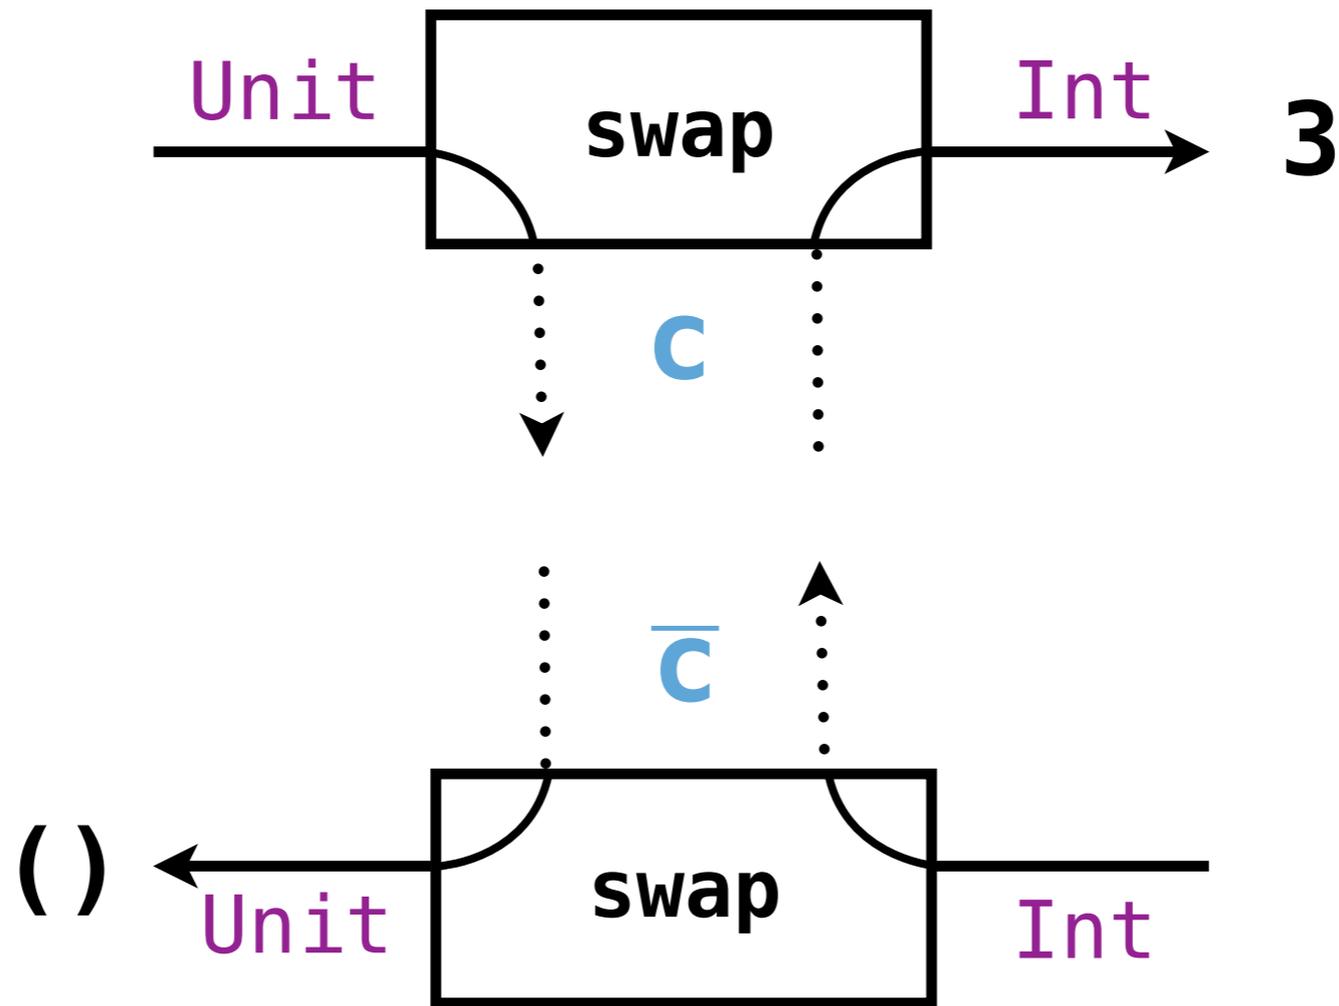
**c**: Chan[Unit, Int]



`c: Chan[Unit, Int]`

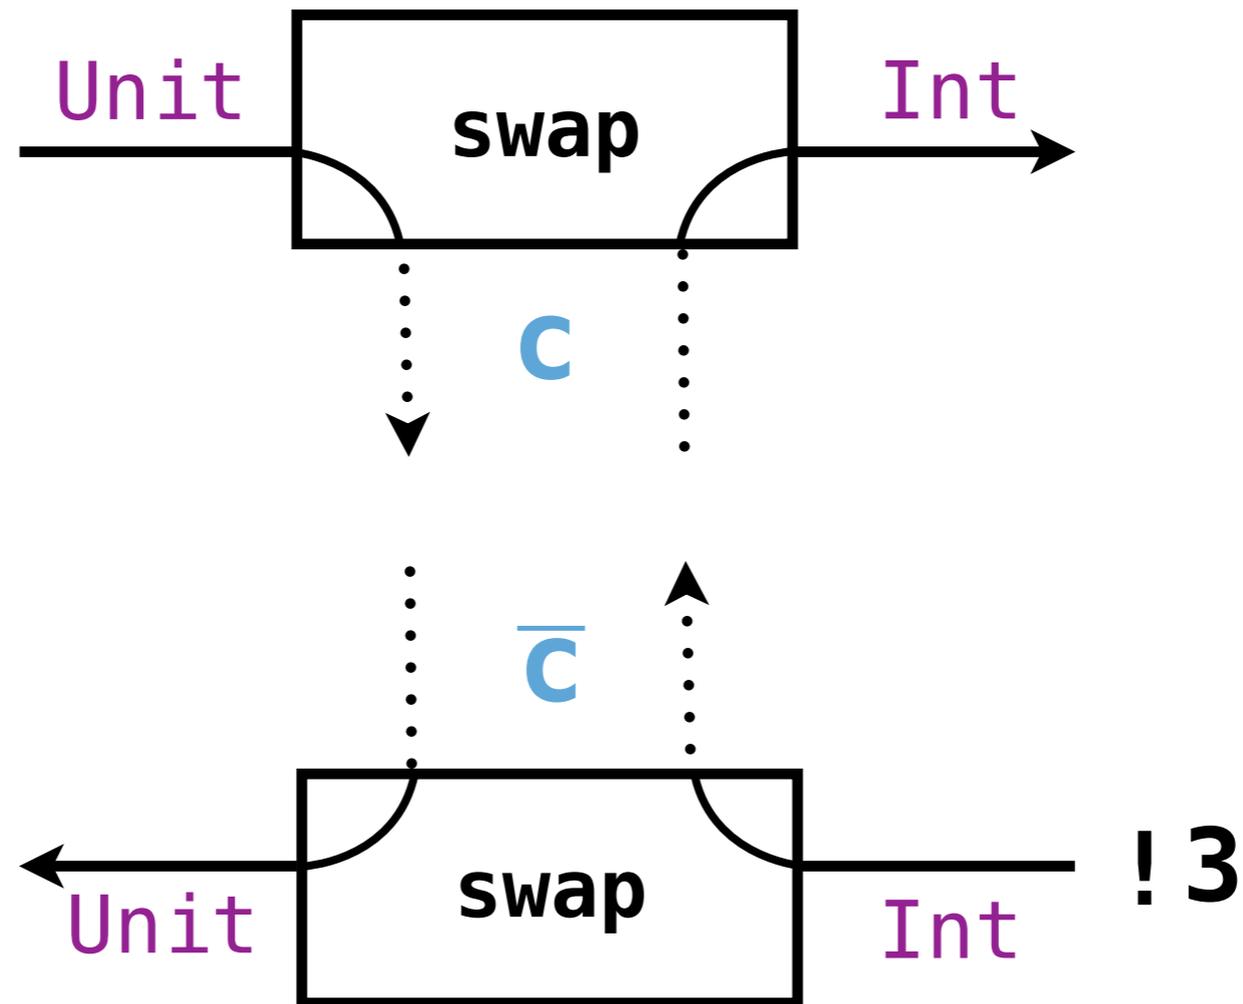


**c**: Chan[Unit, Int]

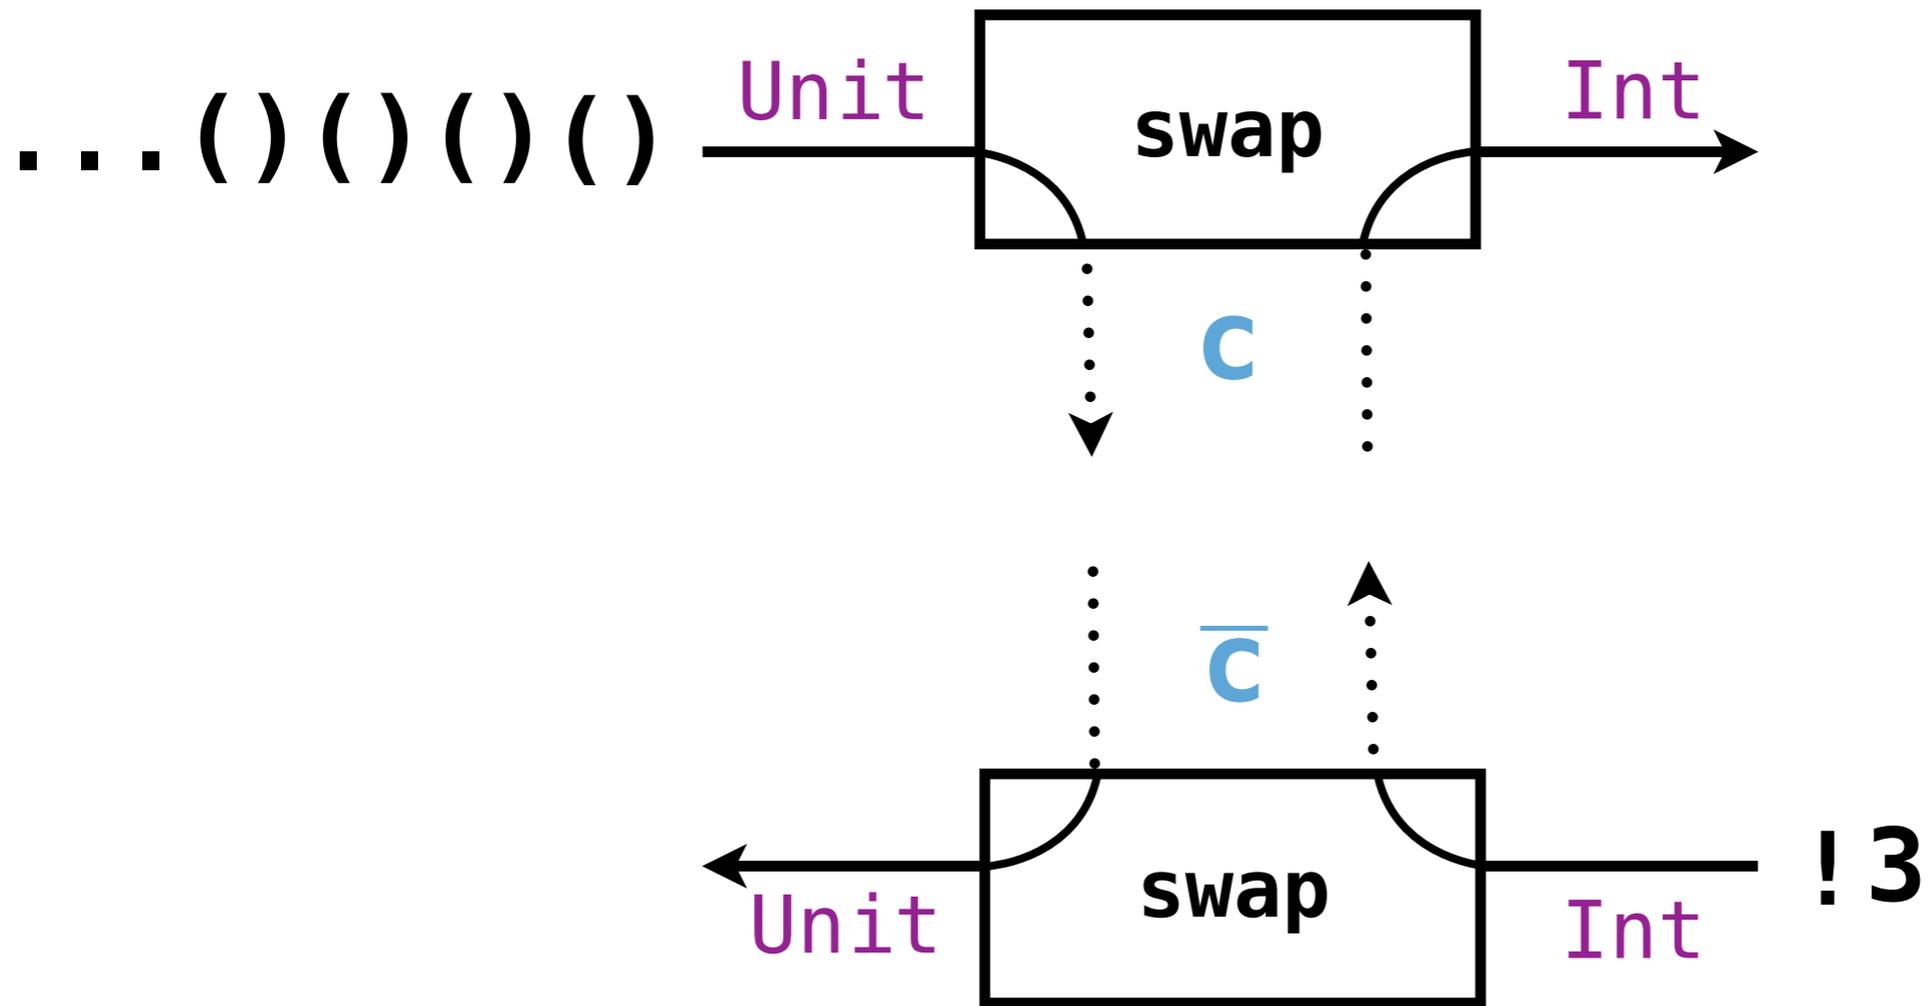


**c**: Chan[Unit, Int]

**dissolve**

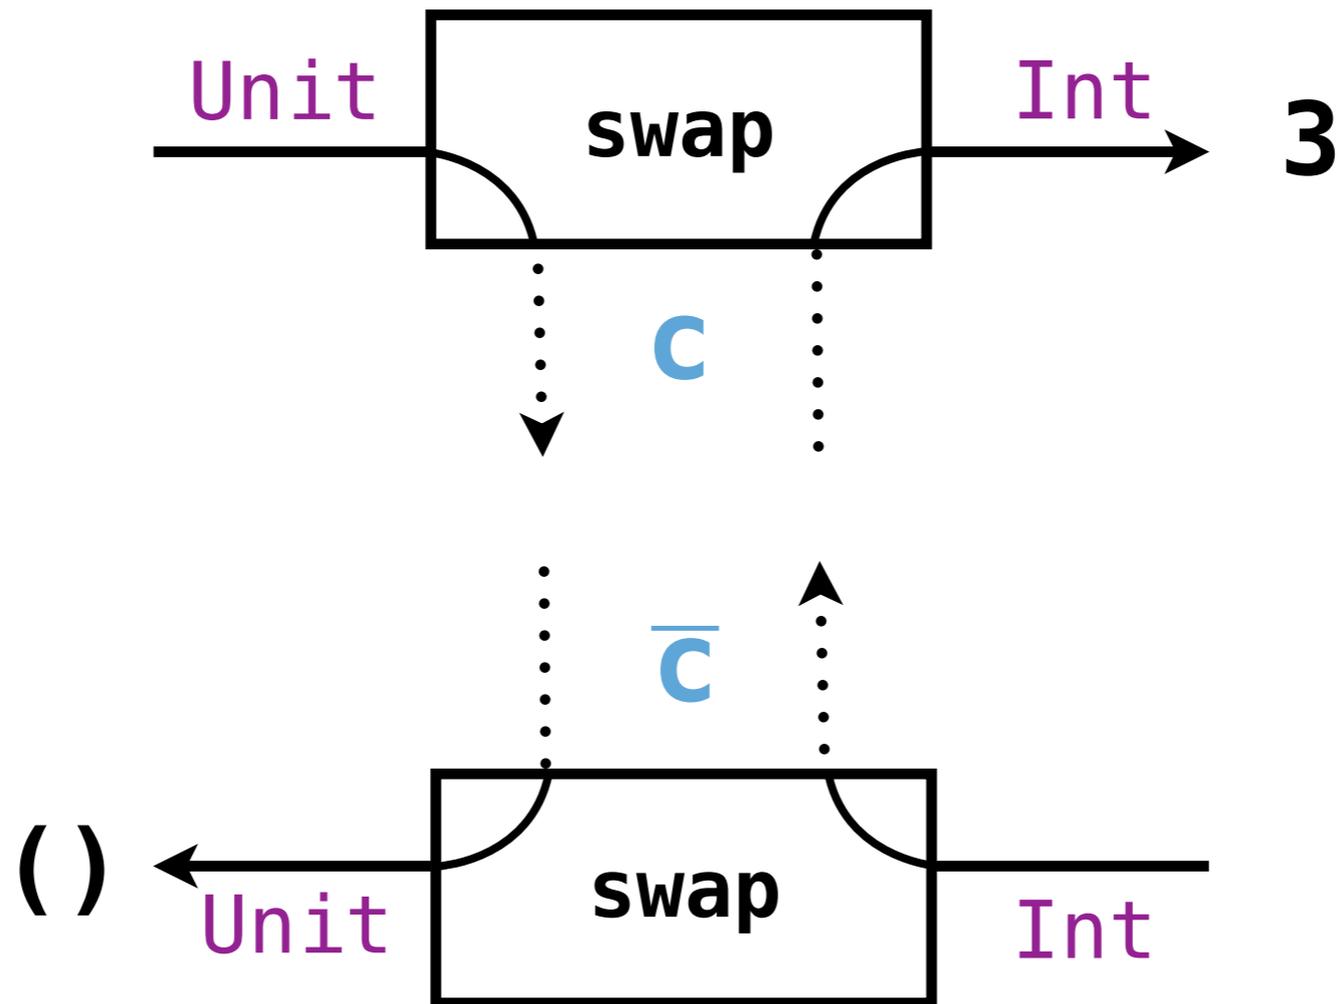


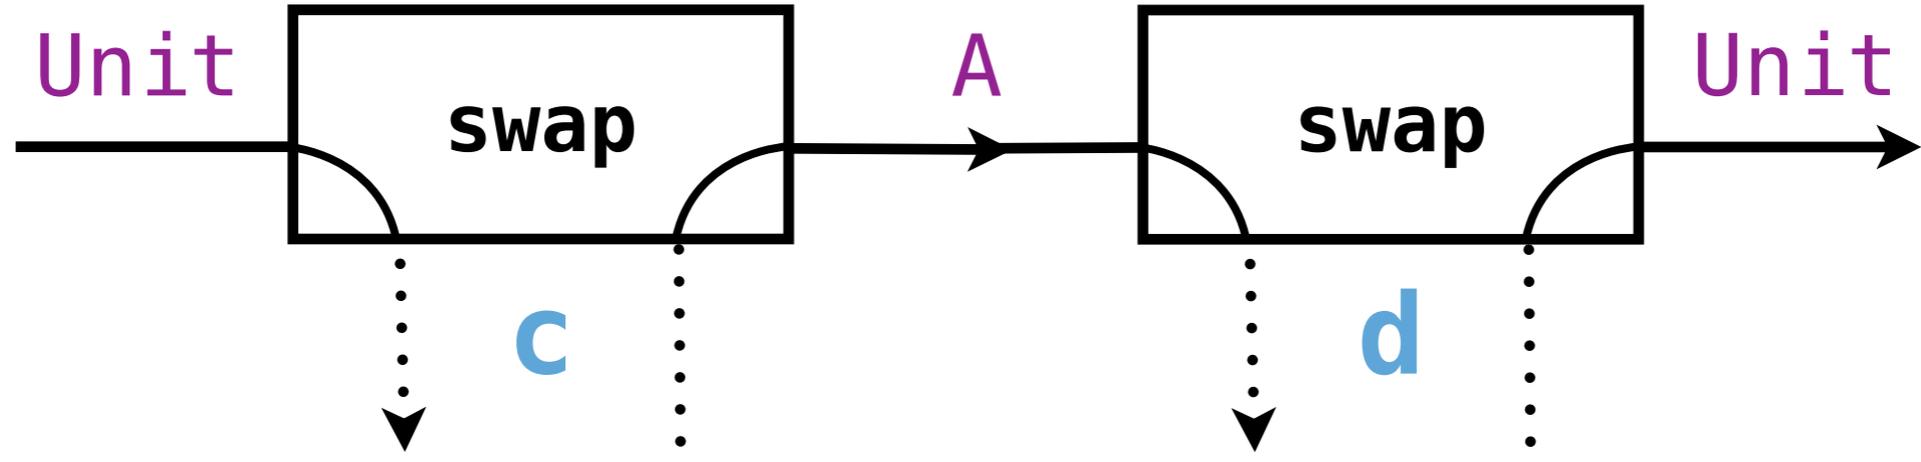
`c: Chan[Unit, Int]`

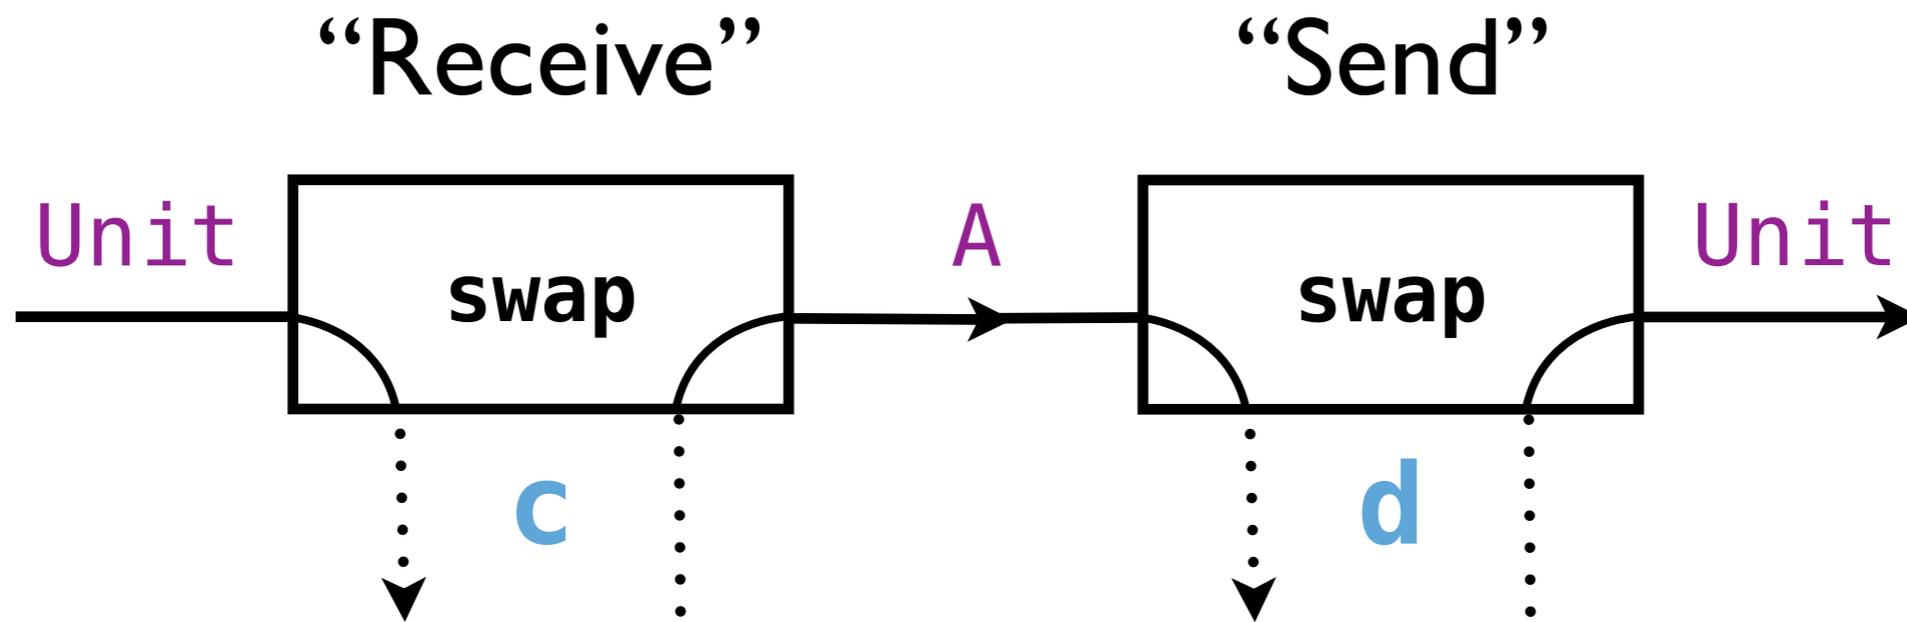


**c**: Chan[Unit, Int]

... () () ()

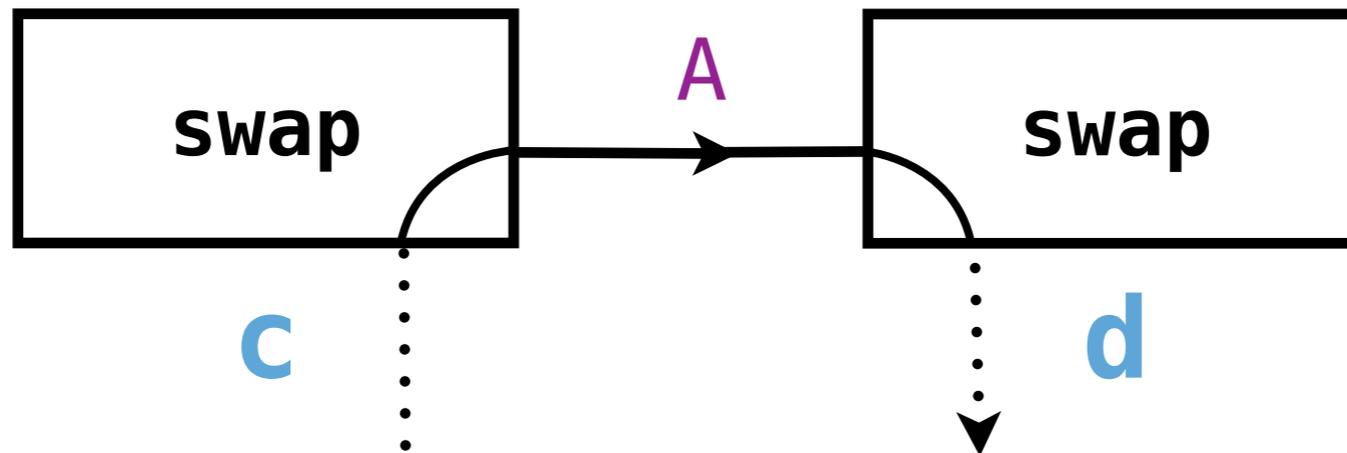




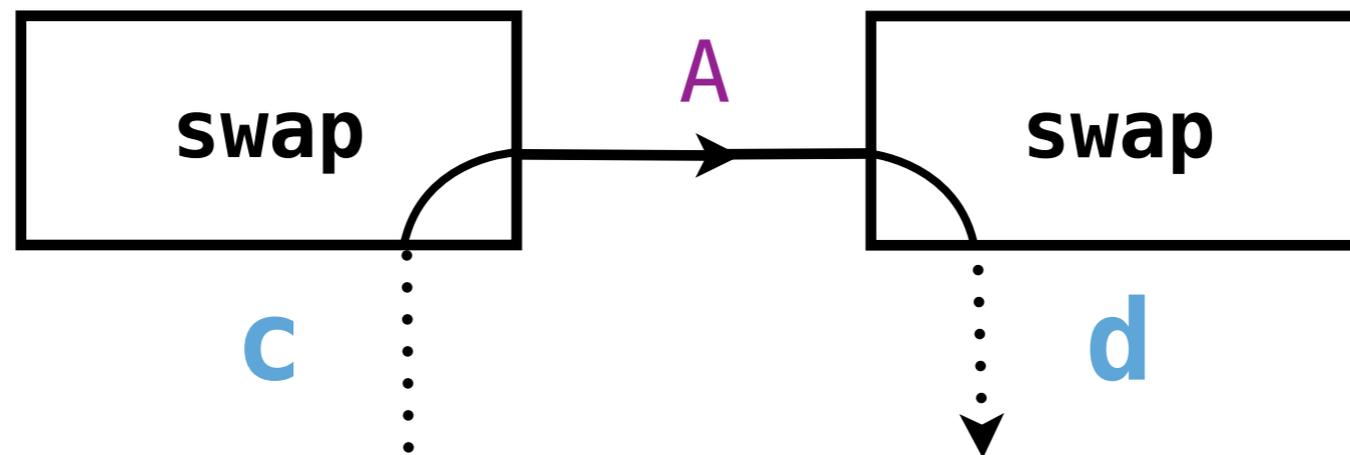


“Receive”

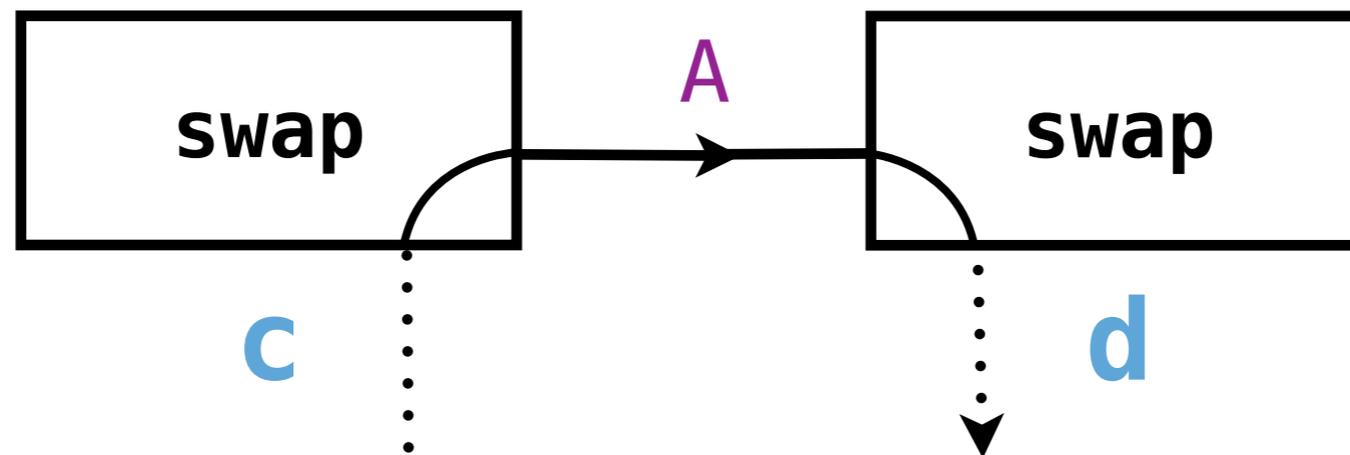
“Send”



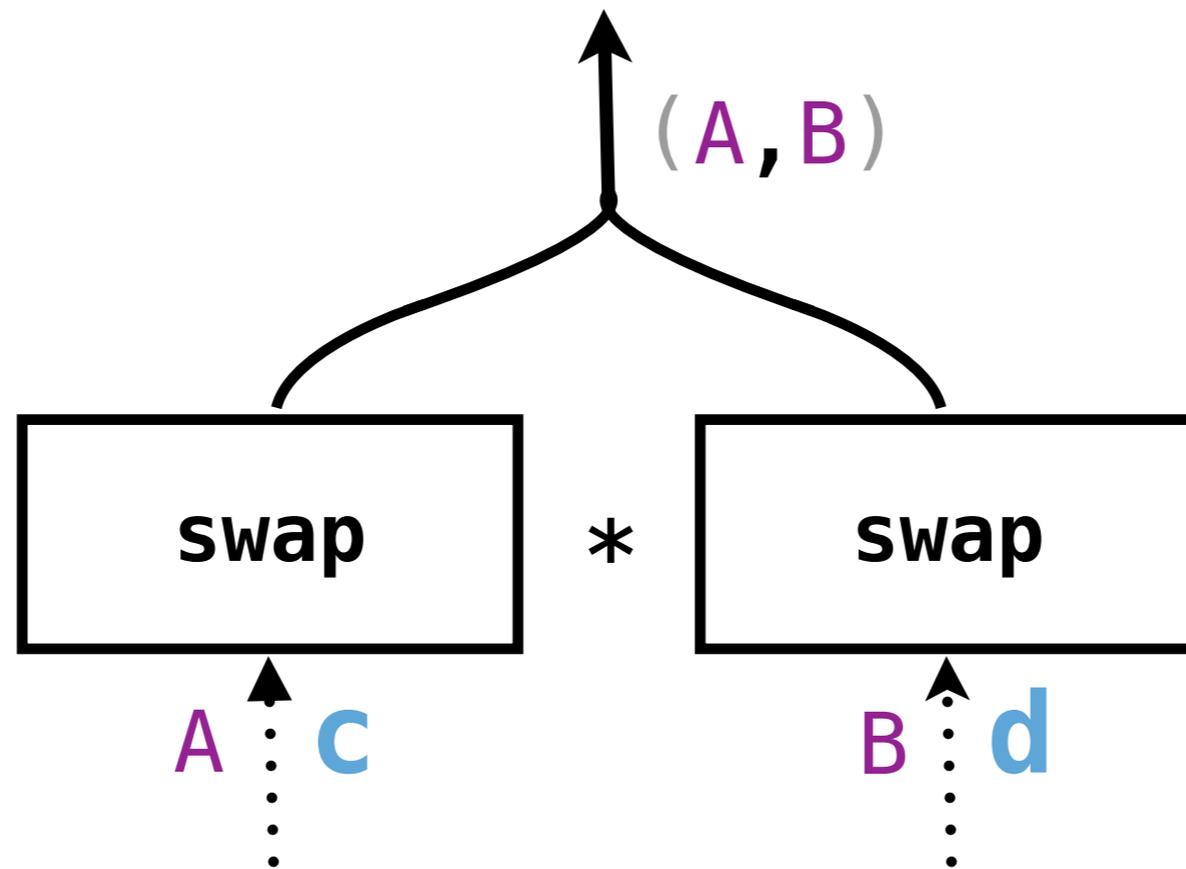
# Pipeline catalyst



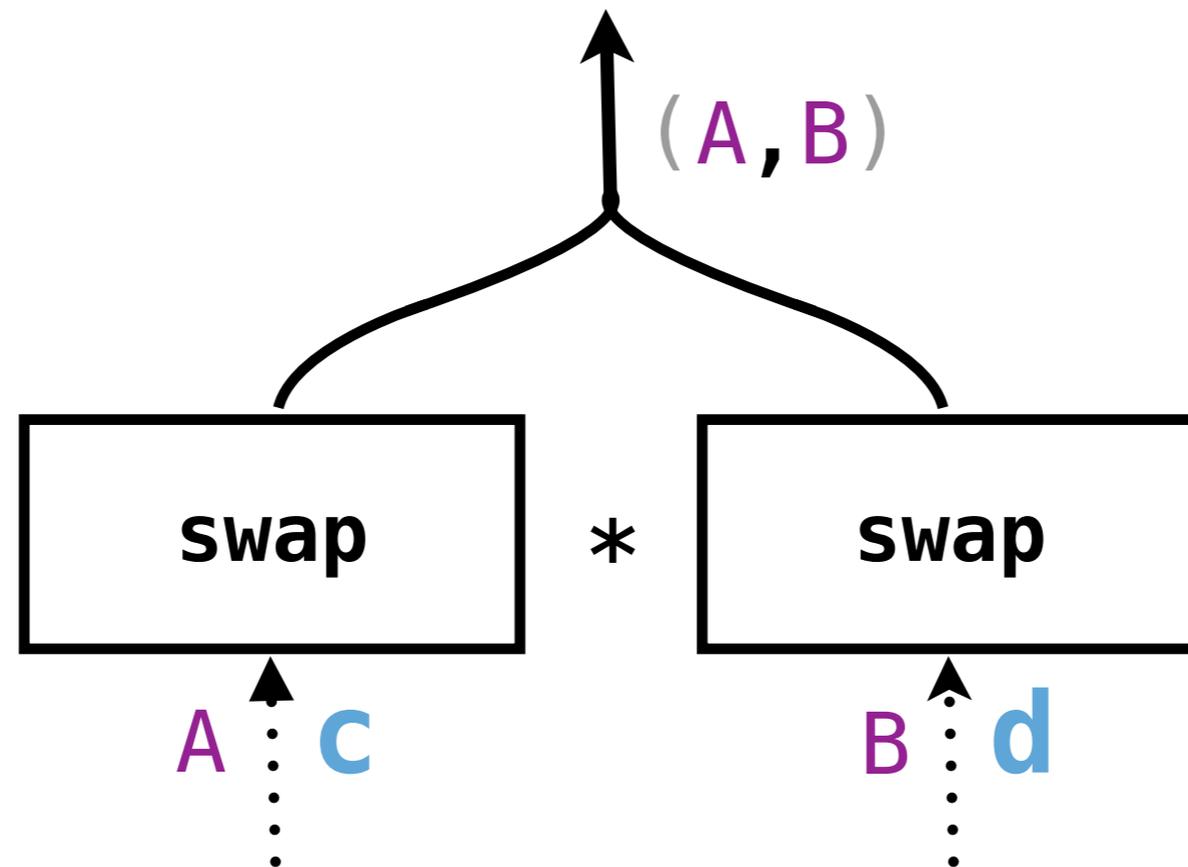
# Pipeline catalyst



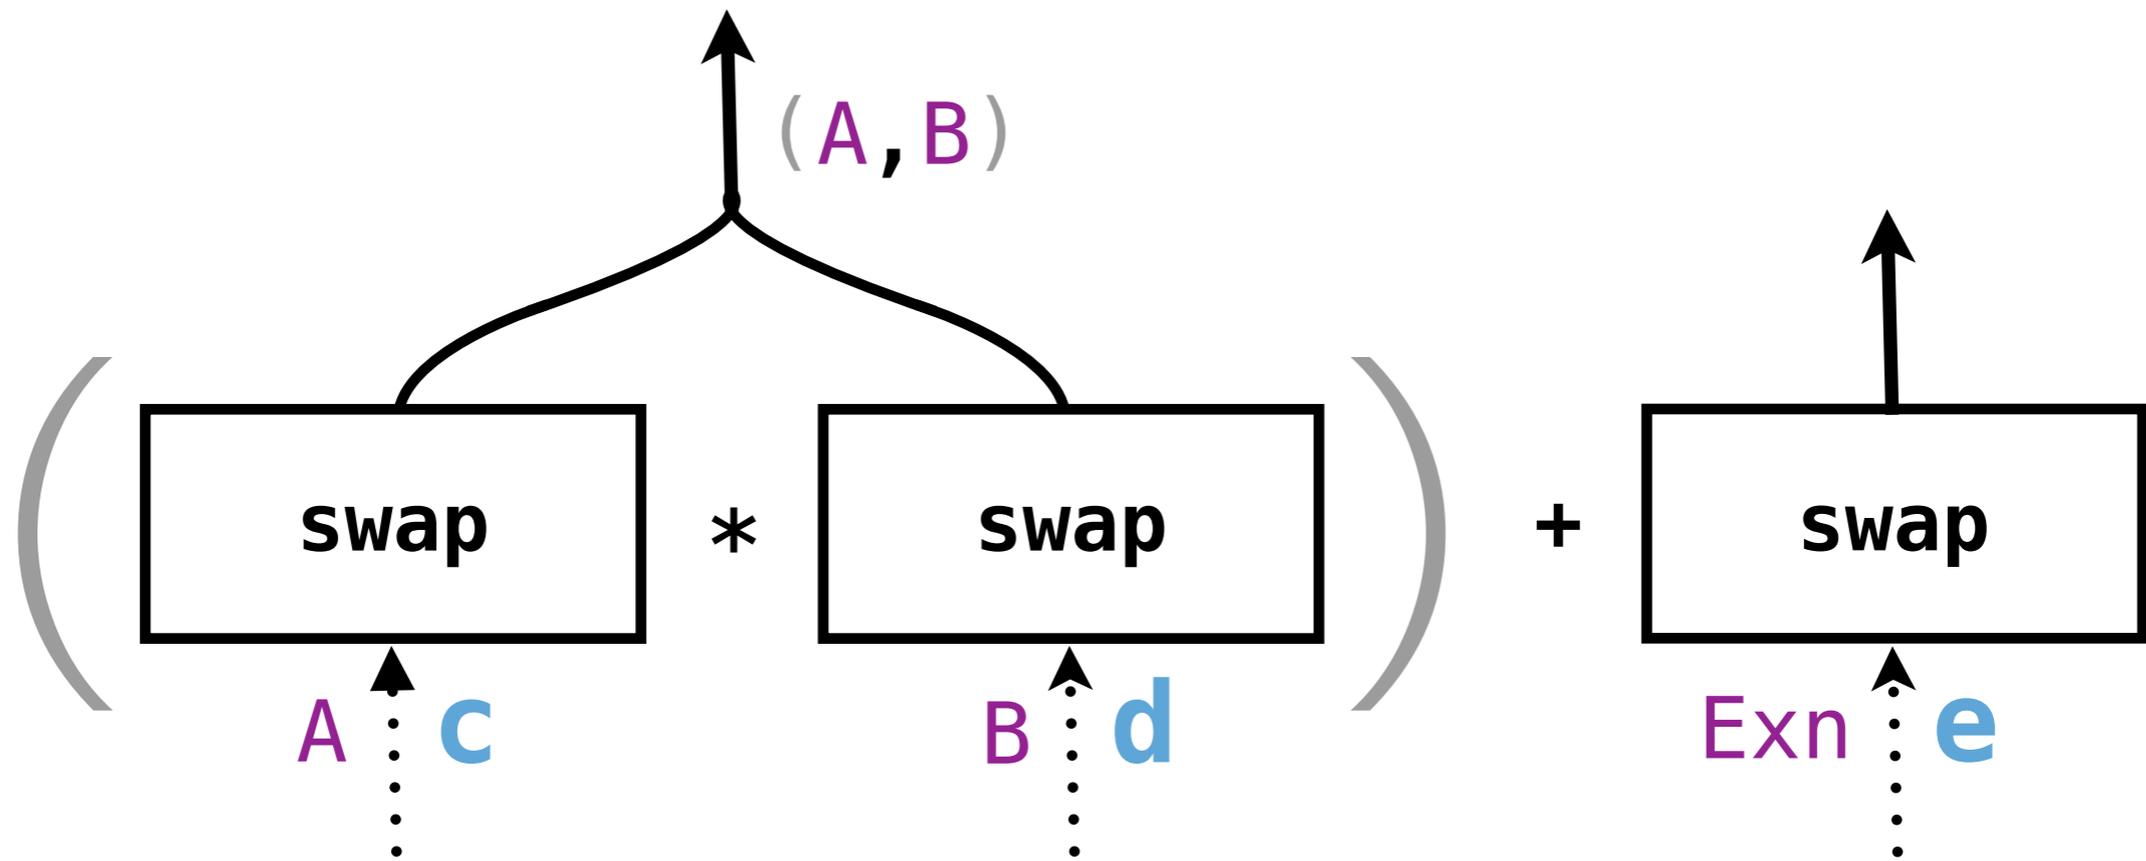
NB: transfer is atomic



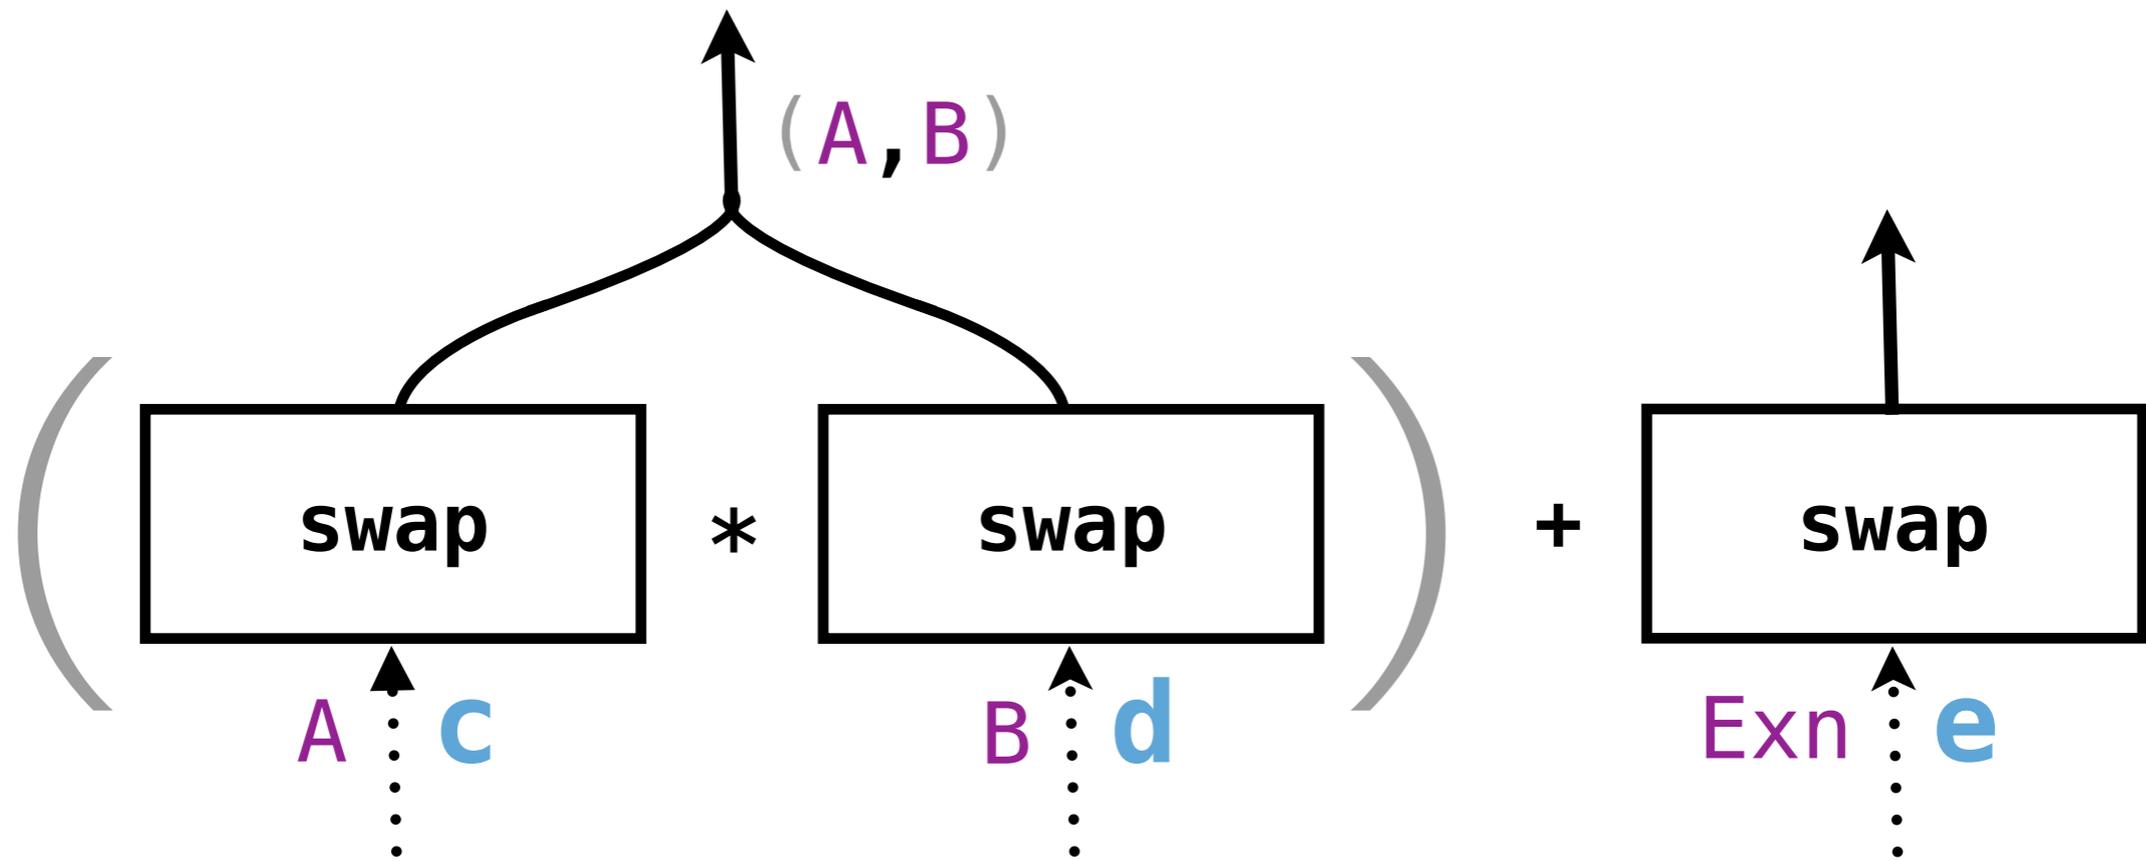
# 2-way join



# 2-way join



# Abortable 2-way join



# Join *Calculus*

$$c_1(x_1) \ \& \ \cdot \cdot \cdot \ \& \ c_n(x_n) \ \Rightarrow \ e$$

# Join *Calculus*

$c_1(x_1) \& \dots \& c_n(x_n) \Rightarrow e$

*becomes*

**(swap  $c_1$  \*  $\dots$  \* swap  $c_n$ )**

**>>> postCommit e**

# Join *Calculus*

$c_1(x_1) \& \dots \& c_n(x_n) \Rightarrow e$

*becomes*

**dissolve**  $\left( \begin{array}{l} \text{swap } c_1 * \dots * \text{swap } c_n \\ \gg \gg \text{postCommit } e \end{array} \right)$

```

class TreiberStack [A] {
  private val head = new Ref[List[A]](Nil)
  val push : A => () = upd(head)(cons)
  val tryPop : () => A? = upd(head) {
    case (x :: xs) => (xs, Some(x))
    case Nil => (Nil, None)
  }
}

```

```

class TreiberStack [A] {
  private val head = new Ref[List[A]](Nil)
  val push : A => () = upd(head)(cons)
  val tryPop : () => A? = upd(head) {
    case (x :: xs) => (xs, Some(x))
    case Nil => (Nil, None)
  }
  val pop : () => A = upd(head) {
    case (x :: xs) => (xs, x)
  }
}

```

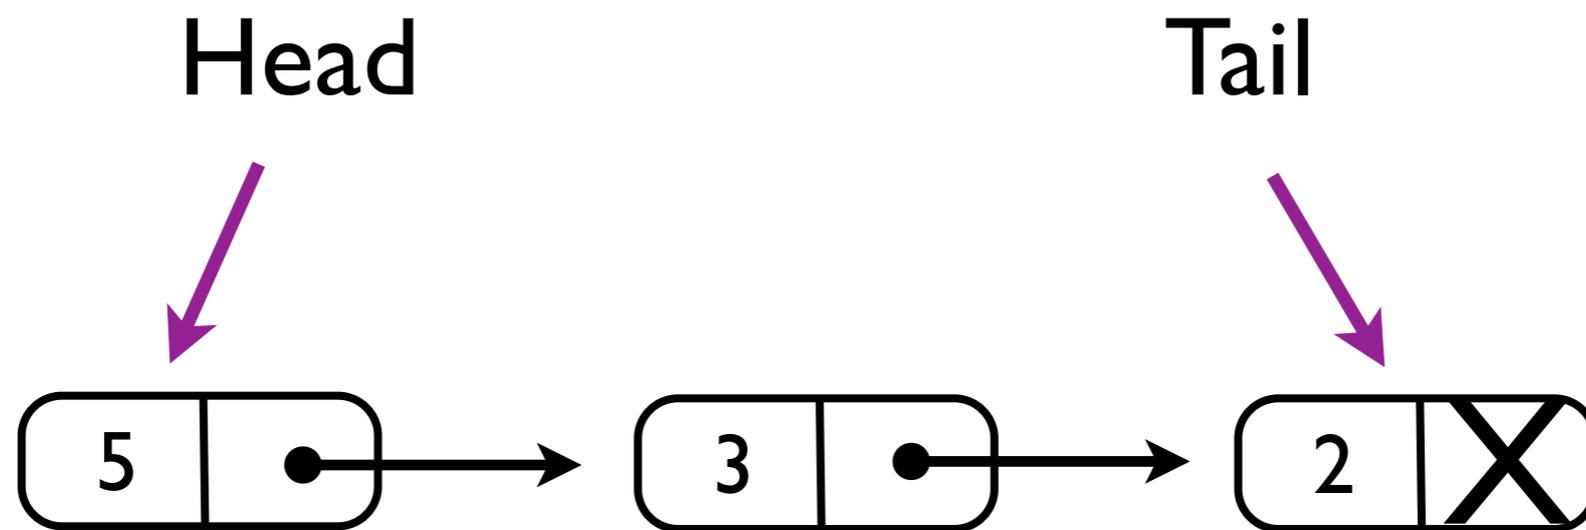
```
class TreiberStack [A] {  
  private val head = new Ref[List[A]](Nil)  
  val push      = upd(head)(cons)  
  val tryPop    = upd(head)(trySplit)  
  val pop       = upd(head)(split)  
}
```

```
class TreiberStack [A] {  
  private val head = new Ref[List[A]](Nil)  
  val push      = upd(head)(cons)  
  val tryPop    = upd(head)(trySplit)  
  val pop       = upd(head)(split)  
}
```

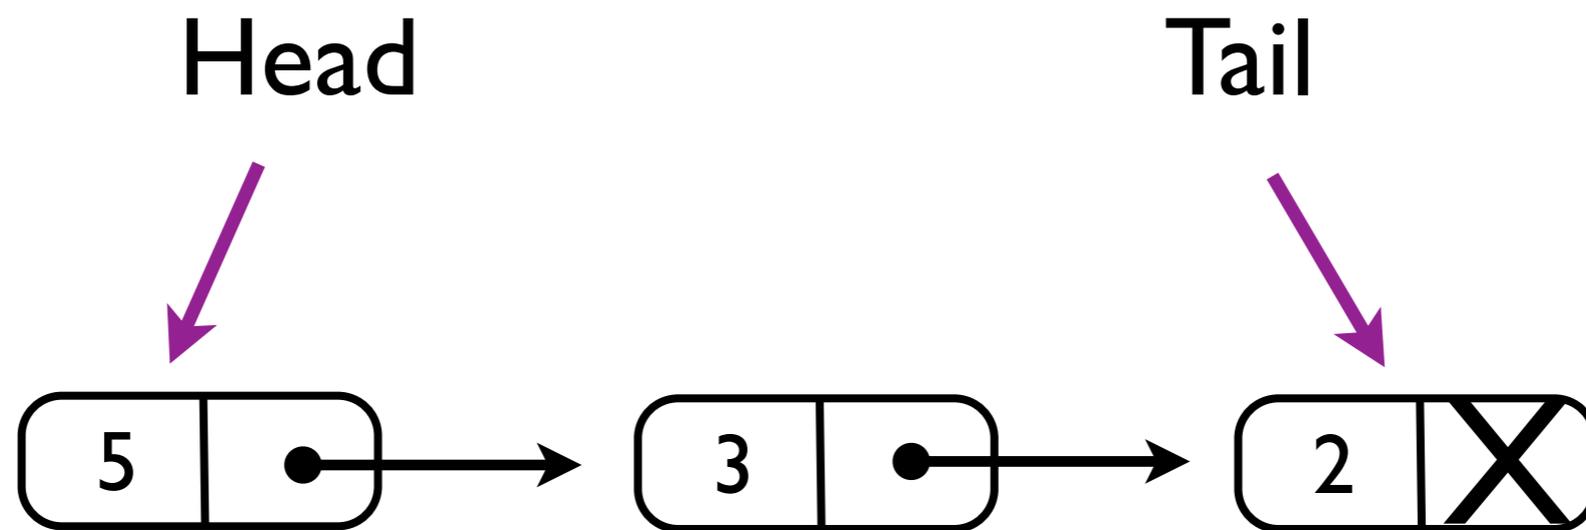
```
class EliminationStack [A] {  
  private val stack = new TreiberStack[A]  
  private val (send, recv) = new Chan[A]  
  val push = stack.push + swap(send)  
  val pop  = stack.pop  + swap(recv)  
}
```

**stack1.pop >>> stack2.push**

# Going Monadic



# Going Monadic



computed:  $A \rightarrow (( ) \Rightarrow B) \rightarrow (A \Rightarrow B)$



Use *invisible side-effects* to traverse the queue while computing the `upd` operation to perform

# Implementation

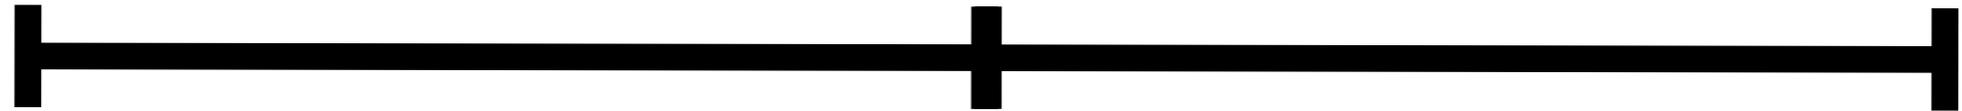
Phase 1

Phase 2

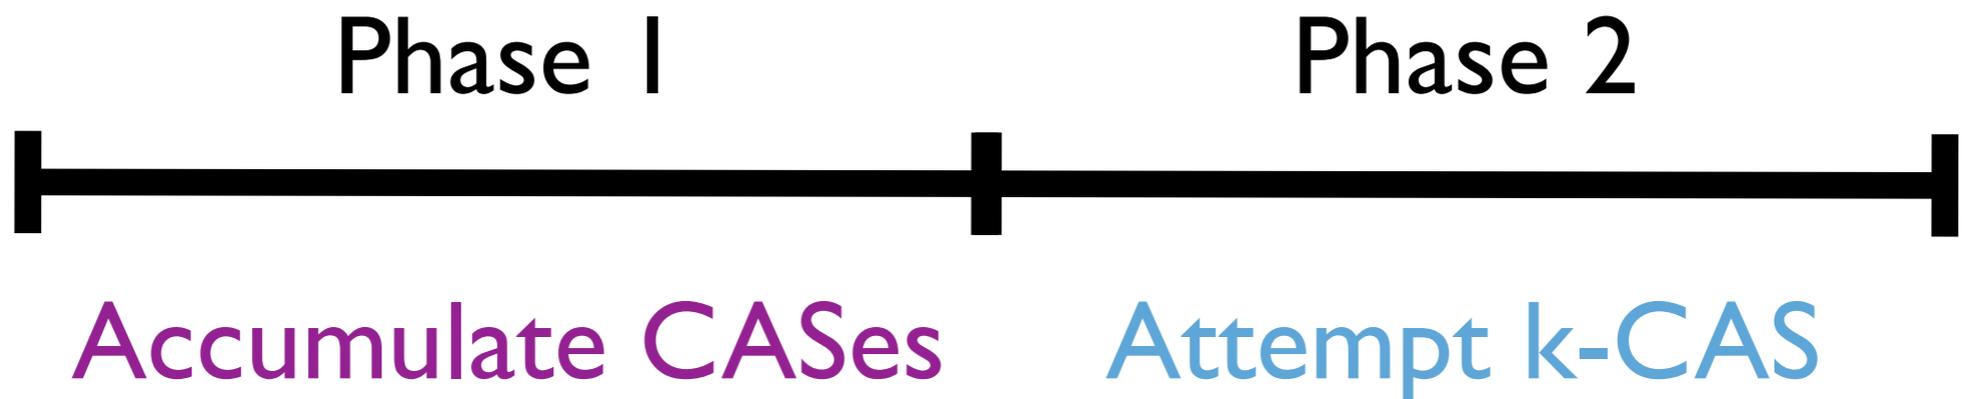


Phase 1

Phase 2



Accumulate CASes

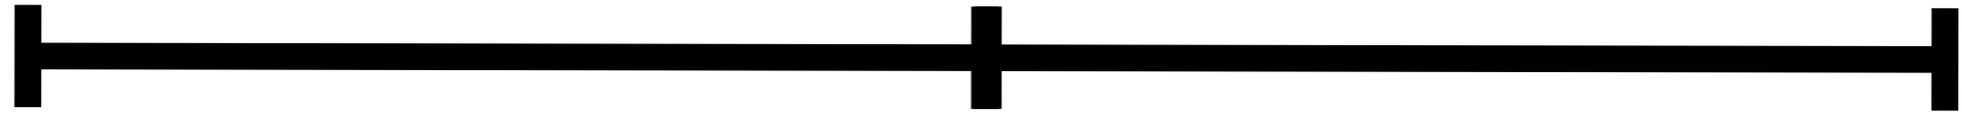


Phase 1

Phase 2

Accumulate CASes

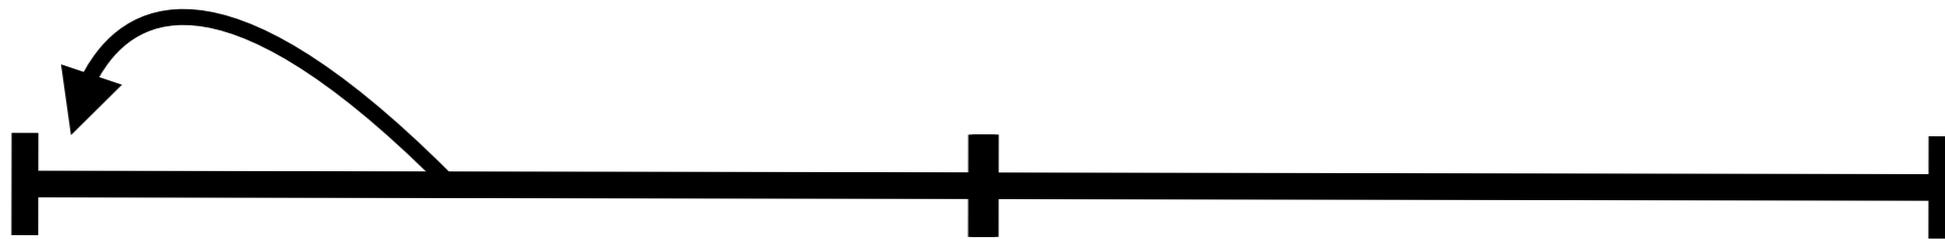
Attempt k-CAS



Accumulate CASes

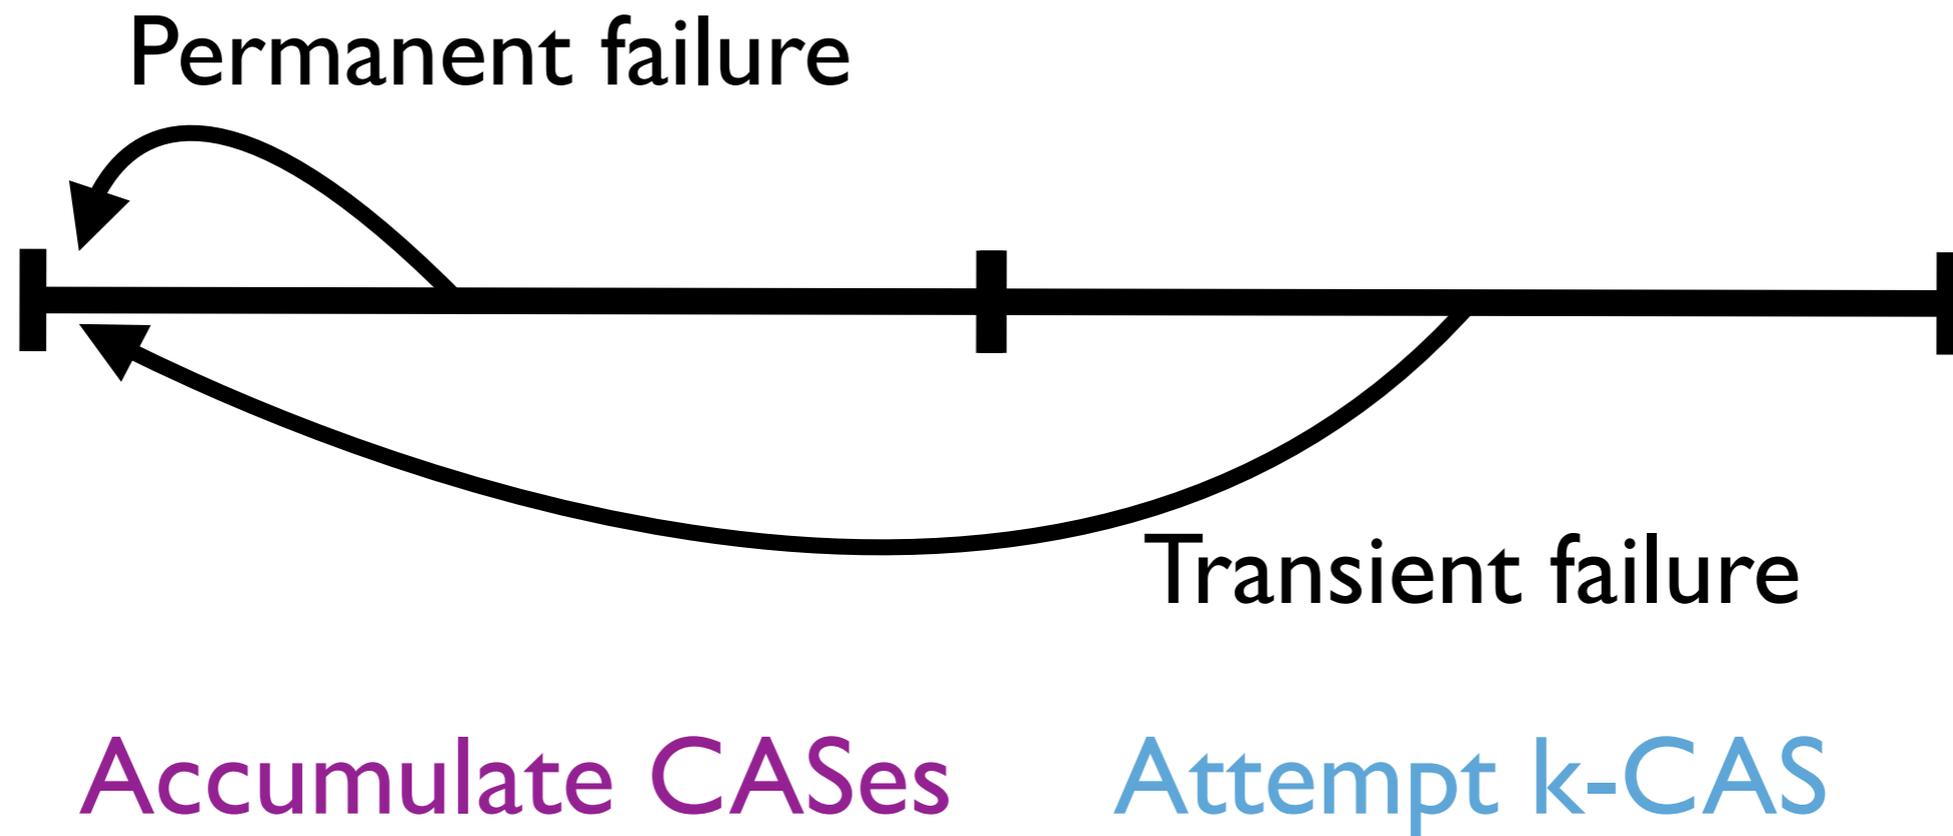
Attempt k-CAS

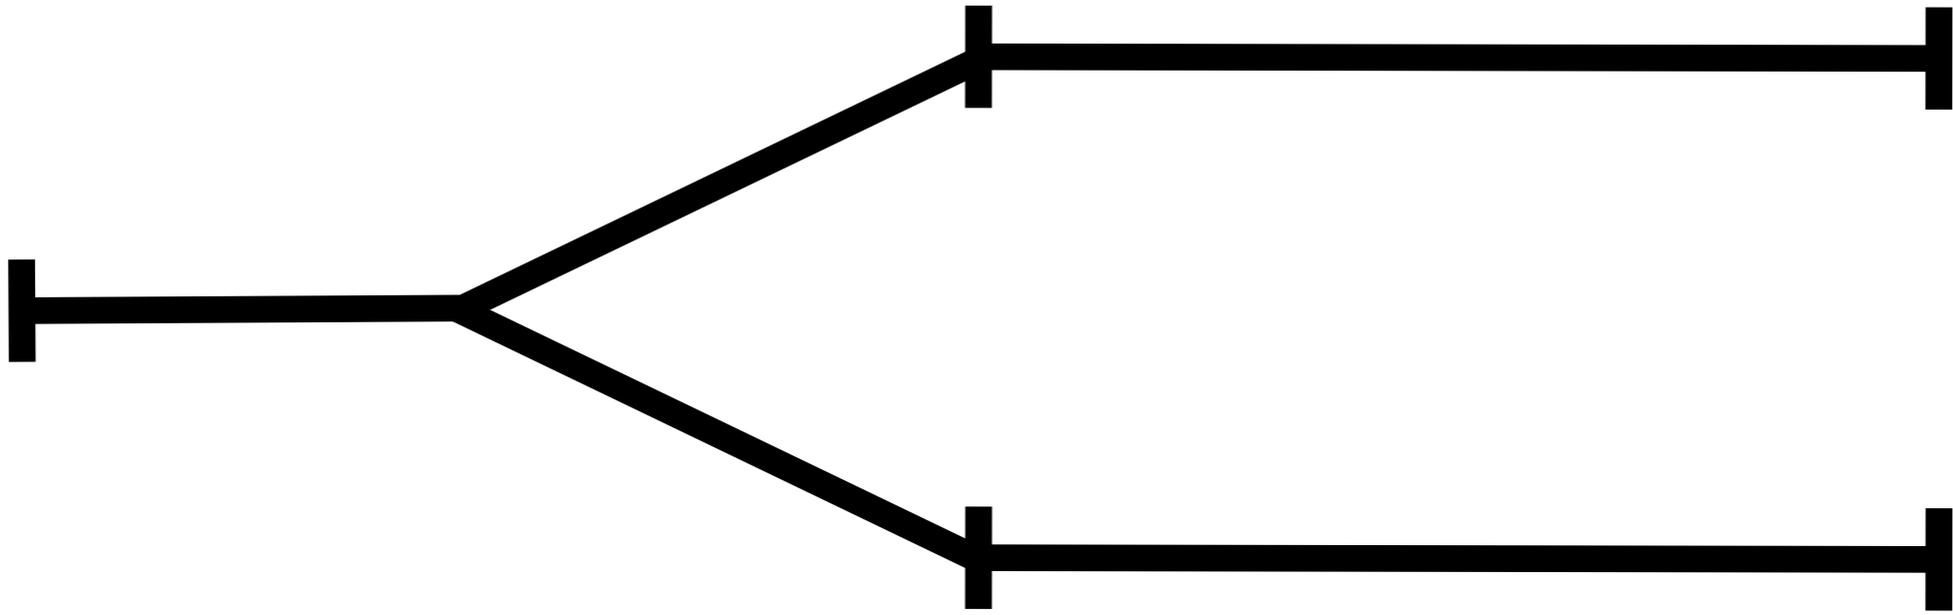
Permanent failure

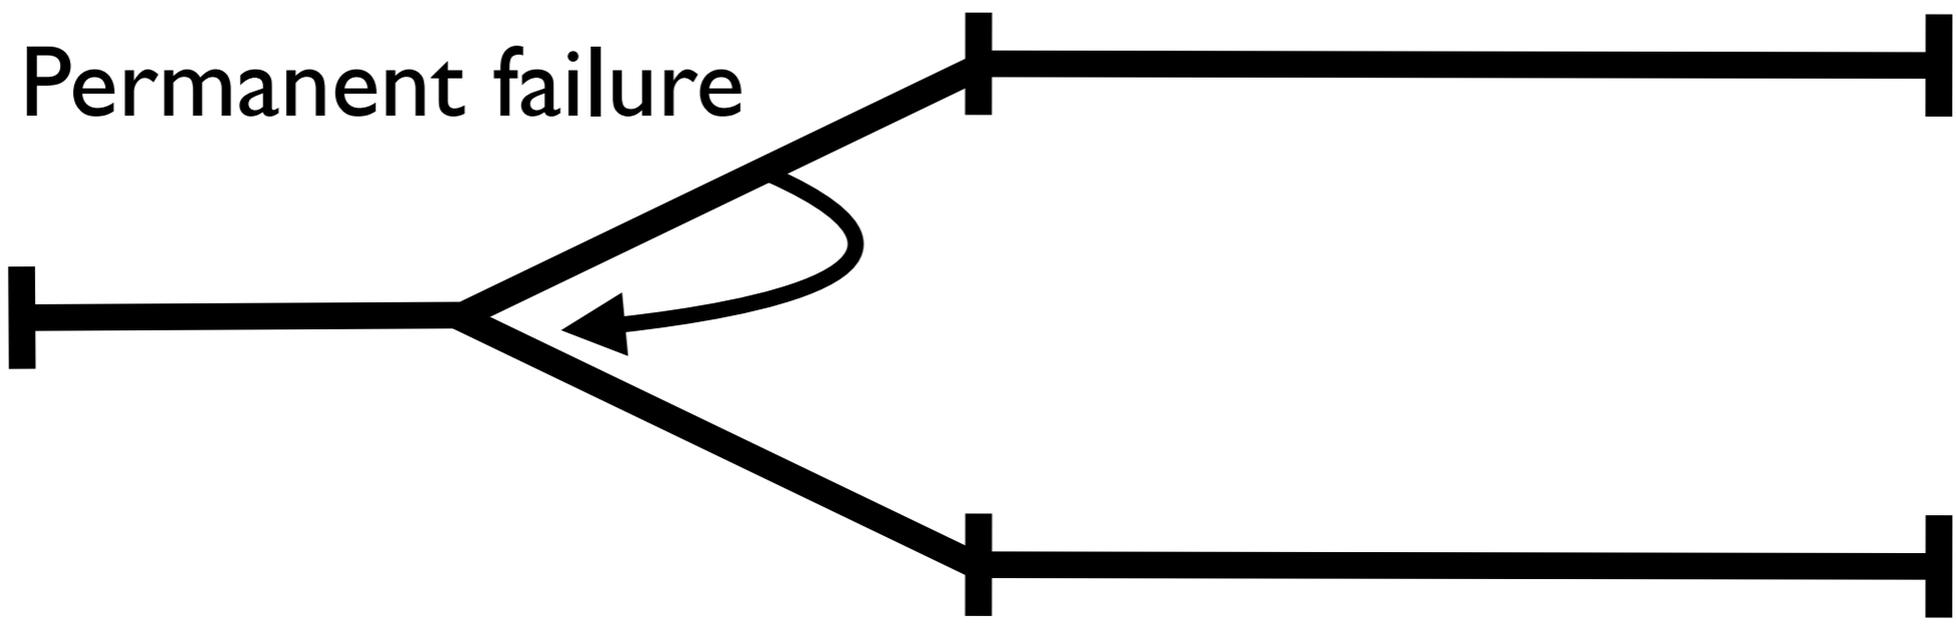


Accumulate CASes

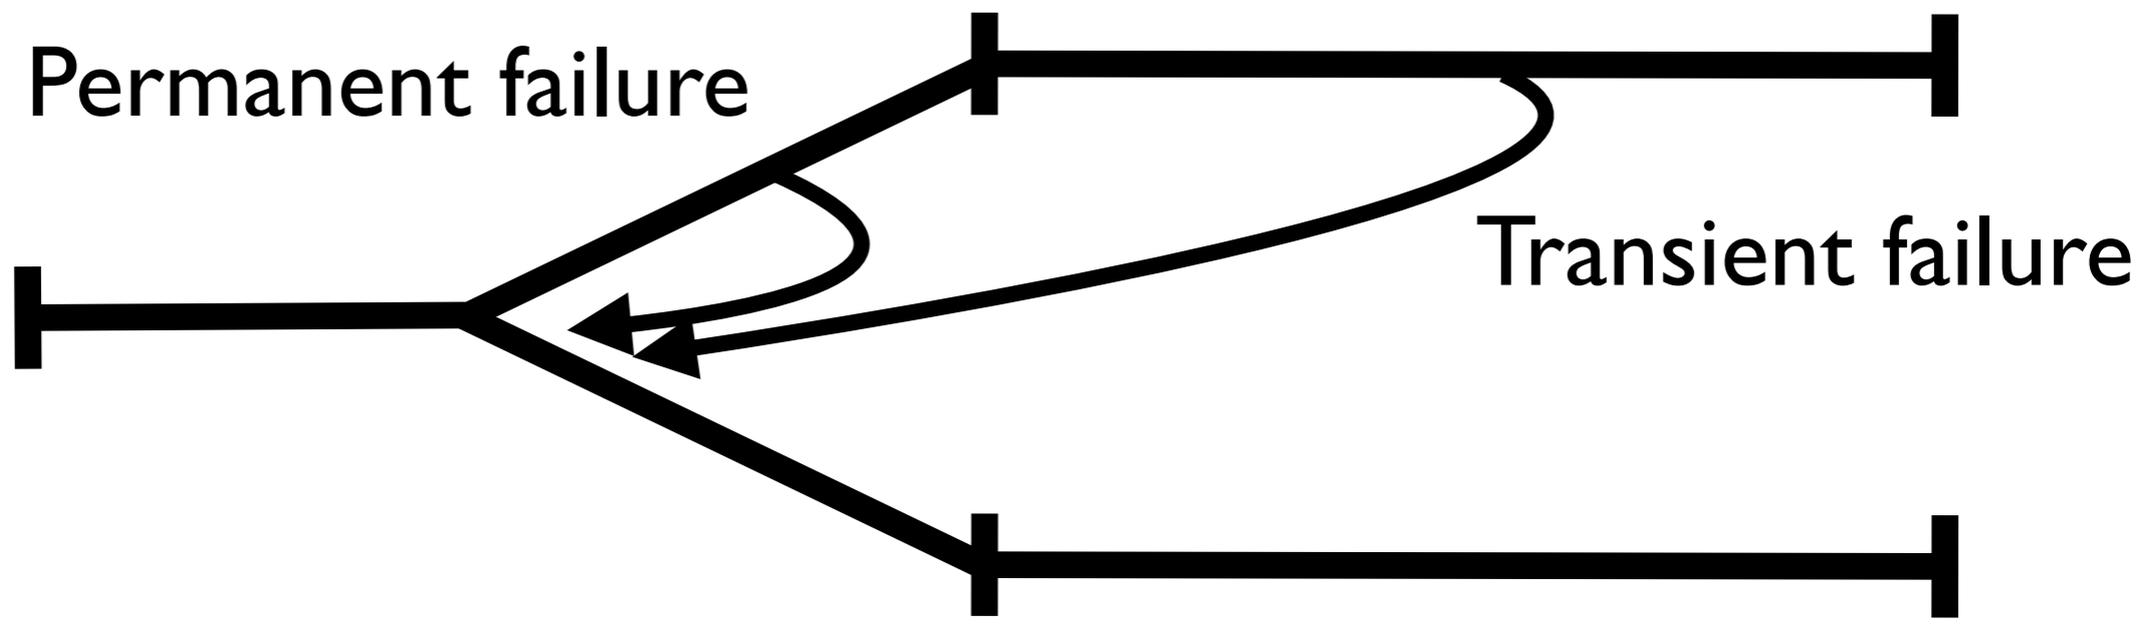
Attempt k-CAS

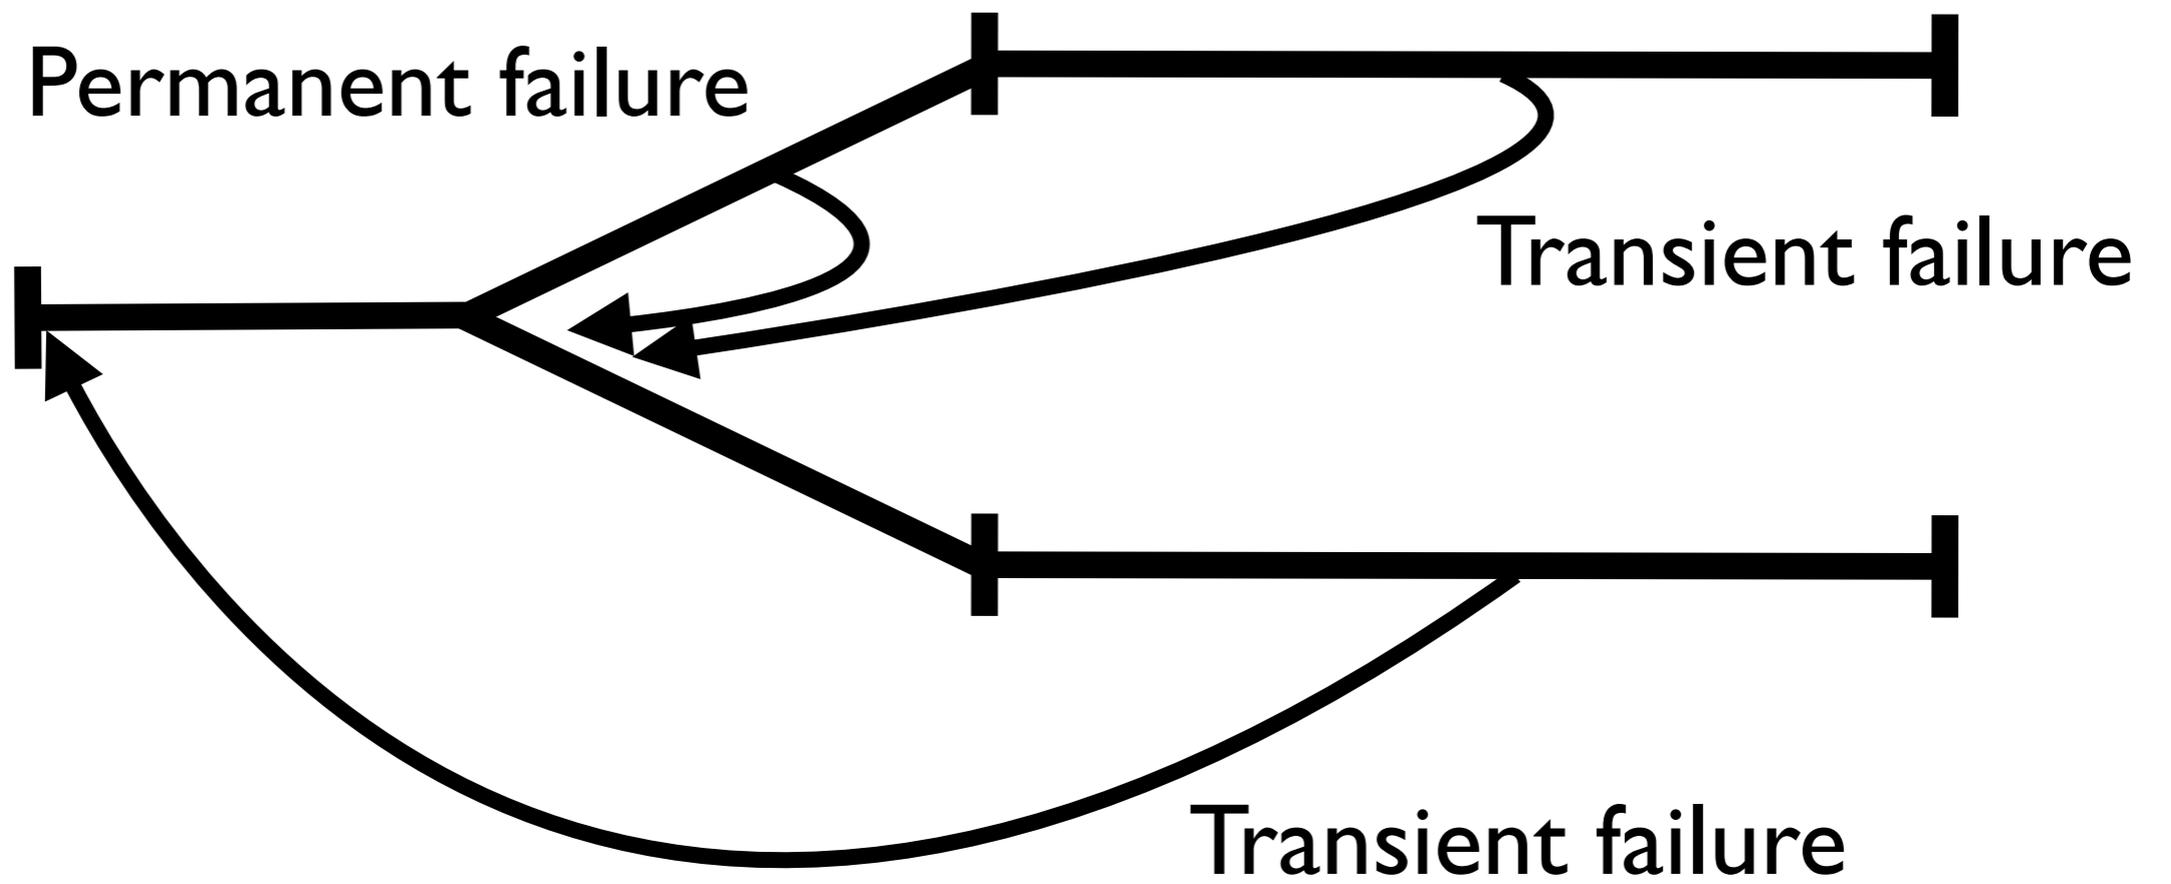


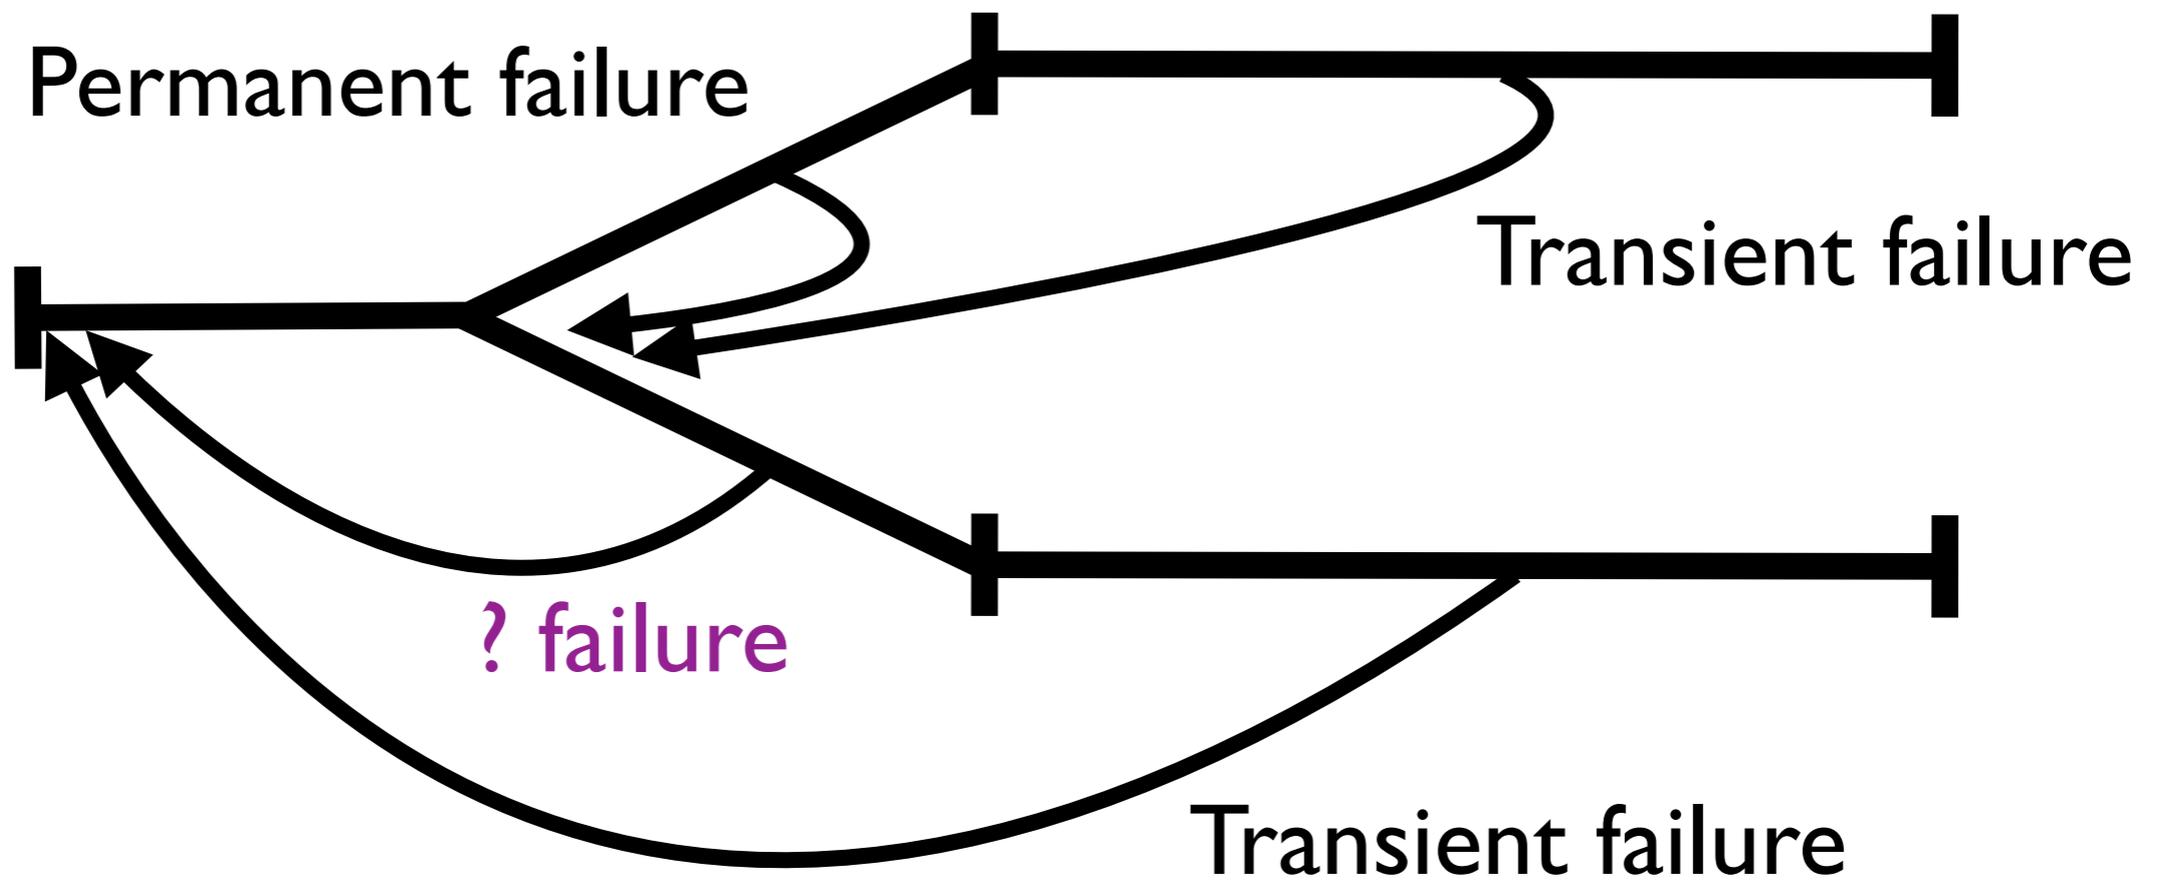


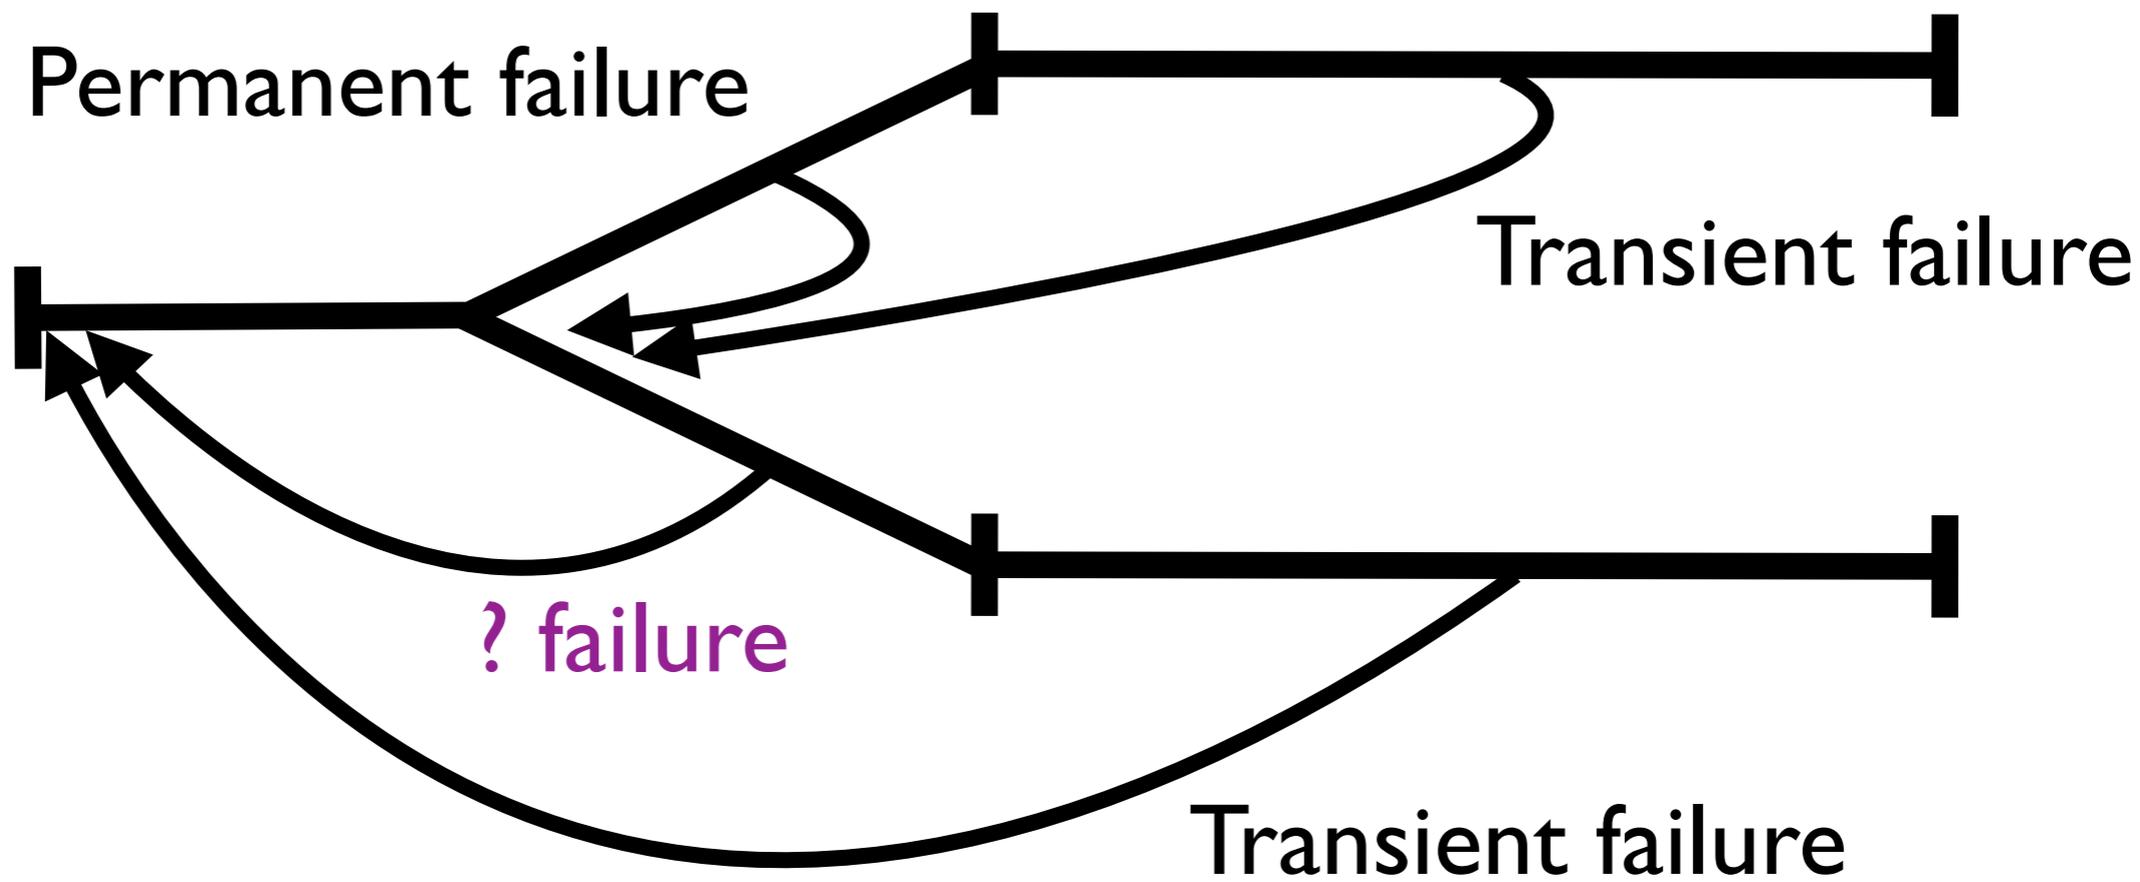


Permanent failure









$$P \ \& \ P = P$$

$$P \ \& \ T = T$$

$$T \ \& \ T = T$$

$$T \ \& \ P = T$$

**Is this just STM?**

# Is this just STM?

No:

- Single CAS collapses to single phase
  - Multiple CASes to single location forbidden
- So the “redo log” is write-only for phase 1!*

Therefore: **pay-as-you-go**

- Treiber stack is really a Treiber stack
- Pay for kCAS only for compositions

# Is this just STM?

**Isolation**  
Shared state

**Interaction**  
Message passing

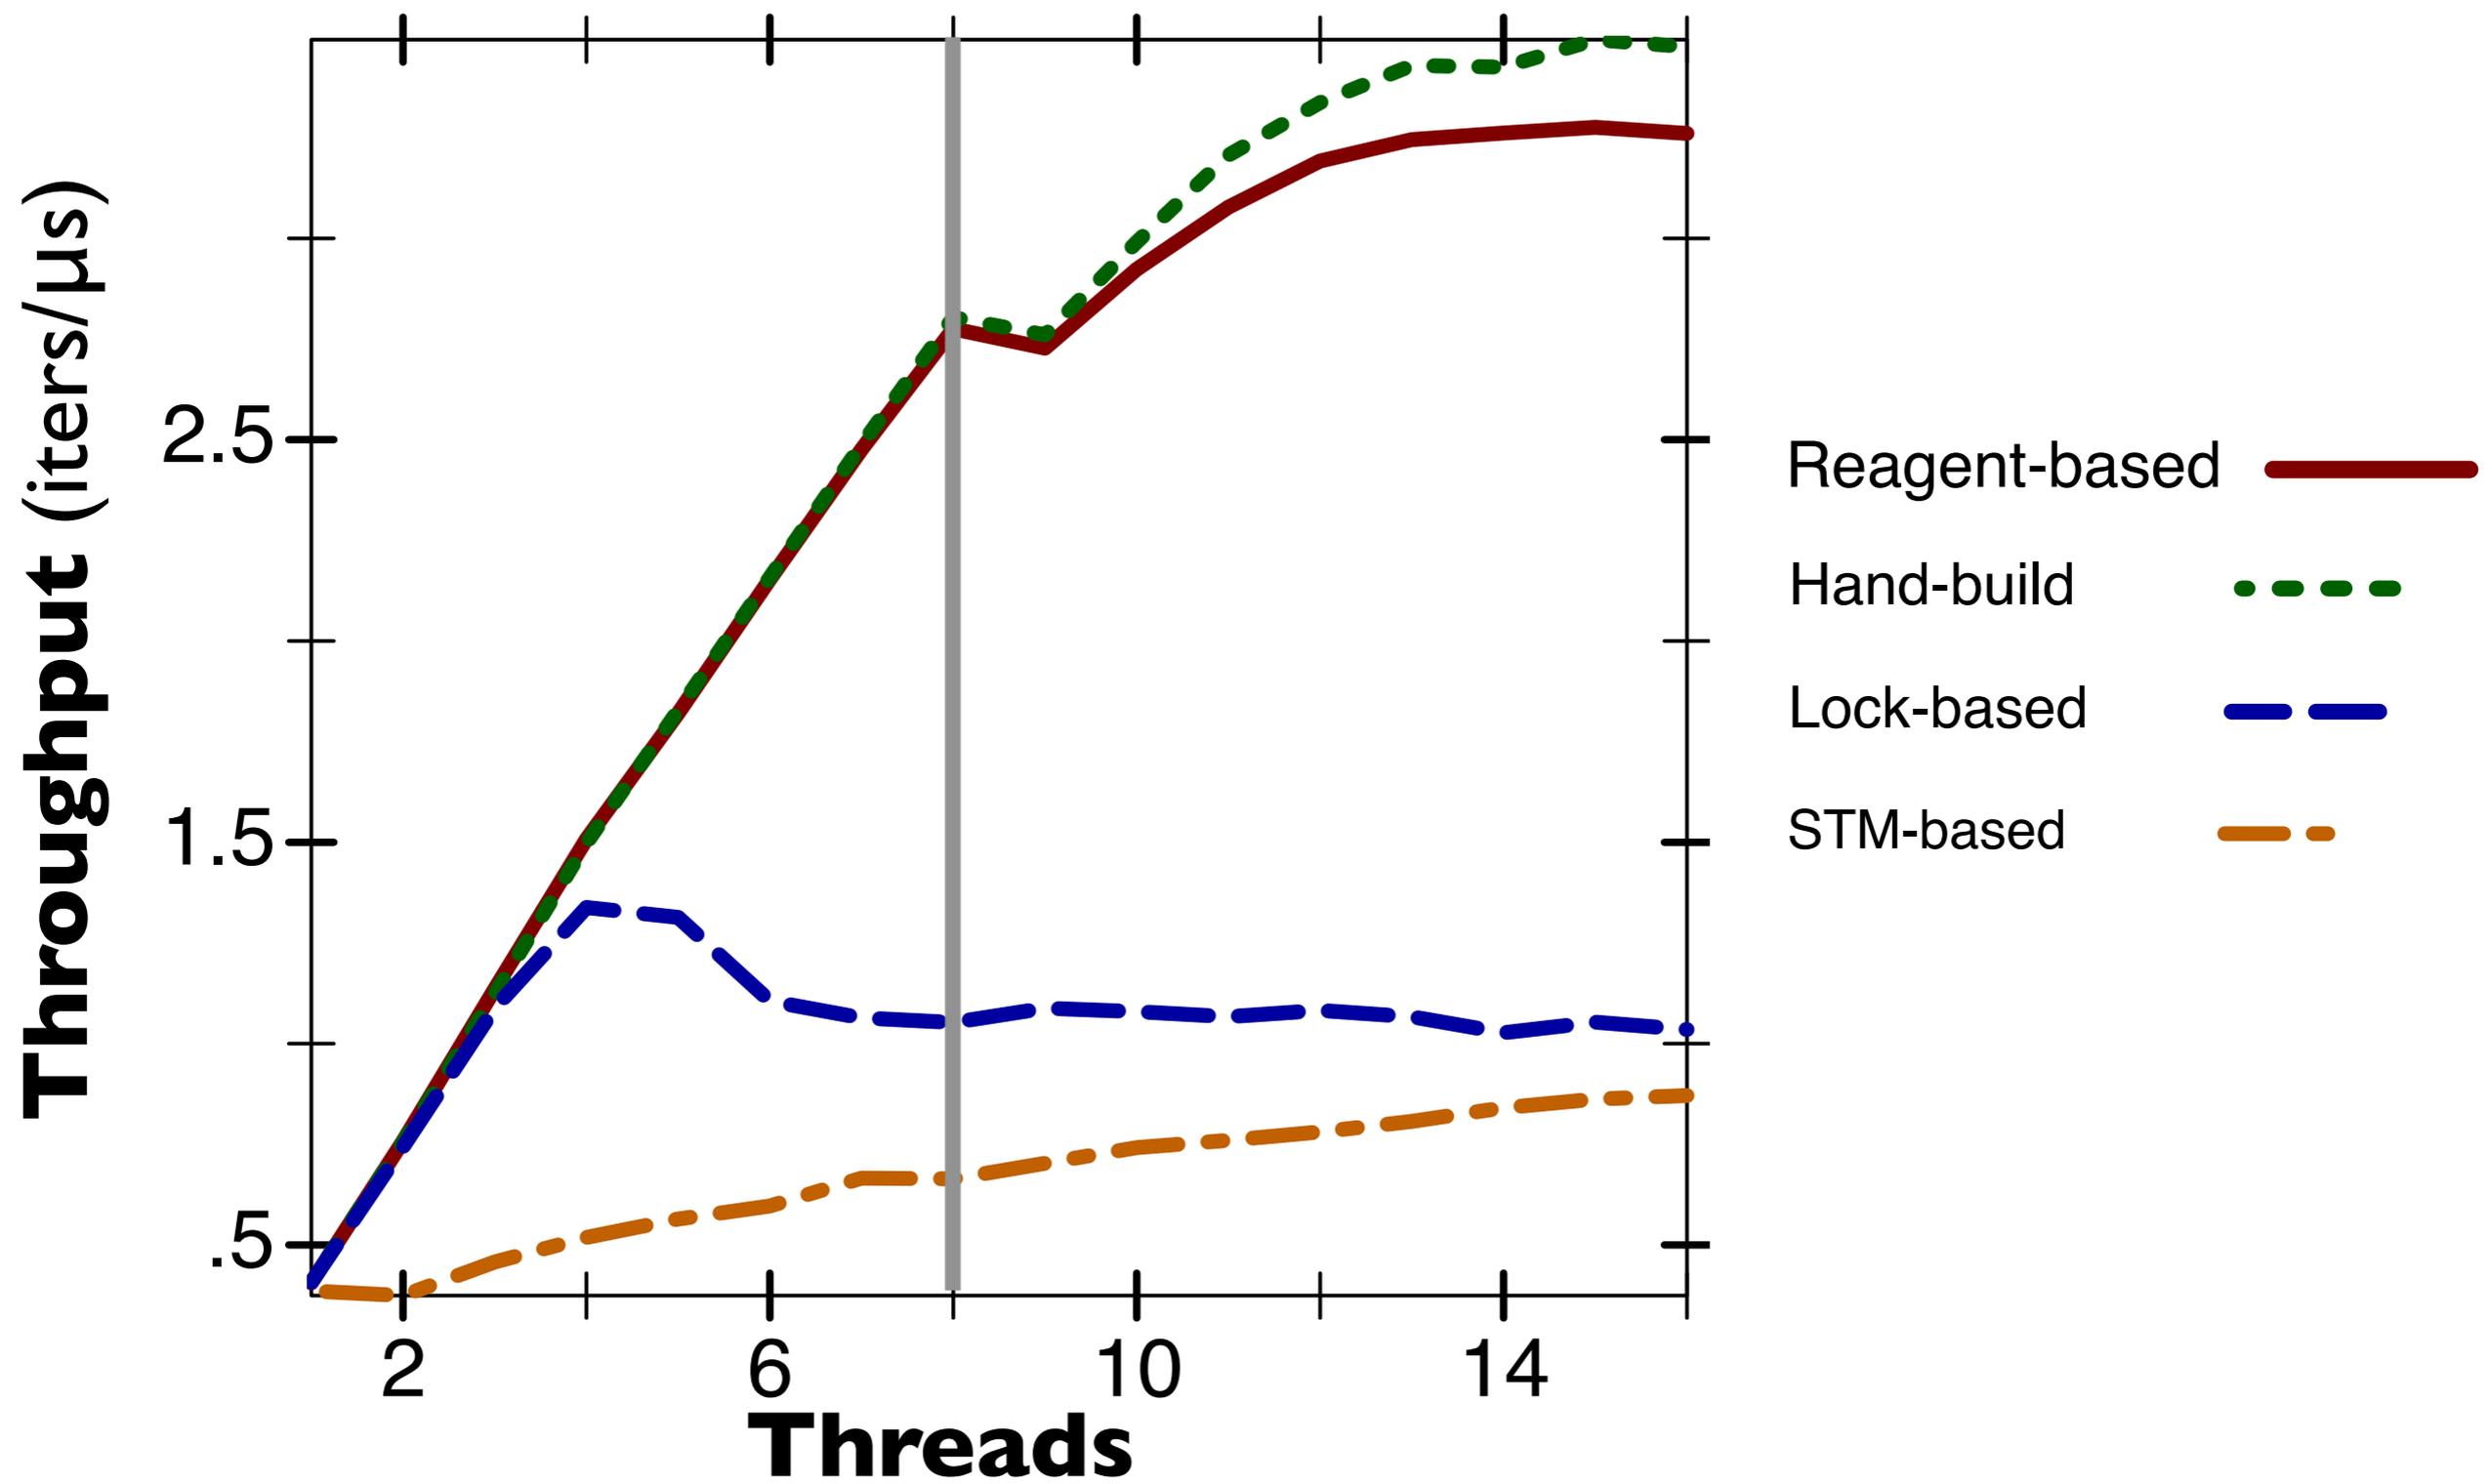
# Is this just STM?

**Isolation**  
Shared state

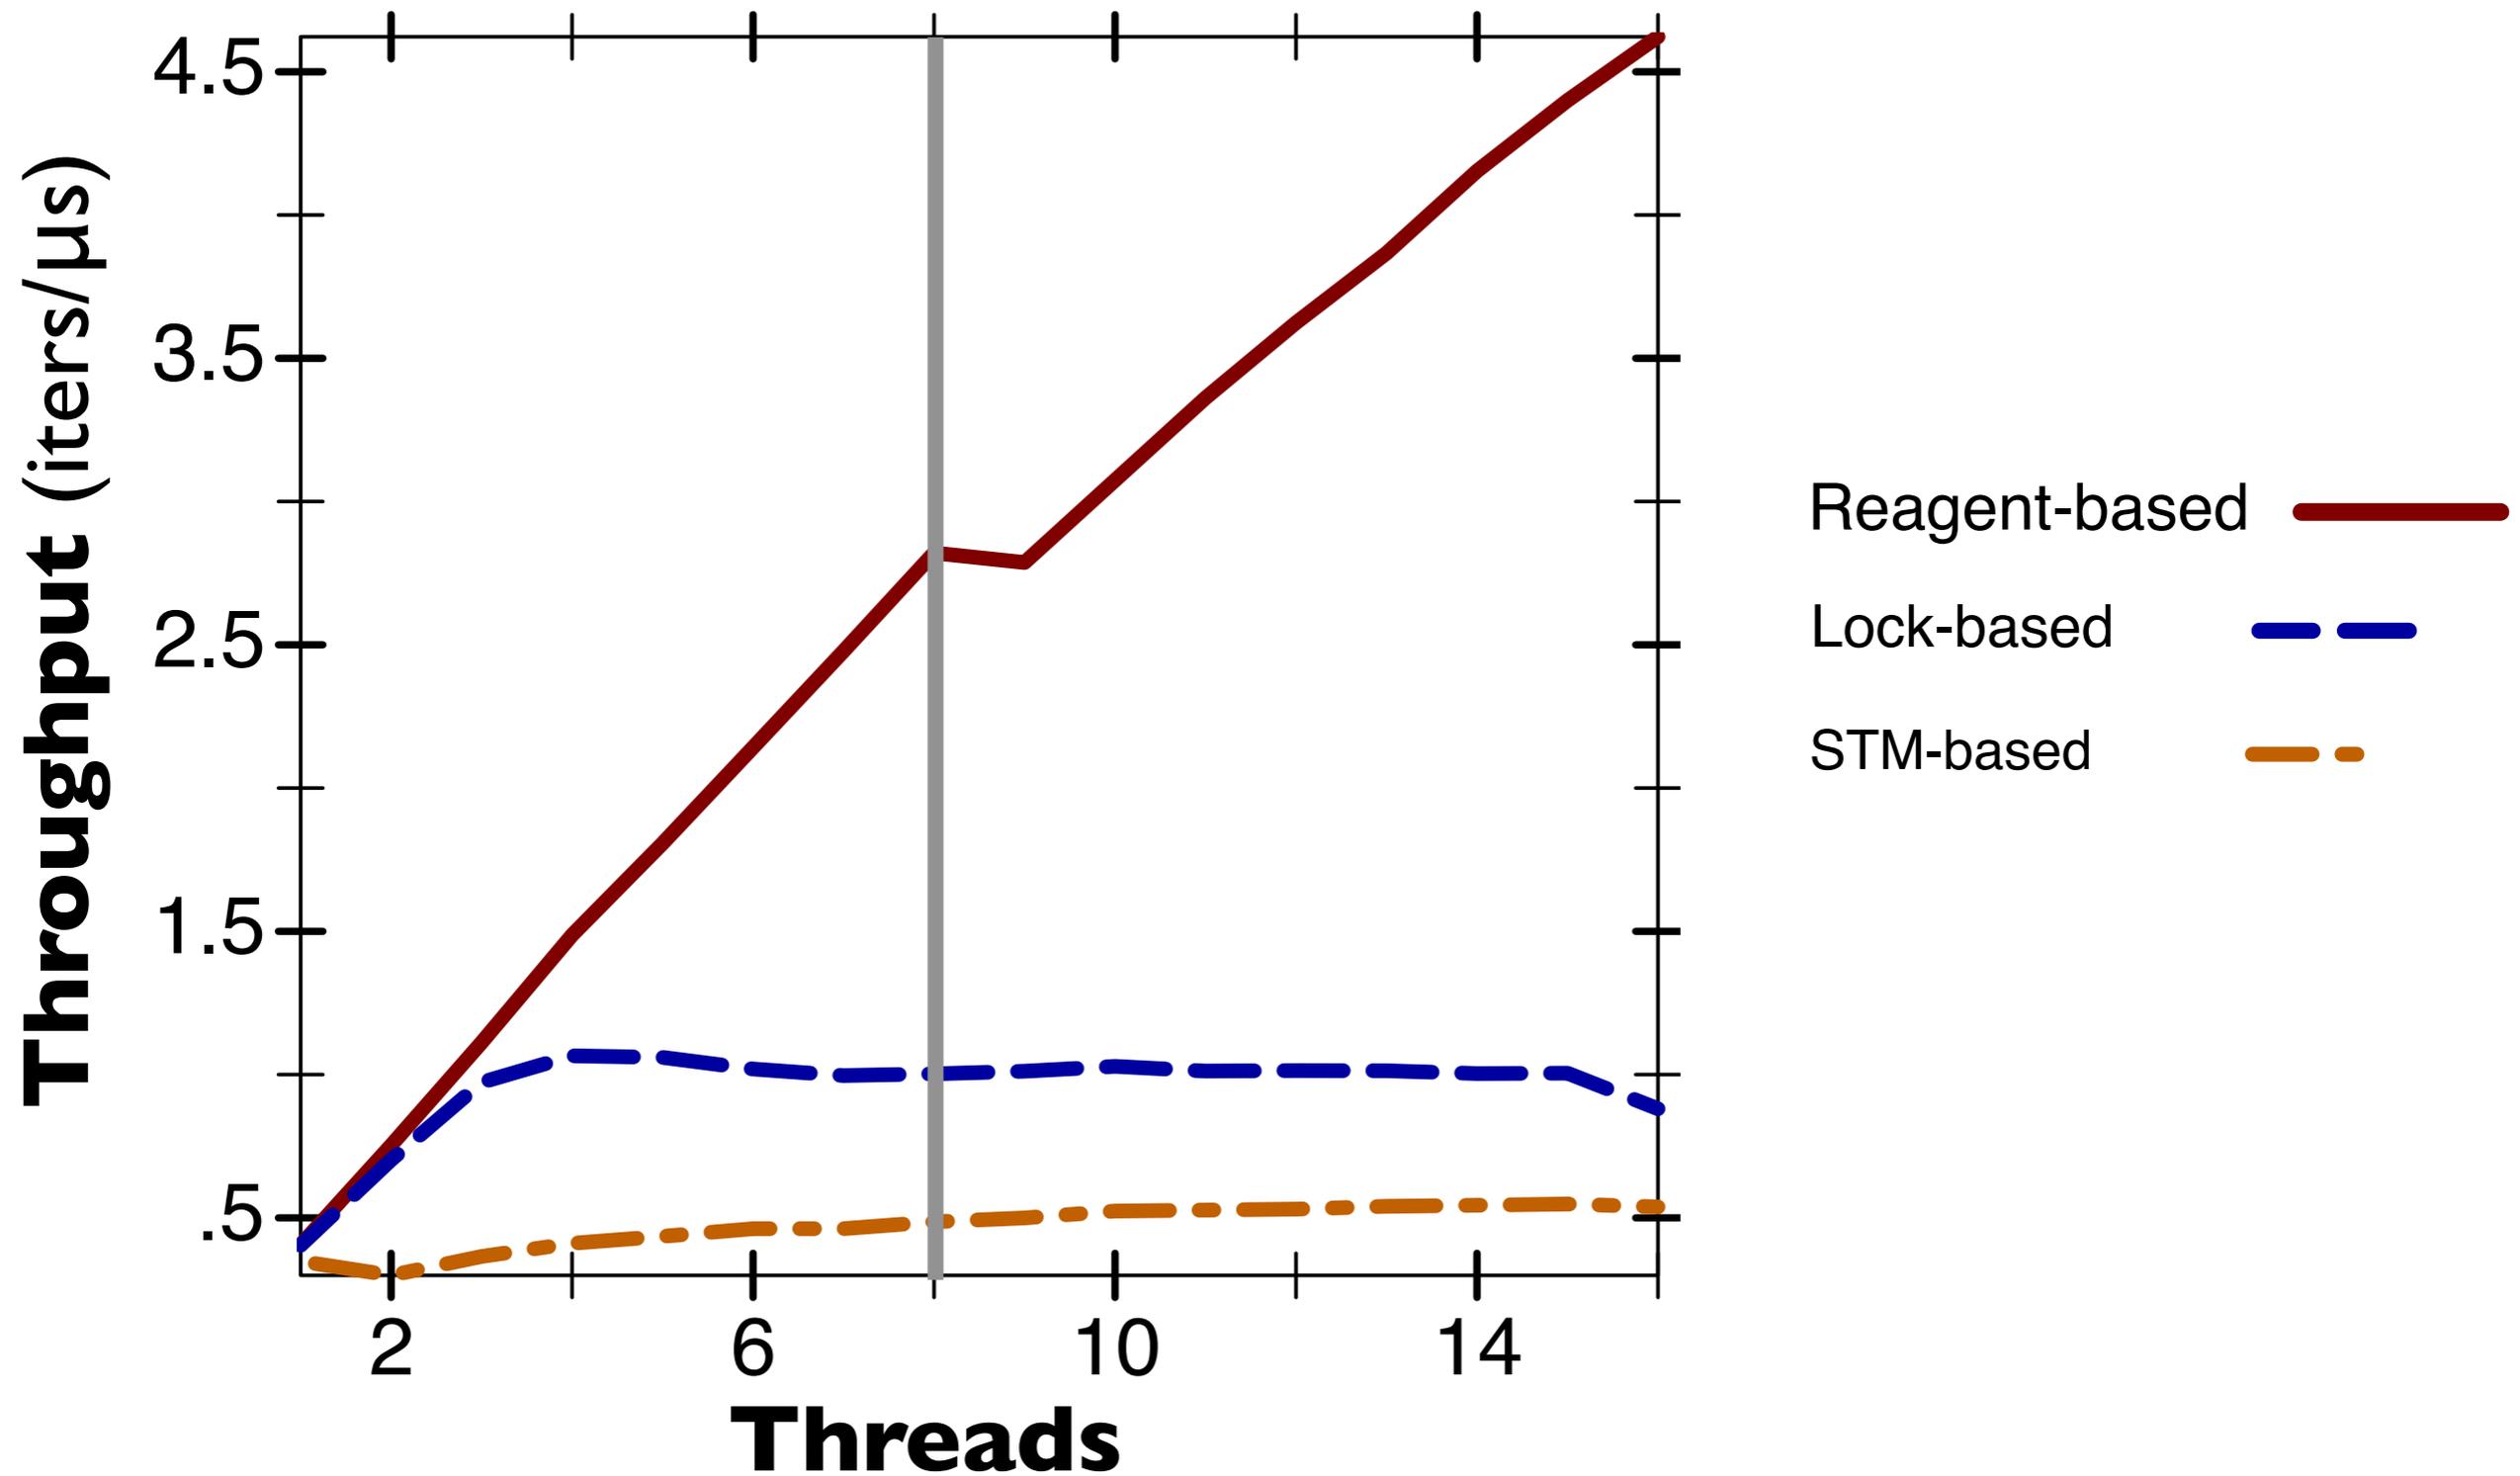
**Interaction**  
Message passing

Using lock-free bags,  
based on earlier work  
with Russo [OOPSLA'11]

# Treiber stack



# Stack transfer



# Open Questions

- Composition and invisible read/writes
  - Find a better rule?
  - Statically detect bad cases?
- Composition with *lock-based* algorithms?
- Conflicts between interaction and isolation?

# Open Questions 2

- Guaranteed inlining
  - Read/CAS windows must be short
  - “CAPER” with Sam Tobin-Hochstadt
- Formal semantics
  - Integrate Haskell’s STM semantics with message-passing?

# Related work

**Joins**

**CML**

**STM**

# Related work

**Joins**

**CML**

**STM**

**Transactional events**  
**Communicating transactions**