


Concurrent Programming: Critical Sections and Locks

CS 4410
Operating Systems



Cornell CIS
COMPUTING AND INFORMATION SCIENCE

[Robbert van Renesse]

An Operating System is a Concurrent Program

- The "kernel contexts" of each of the processes share many data structures
 - ready queue, wait queues, file system cache, and much more
- Sharing is further complicated by interrupt handlers that also access those data structures

Synchronization Lectures Outline

- What is the problem?
 - no determinism, no atomicity
- What is the solution?
 - some form of locks
- How to implement locks?
 - there are multiple ways
- How to specify concurrent problems?
 - atomic operations
- How to construct correct concurrent code?
 - invariants
- How to test concurrent programs
 - comparing behaviors

Concurrent Programming is Hard

Why?

- Concurrent programs are *non-deterministic*
 - run them twice with same input, get two different answers
 - or worse, one time it works and the second time it fails
- Program statements are executed *non-atomically*
 - **x += 1** compiles to something like
 - **LOAD x**
 - **ADD 1**
 - **STORE x**

Non-Determinism

```
1  shared = True
2
3  def f(): assert shared
4  def g(): shared = False
5
6  f()
7  g()
```

(a) [[code/prog1.hny](#)] Sequential

```
1  shared = True
2
3  def f(): assert shared
4  def g(): shared = False
5
6  spawn f()
7  spawn g()
```

(b) [[code/prog2.hny](#)] Concurrent

Figure 3.1: A sequential and a concurrent program.

Non-Determinism

```
1  shared = True
2
3  def f(): assert shared
4  def g(): shared = False
5
6  f()
7  g()
```

(a) [[code/prog1.hny](#)] Sequential

```
1  shared = True
2
3  def f(): assert shared
4  def g(): shared = False
5
6  spawn f()
7  spawn g()
```

(b) [[code/prog2.hny](#)] Concurrent

Figure 3.1: A sequential and a concurrent program.



#states 2
2 components, 0 bad states
No issues

- Schedule thread T0: **init()**
 - Line 1: Initialize shared to True
 - **Thread terminated**
- Schedule thread T2: **g()**
 - Line 4: Set shared to False (was True)
 - **Thread terminated**
- Schedule thread T1: **f()**
 - Line 3: Harmony assertion failed

Non-Atomicity

2 threads updating a shared variable **amount**

- One thread (you) wants to decrement amount by \$10K
- Other thread (IRS) wants to decrement amount by 50%



Memory



amount 100,000

What happens when both threads are running?

Non-Atomicity

2 threads updating a shared variable **amount**

- One thread (you) wants to decrement amount by \$10K
- Other thread (IRS) wants to decrement amount by 50%



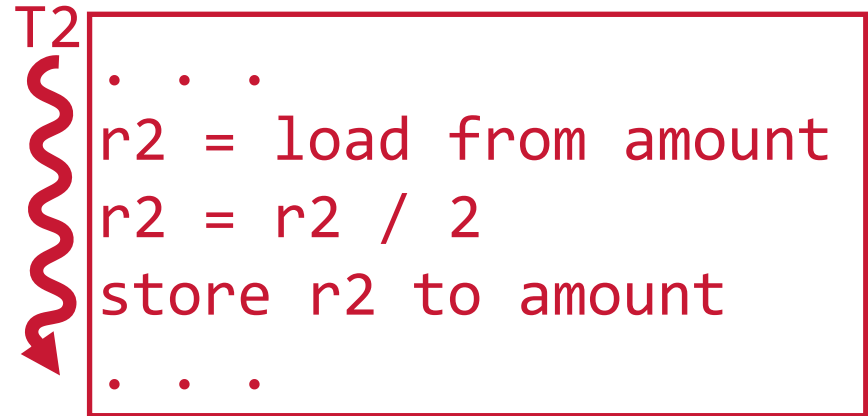
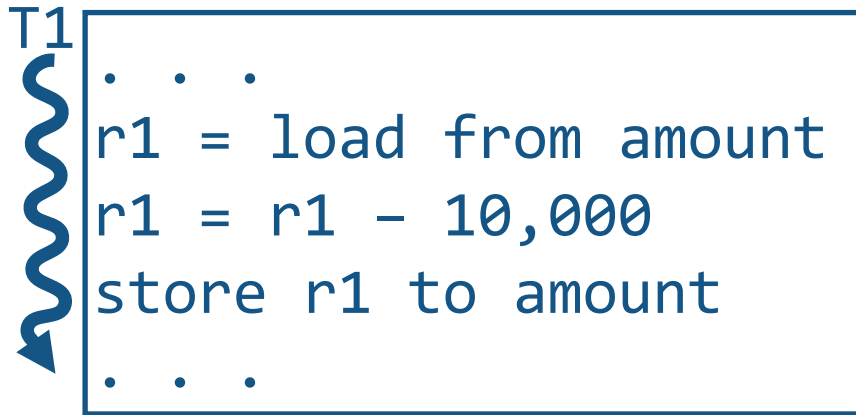
Memory

amount 100,000

What happens when both threads are running?

Non-Atomicity

Might execute like this:



Memory

amount

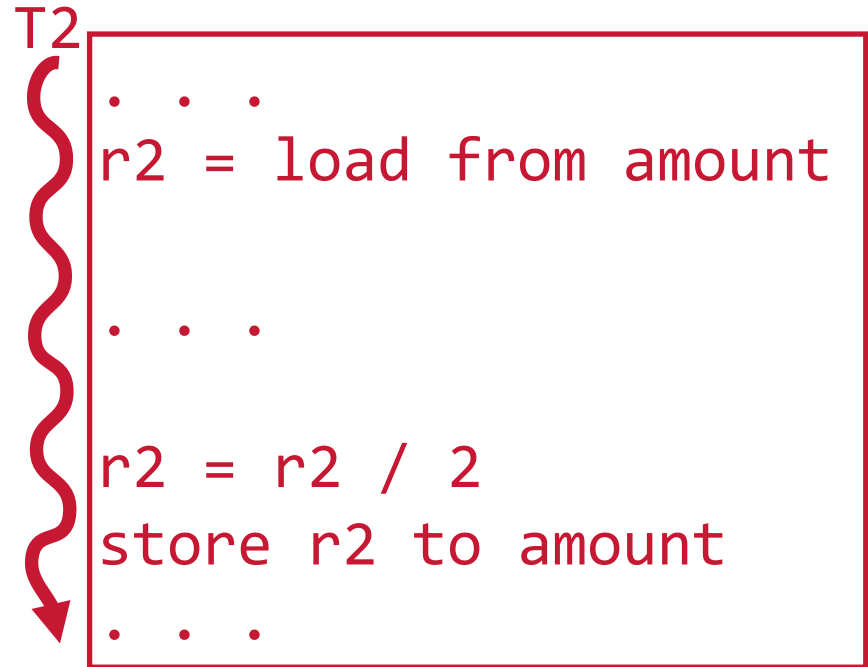
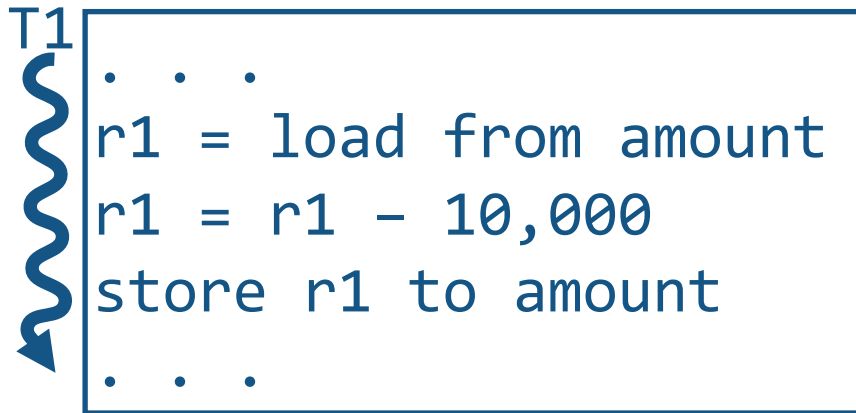
40,000

Or vice versa (T1 then T2 → 45,000)...

either way is fine...

Non-Atomicity

Or it might execute like this:



Memory

amount


50,000

Lost Update!

Wrong ..and very difficult to debug

Example: Races with Shared Queue

- 2 concurrent enqueue() operations?
- 2 concurrent dequeue() operations?



What could possibly go wrong?

Race Conditions

= ***timing dependent error involving shared state***

- Once thread A starts, it needs to “race” to finish
- Whether race condition happens depends on thread schedule
 - Different “schedules” or “interleavings” exist
(a schedule is a total order on machine instructions)

***All possible interleavings
should be safe!***

Race Conditions are Hard to Debug

- Number of possible interleavings is huge
- Some interleavings are good
- Some interleavings are bad
 - But bad interleavings may rarely happen!
 - Works 100x \neq no race condition
- Timing dependent: small changes hide bugs
 - add print statement \rightarrow bug no longer seems to happen

My experience until spring 2020

1. Students develop their code in Python or C
2. They test by running code many times
3. They submit their code, confident that it is correct
4. RVR tests the code with his secret and evil methods
 - uses homebrew library that randomly samples from possible interleavings (“fuzzing”)
5. Finds most submissions are broken
6. RVR unhappy, students unhappy

Why is that?

- Several studies show that heavily used code implemented, reviewed, and tested by expert programmers have lots of concurrency bugs
- Even professors who teach concurrency or write books and papers about concurrency get it wrong sometimes

Enter *Harmony*

- A new concurrent programming language
 - heavily based on Python syntax to reduce learning curve for many
- A new underlying virtual machine
 - quite different from any other:

it tries *all* possible executions of a program until it finds a problem, if any
(this is called “*model checking*”)

Example (same as before)

```
def T1():  
    amount -= 10000
```

```
def T2():  
    amount /= 2
```

```
spawn T1()  
spawn T2()
```

Harmony Machine Code

0 Jump 40

1 Frame T1 ()	
2 Load amount	T1a: LOAD amount
3 Push 10000	T1b: SUB 10000
4 2-ary -	
5 Store amount	T1c: STORE amount
6 Return	

```
def T1():  
    amount -= 10000
```

7 Frame T2 ()	T2a: LOAD amount
8 Load amount	T2b: DIV 2
9 Push 2	
10 2-ary /	T2c: STORE amount
11 Store amount	
12 Return	

```
def T2():  
    amount /= 2
```

Harmony Virtual Machine *State*

Three parts:

1. code (which never changes)
2. values of the shared variables
3. states of each of the running threads
 - “contexts”
 - PC, stack

State represents one vertex in the graph model

Simplified model

T1a: LOAD amount

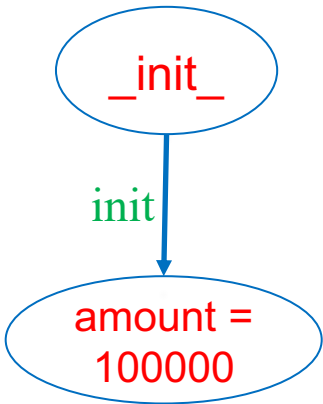
T1b: SUB 10000

T1c: STORE amount

T2a: LOAD amount

T2b: DIV 2

T2c: STORE amount



Simplified model (ignoring main)

T1a: LOAD amount


T1b: SUB 10000

T1c: STORE amount

T2a: LOAD amount

T2b: DIV 2

T2c: STORE amount



Simplified model (ignoring main)

T1a: LOAD amount

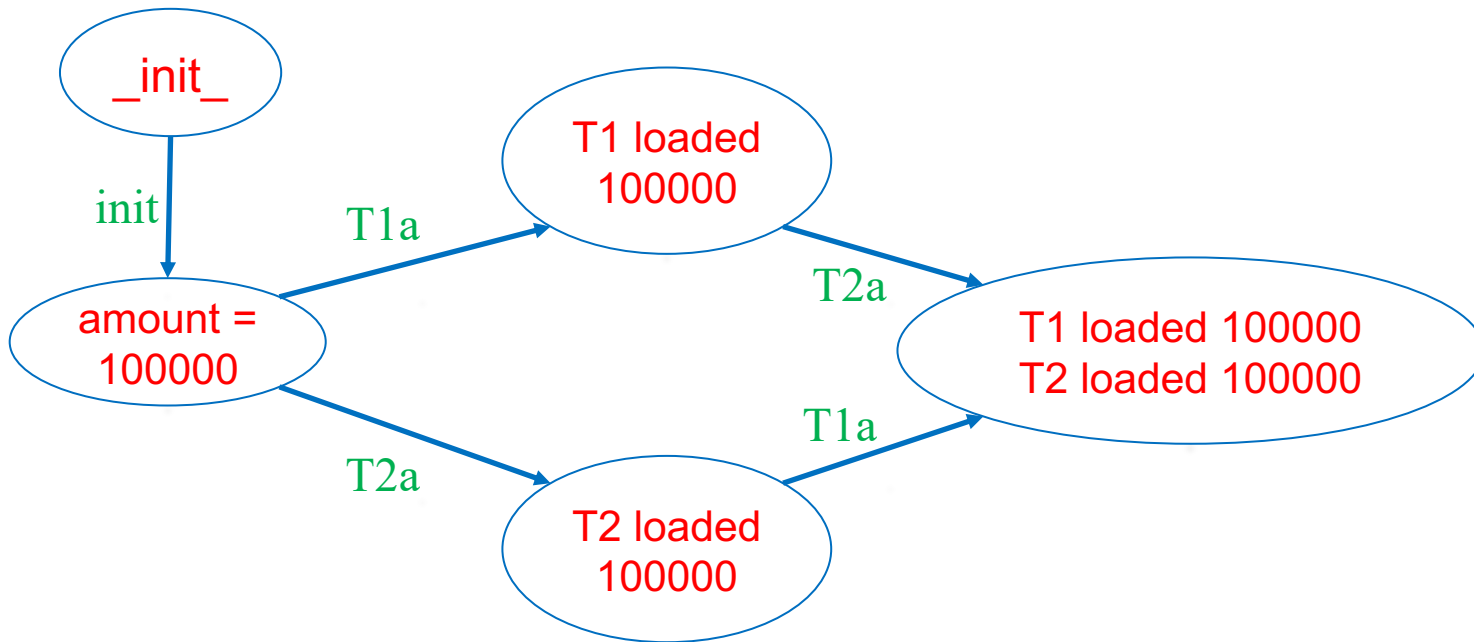
T1b: SUB 10000

T1c: STORE amount

T2a: LOAD amount

T2b: DIV 2

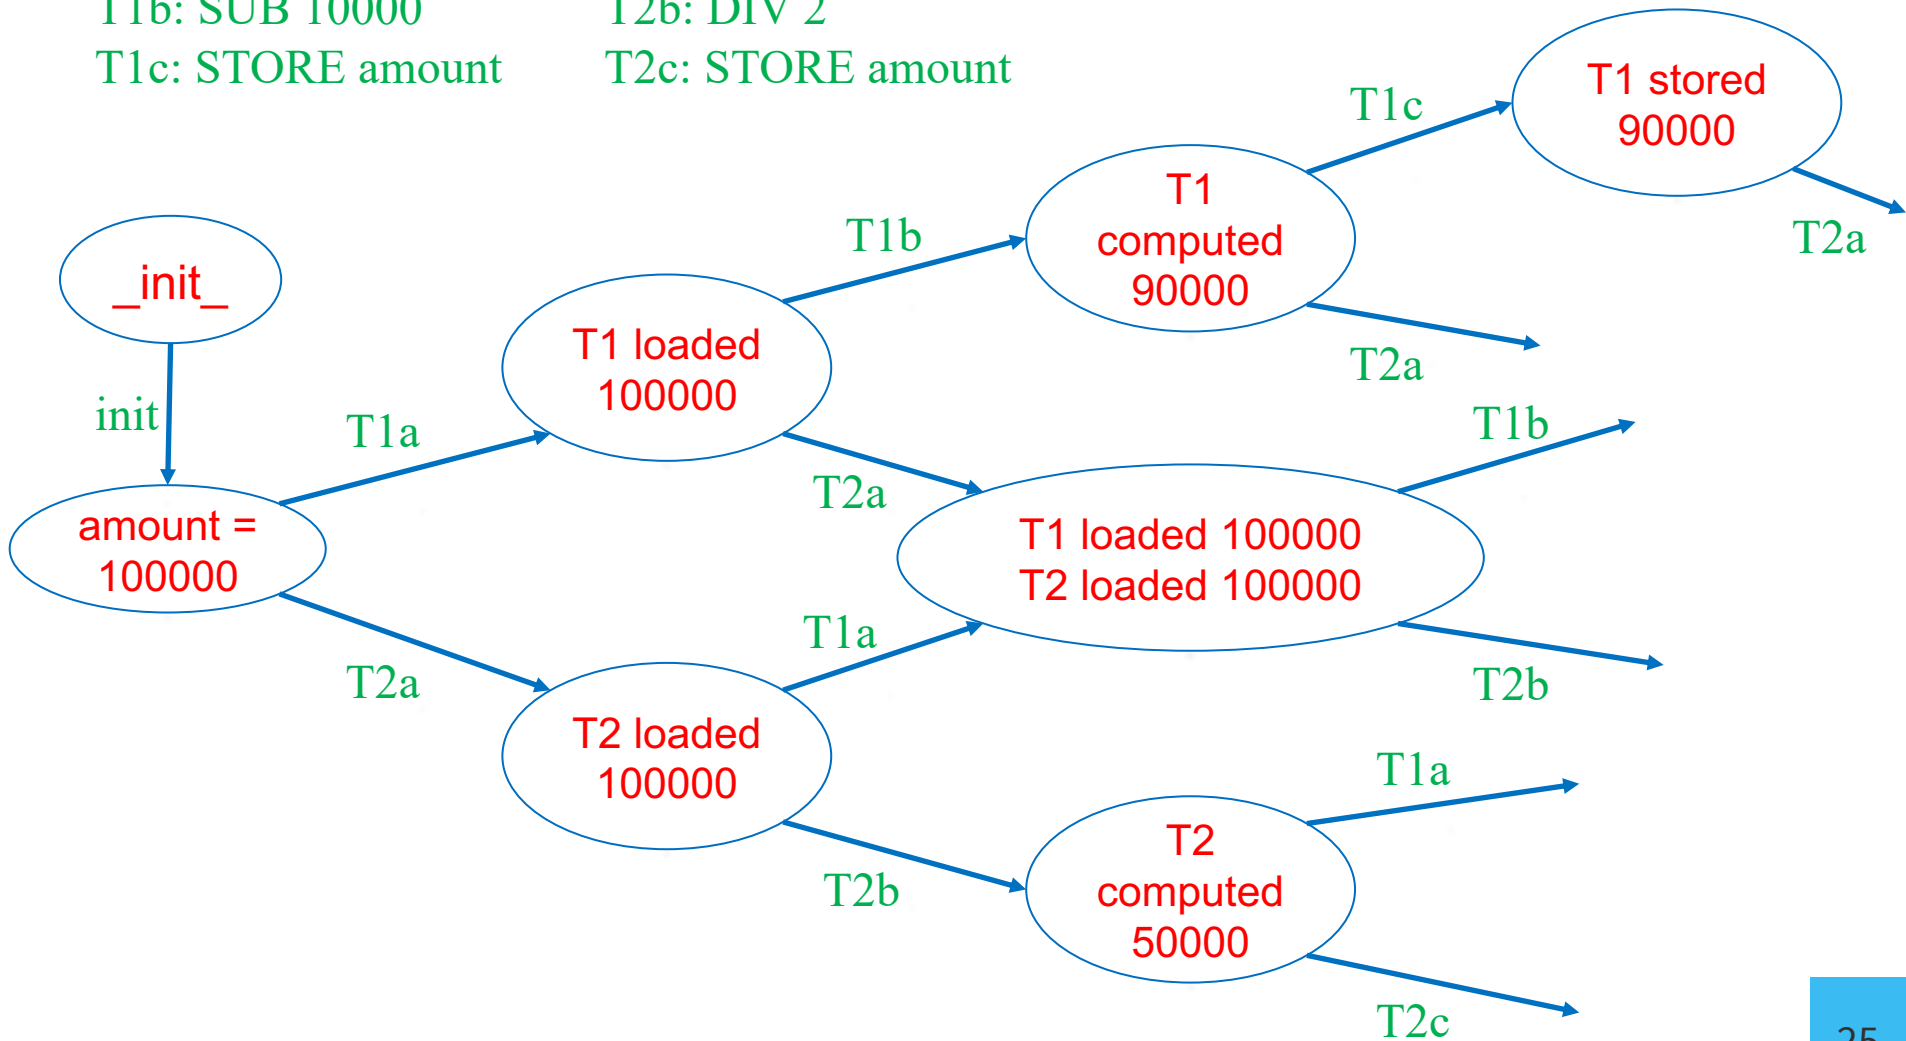
T2c: STORE amount



Simplified model (ignoring main)

T1a: LOAD amount
T1b: SUB 10000
T1c: STORE amount

T2a: LOAD amount
T2b: DIV 2
T2c: STORE amount



Harmony != Python

Harmony	Python
tries all possible executions	executes just one
(...) == [...] == ...	1 != [1] != (1)
1, == [1,] == (1,) != (1) == [1] == 1	[1,] == [1] != (1) == 1 != (1,)
f(1) == f 1 == f[1]	f 1 and f[1] are illegal (if f is method)
{ } is empty set	{ } is empty dictionary
few operator precedence rules --- use parentheses often	many operator precedence rules
variables global unless declared otherwise	depends... Sometimes must be explicitly declared global
no return , break , continue	various flow control escapes
no classes	object-oriented
...	...

I/O in Harmony?

- Input:
 - **choose** expression
 - $x = \mathbf{choose}(\{ 1, 2, 3 \})$
 - allows Harmony to know all possible inputs
 - **const** expression
 - **const** $x = 3$
 - can be overridden with “-c x=4” flag to harmony
 - Output:
 - **print** $x + y$
 - **assert** $x + y < 10, (x, y)$

I/O in Harmony?

- Input:
 - **choose** expression
 - $x = \mathbf{choose}(\{ 1, 2, 3 \})$
 - allows Harmony to choose between multiple inputs
 - **const** expression
 - **c** = constant
 - can be used to restrict the number of choices with “-c x=4” flag to harmony
 - Output:
 - **print** $x + y$
 - **assert** $x + y < 10, (x, y)$

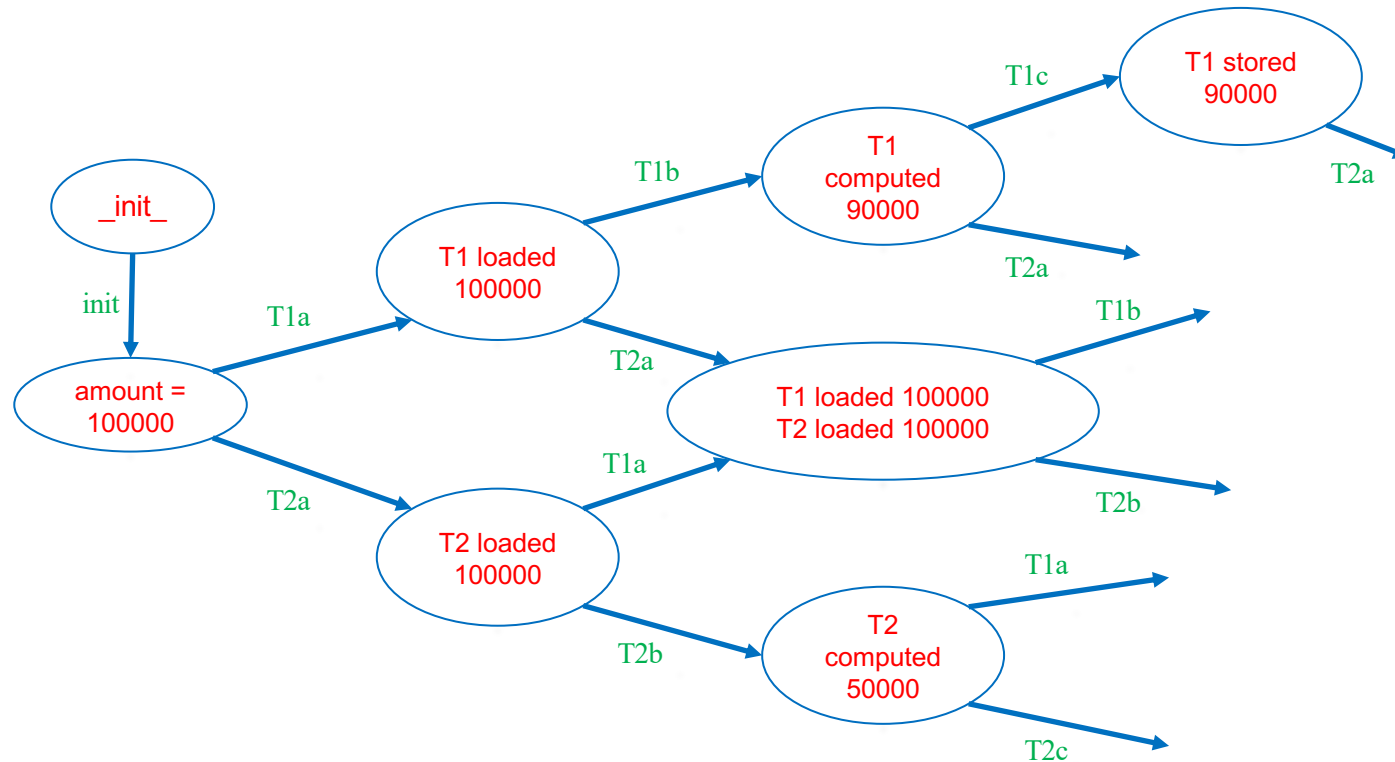
No open(), read(), or
input() statements

Non-determinism in Harmony

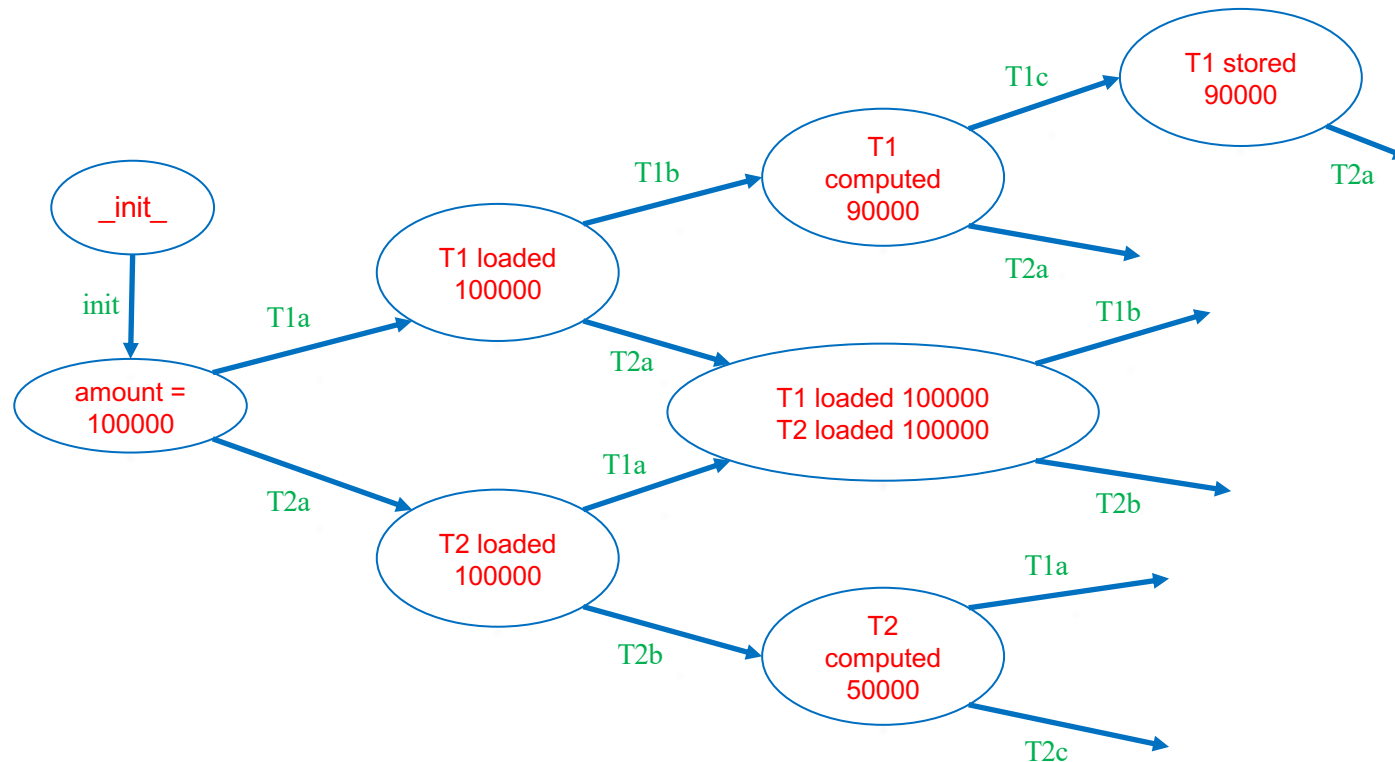
Three sources:

1. **choose** expressions
2. thread interleavings
3. Interrupts

Limitation: models must be finite!



Limitation: models must be finite!





- But models are allowed to have cycles.
- Executions are allowed to be unbounded!
- Harmony checks for *possibility* of termination

Back to our problem...

2 threads updating a shared variable **amount**

- One thread wants to decrement amount by \$10K
- Other thread wants to decrement amount by 50%



Memory


amount

100000

How to “serialize” these executions?

Critical Section

Must be serialized due to shared memory access



Goals

Mutual Exclusion: 1 thread in a critical section at time

Progress: a thread can get in when there's no other thread

Fairness: equal chances of getting into CS

... in practice, fairness rarely guaranteed or needed

Critical Section

Must be serialized due to shared memory access



Goals

Mutual Exclusion: 1 thread in a critical section at time

Progress: a thread can get in when there's no other thread

Fairness: equal chances of getting into CS

... in practice, fairness rarely guaranteed or needed

Mutual Exclusion and Progress

- Need both:
 - either one is trivial to achieve by itself

Critical Sections in Harmony

```
1  def thread():
2      while True:
3          # Critical section is here
4          pass
5
6  spawn thread()
7  spawn thread()
```

- How do we check mutual exclusion?
- How do we check progress?

Specifying Critical Sections in Harmony

```
1 # number of threads in the critical section
2 in_cs = 0
3 invariant in_cs in { 0, 1 }
4
5 def thread():
6     while choose { False, True }:
7         # Enter critical section
8         atomically in_cs += 1
9
10        # Critical section is here
11        pass
12
13        # Exit critical section
14        atomically in_cs -= 1
15
16 spawn thread()
17 spawn thread()
```

- How do we check mutual exclusion?
- How do we check progress?

Specifying Critical Sections in Harmony

```
1 # number of threads in the critical section
2 in_cs = 0
3 invariant in_cs in { 0, 1 }
4
5 def thread():
6     while choose { False, True }:
7         # Enter critical section
8         atomically in_cs += 1
9
10        # Critical section is here
11        pass
12
13        # Exit critical section
14        atomically in_cs -= 1
15
16 spawn thread()
17 spawn thread()
```

mutual exclusion

- How do we check mutual exclusion?
- How do we check progress?

Specifying Critical Sections in Harmony

```
1 # number of threads in the critical section
2 in_cs = 0
3 invariant in_cs in { 0, 1 }
4
5 def thread():
6     while choose { False, True }:
7         # Enter critical section
8         atomically in_cs += 1
9
10        # Critical section is here
11        pass
12
13        # Exit critical section
14        atomically in_cs -= 1
15
16 spawn thread()
17 spawn thread()
```

mutual exclusion

do zero or more times

- How do we check mutual exclusion?
- How do we check progress?

Specifying Critical Sections in Harmony

```
1 # number of threads in the critical section
2 in_cs = 0
3 invariant in_cs in { 0, 1 }
4
5 def thread():
6     while choose { False, True }:
7         # Enter critical section
8         atomically in_cs += 1
9
10        # Critical section is here
11        pass
12
13        # Exit critical section
14        atomically in_cs -= 1
15
16 spawn thread()
17 spawn thread()
```

mutual exclusion

do zero or more times

increment in_cs

- How do we check mutual exclusion?
- How do we check progress?

Specifying Critical Sections in Harmony

```
1 # number of threads in the critical section
2 in_cs = 0
3 invariant in_cs in { 0, 1 }
4
5 def thread():
6     while choose { False, True }:
7         # Enter critical section
8         atomically in_cs += 1
9
10        # Critical section is here
11        pass
12
13        # Exit critical section
14        atomically in_cs -= 1
15
16 spawn thread()
17 spawn thread()
```

mutual exclusion

do zero or more times

increment in_cs

execute critical section

- How do we check mutual exclusion?
- How do we check progress?

Specifying Critical Sections in Harmony

```
1 # number of threads in the critical section
2 in_cs = 0
3 invariant in_cs in { 0, 1 }
4
5 def thread():
6     while choose { False, True }:
7         # Enter critical section
8         atomically in_cs += 1
9
10        # Critical section is here
11        pass
12
13        # Exit critical section
14        atomically in_cs -= 1
15
16 spawn thread()
17 spawn thread()
```

mutual exclusion

do zero or more times

increment in_cs

execute critical section

decrement in_cs

Progress: Harmony checks that all thread *can* terminate

Specification vs implementation

- Spec is fine, but this is an O.S. class!
- Sounds like we need a lock
- The question is:

How does one build a lock?

- Harmony is a concurrent programming language. *Really, doesn't Harmony have locks?*

You have to build them too!

First attempt: a naïve lock

```
1  in_cs = 0
2  invariant in_cs in { 0, 1 }
3
4  lockTaken = False
5
6  def thread(self):
7      while choose({ False, True }):
8          # Enter critical section
9          await not lockTaken
10         lockTaken = True
11
12         atomically in_cs += 1
13         # Critical section
14         atomically in_cs -= 1
15
16         # Leave critical section
17         lockTaken = False
18
19  spawn thread(0)
20  spawn thread(1)
```

First attempt: a naïve lock

```
1  in_cs = 0
2  invariant in_cs in { 0, 1 }
3
4  lockTaken = False
5
6  def thread(self):
7      while choose({ False, True }):
8          # Enter critical section
9          await not lockTaken
10         lockTaken = True
11
12         atomically in_cs += 1
13         # Critical section
14         atomically in_cs -= 1
15
16         # Leave critical section
17         lockTaken = False
18
19  spawn thread(0)
20  spawn thread(1)
```



wait till lock is free, then take it

First attempt: a naïve lock

```
1  in_cs = 0
2  invariant in_cs in { 0, 1 }
3
4  lockTaken = False
5
6  def thread(self):
7      while choose({ False, True }):
8          # Enter critical section
9          await not lockTaken
10         lockTaken = True
11
12         atomically in_cs += 1
13         # Critical section
14         atomically in_cs -= 1
15
16         # Leave critical section
17         lockTaken = False
18
19  spawn thread(0)
20  spawn thread(1)
```

- Schedule thread T0: `init()`
 - Line 1: Initialize `in_cs` to 0
 - Line 4: Initialize `lockTaken` to False
 - **Thread terminated**
- Schedule thread T3: `thread(1)`
 - Line 7: Choose True
 - Preempted in `thread(1)` about to store True into `lockTaken` in line 10
- Schedule thread T2: `thread(0)`
 - Line 7: Choose True
 - Line 10: Set `lockTaken` to True (was False)
 - Line 12: Set `in_cs` to 1 (was 0)
 - Preempted in `thread(0)` about to execute atomic section in line 14
- Schedule thread T3: `thread(1)`
 - Line 10: Set `lockTaken` to True (unchanged)
 - Line 12: Set `in_cs` to 2 (was 1)
 - Preempted in `thread(1)` about to execute atomic section in line 14
- Schedule thread T1: `invariant()`
 - Line 2: Harmony assertion failed

Second attempt: *flags*

```
1 in_cs = 0
2 invariant in_cs in { 0, 1 }
3
4 flags = [ False, False ]
5
6 def thread(self):
7     while choose({ False, True }):
8         # Enter critical section
9         flags[self] = True
10        await not flags[1 - self]
11
12        atomically in_cs += 1
13        # Critical section
14        atomically in_cs -= 1
15
16        # Leave critical section
17        flags[self] = False
18
19 spawn thread(0)
20 spawn thread(1)
```

Second attempt: *flags*

```
1 in_cs = 0
2 invariant in_cs in { 0, 1 }
3
4 flags = [ False, False ]
5
6 def thread(self):
7     while choose({ False, True }):
8         # Enter critical section
9         flags[self] = True
10        await not flags[1 - self]
11
12        atomically in_cs += 1
13        # Critical section
14        atomically in_cs -= 1
15
16        # Leave critical section
17        flags[self] = False
18
19 spawn thread(0)
20 spawn thread(1)
```



Second attempt: *flags*

```
1 in_cs = 0
2 invariant in_cs in { 0, 1 }
3
4 flags = [ False, False ]
5
6 def thread(self):
7     while choose({ False, True }):
8         # Enter critical section
9         flags[self] = True
10        await not flags[1 - self]
11
12        atomically in_cs += 1
13        # Critical section
14        atomically in_cs -= 1
15
16        # Leave critical section
17        flags[self] = False
18
19 spawn thread(0)
20 spawn thread(1)
```



show intent to enter critical section



wait until there's no one else

Second attempt: *flags*

```
1 in_cs = 0
2 invariant in_cs in { 0, 1 }
3
4 flags = [ False, False ]
5
6 def thread(self):
7     while choose({ False, True }):
8         # Enter critical section
9         flags[self] = True
10        await not flags[1 - self]
11
12        atomically in_cs += 1
13        # Critical section
14        atomically in_cs -= 1
15
16        # Leave critical section
17        flags[self] = False
18
19 spawn thread(0)
20 spawn thread(1)
```

Summary: some execution cannot terminate

Here is a summary of an execution that exhibits the issue:

- Schedule thread T0: **init()**
 - Line 1: Initialize in_cs to 0
 - Line 4: Initialize flags to [False, False]
 - **Thread terminated**
- Schedule thread T1: thread(0)
 - Line 7: Choose True
 - Line 9: Set flags[0] to True (was False)
 - Preempted in thread(0) about to load variable flags[1] in line 10
- Schedule thread T2: thread(1)
 - Line 7: Choose True
 - Line 9: Set flags[1] to True (was False)
 - Preempted in thread(1) about to load variable flags[0] in line 10

Final state (all threads have terminated or are blocked):

- Threads:
 - T1: (blocked) thread(0)
 - about to load variable flags[1] in line 10
 - T2: (blocked) thread(1)
 - about to load variable flags[0] in line 10

Third attempt: *turn* variable

```
1  in_cs = 0
2  invariant in_cs in { 0, 1 }
3
4  turn = 0
5
6  def thread(self):
7      while choose({ False, True }):
8          # Enter critical section
9          turn = 1 - self
10         await turn == self
11
12         atomically in_cs += 1
13         # Critical section
14         atomically in_cs -= 1
15
16         # Leave critical section
17
18     spawn thread(0)
19     spawn thread(1)
```

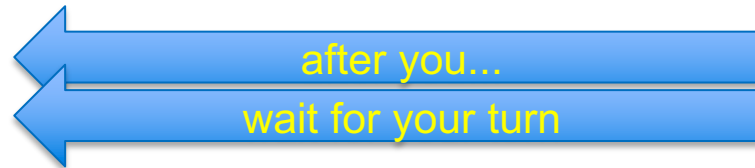
Third attempt: *turn* variable

```
1  in_cs = 0
2  invariant in_cs in { 0, 1 }
3
4  turn = 0
5
6  def thread(self):
7      while choose({ False, True }):
8          # Enter critical section
9          turn = 1 - self
10         await turn == self
11
12         atomically in_cs += 1
13         # Critical section
14         atomically in_cs -= 1
15
16         # Leave critical section
17
18     spawn thread(0)
19     spawn thread(1)
```



Third attempt: *turn* variable

```
1  in_cs = 0
2  invariant in_cs in { 0, 1 }
3
4  turn = 0
5
6  def thread(self):
7      while choose({ False, True }):
8          # Enter critical section
9          turn = 1 - self
10         await turn == self
11
12         atomically in_cs += 1
13         # Critical section
14         atomically in_cs -= 1
15
16         # Leave critical section
17
18     spawn thread(0)
19     spawn thread(1)
```



Third attempt: *turn* variable

```
1 in_cs = 0
2 invariant in_cs in { 0, 1 }
3
4 turn = 0
5
6 def thread(self):
7     while choose({ False, True }):
8         # Enter critical section
9         turn = 1 - self
10        await turn == self
11
12        atomically in_cs += 1
13        # Critical section
14        atomically in_cs -= 1
15
16        # Leave critical section
17
18 spawn thread(0)
19 spawn thread(1)
```

Summary: some execution cannot terminate

Here is a summary of an execution that exhibits the issue:

- Schedule thread T0: **init()**
 - Line 1: Initialize `in_cs` to 0
 - Line 4: Initialize `turn` to 0
 - **Thread terminated**
- Schedule thread T2: `thread(1)`
 - Line 7: Choose False
 - **Thread terminated**
- Schedule thread T1: `thread(0)`
 - Line 7: Choose True
 - Line 9: Set `turn` to 1 (was 0)
 - Preempted in `thread(0)` about to load variable `turn` in line 10

Final state (all threads have terminated or are blocked):

- Threads:
 - T1: (blocked) `thread(0)`
 - about to load variable `turn` in line 10
 - T2: (terminated) `thread(1)`

Peterson's Algorithm: *flags & turn*

```
1  in_cs = 0
2  invariant in_cs in { 0, 1 }
3
4  sequential flags, turn
5  flags = [ False, False ]
6  turn = choose({0, 1})
7
8  def thread(self):
9      while choose({ False, True }):
10         # Enter critical section
11         flags[self] = True
12         turn = 1 - self
13         await (not flags[1 - self]) or (turn == self)
14
15         atomically in_cs += 1
16         # Critical section
17         atomically in_cs -= 1
18
19         # Leave critical section
20         flags[self] = False
21
22  spawn thread(0)
23  spawn thread(1)
```

Peterson's Algorithm: *flags & turn*

```
1  in_cs = 0
2  invariant in_cs in { 0, 1 }
3
4  sequential flags, turn
5  flags = [ False, False ]
6  turn = choose({0, 1})
7
8  def thread(self):
9      while choose({ False, True }):
10         # Enter critical section
11         flags[self] = True
12         turn = 1 - self
13         await (not flags[1 - self]) or (turn == self)
14
15         atomically in_cs += 1
16         # Critical section
17         atomically in_cs -= 1
18
19         # Leave critical section
20         flags[self] = False
21
22  spawn thread(0)
23  spawn thread(1)
```



in critical section

Peterson's Algorithm: *flags & turn*

```
1 in_cs = 0
2 invariant in_cs in { 0, 1 }
3
4 sequential flags, turn
5 flags = [ False, False ]
6 turn = choose({0, 1})
7
8 def thread(self):
9     while choose({ False, True }):
10         # Enter critical section
11         flags[self] = True
12         turn = 1 - self
13         await (not flags[1 - self]) or (turn == self)
14
15         atomically in_cs += 1
16         # Critical section
17         atomically in_cs -= 1
18
19         # Leave critical section
20         flags[self] = False
21
22 spawn thread(0)
23 spawn thread(1)
```

load and store instructions are atomic

in critical section

Peterson's Algorithm: *flags & turn*

```
1 in_cs = 0
2 invariant in_cs in { 0, 1 }
3
4 sequential flags, turn
5 flags = [ False, False ]
6 turn = choose({0, 1})
7
8 def thread(self):
9     while choose({ False, True }):
10         # Enter critical section
11         flags[self] = True
12         turn = 1 - self
13         await (not flags[1 - self]) or (turn == self)
14
15         atomically in_cs += 1
16         # Critical section
17         atomically in_cs -= 1
18
19         # Leave critical section
20         flags[self] = False
21
22 spawn thread(0)
23 spawn thread(1)
```

load and store instructions are atomic

uses **flags** and **turn** variable (3 bits total)

in critical section

Peterson's Algorithm: *flags & turn*

```
1 in_cs = 0
2 invariant in_cs in { 0, 1 }
3
4 sequential flags, turn ← load and store instructions are atomic
5 flags = [ False, False ] ← uses flags and turn variable (3 bits total)
6 turn = choose({0, 1})
7
8 def thread(self):
9     while choose({ False, True }):
10        # Enter critical section
11        flags[self] = True ← first indicate intention to enter critical section
12        turn = 1 - self
13        await (not flags[1 - self]) or (turn == self)
14
15        atomically in_cs += 1 ← in critical section
16        # Critical section
17        atomically in_cs -= 1
18
19        # Leave critical section
20        flags[self] = False ← no longer in critical section
21
22 spawn thread(0)
23 spawn thread(1)
```

Peterson's Algorithm: *flags & turn*

```
1 in_cs = 0
2 invariant in_cs in { 0, 1 }
3
4 sequential flags, turn
5 flags = [ False, False ]
6 turn = choose({0, 1})
```

load and store instructions are atomic

uses **flags** and **turn** variable (3 bits total)

```
8 def thread(self):
9     while choose({ False, True }):
10        # Enter critical section
11        flags[self] = True
12        turn = 1 - self
13        await (not flags[1 - self]) or (turn == self)
```

first indicate intention to enter critical section

also give other thread a turn first

```
15 atomically in_cs += 1
16 # Critical section
17 atomically in_cs -= 1
```

in critical section

```
18
19 # Leave critical section
20 flags[self] = False
```

no longer in critical section

```
22 spawn thread(0)
23 spawn thread(1)
```

Peterson's Algorithm: *flags & turn*

```
1 in_cs = 0
2 invariant in_cs in { 0, 1 }
3
4 sequential flags, turn
5 flags = [ False, False ]
6 turn = choose({0, 1})
7
8 def thread(self):
9     while choose({ False, True }):
10        # Enter critical section
11        flags[self] = True
12        turn = 1 - self
13        await (not flags[1 - self]) or (turn == self)
14
15        atomically in_cs += 1
16        # Critical section
17        atomically in_cs -= 1
18
19        # Leave critical section
20        flags[self] = False
21
22 spawn thread(0)
23 spawn thread(1)
```

load and store instructions are atomic

uses **flags** and **turn** variable (3 bits total)

first indicate intention to enter critical section

also give other thread a turn first

wait for one of either conditions

in critical section

no longer in critical section

So, we proved Peterson's Algorithm correct by brute force, enumerating all possible executions. We now know *that* it works.

*But how does one prove it by deduction?
so one understands *why* it works...*

What and how?

- Need to show that, for any execution, all states reached satisfy mutual exclusion
 - in other words, mutual exclusion is *invariant*
invariant = predicate that holds in every reachable state

What is an invariant?

A property that holds in all reachable states

(and possibly in some unreachable states as well)

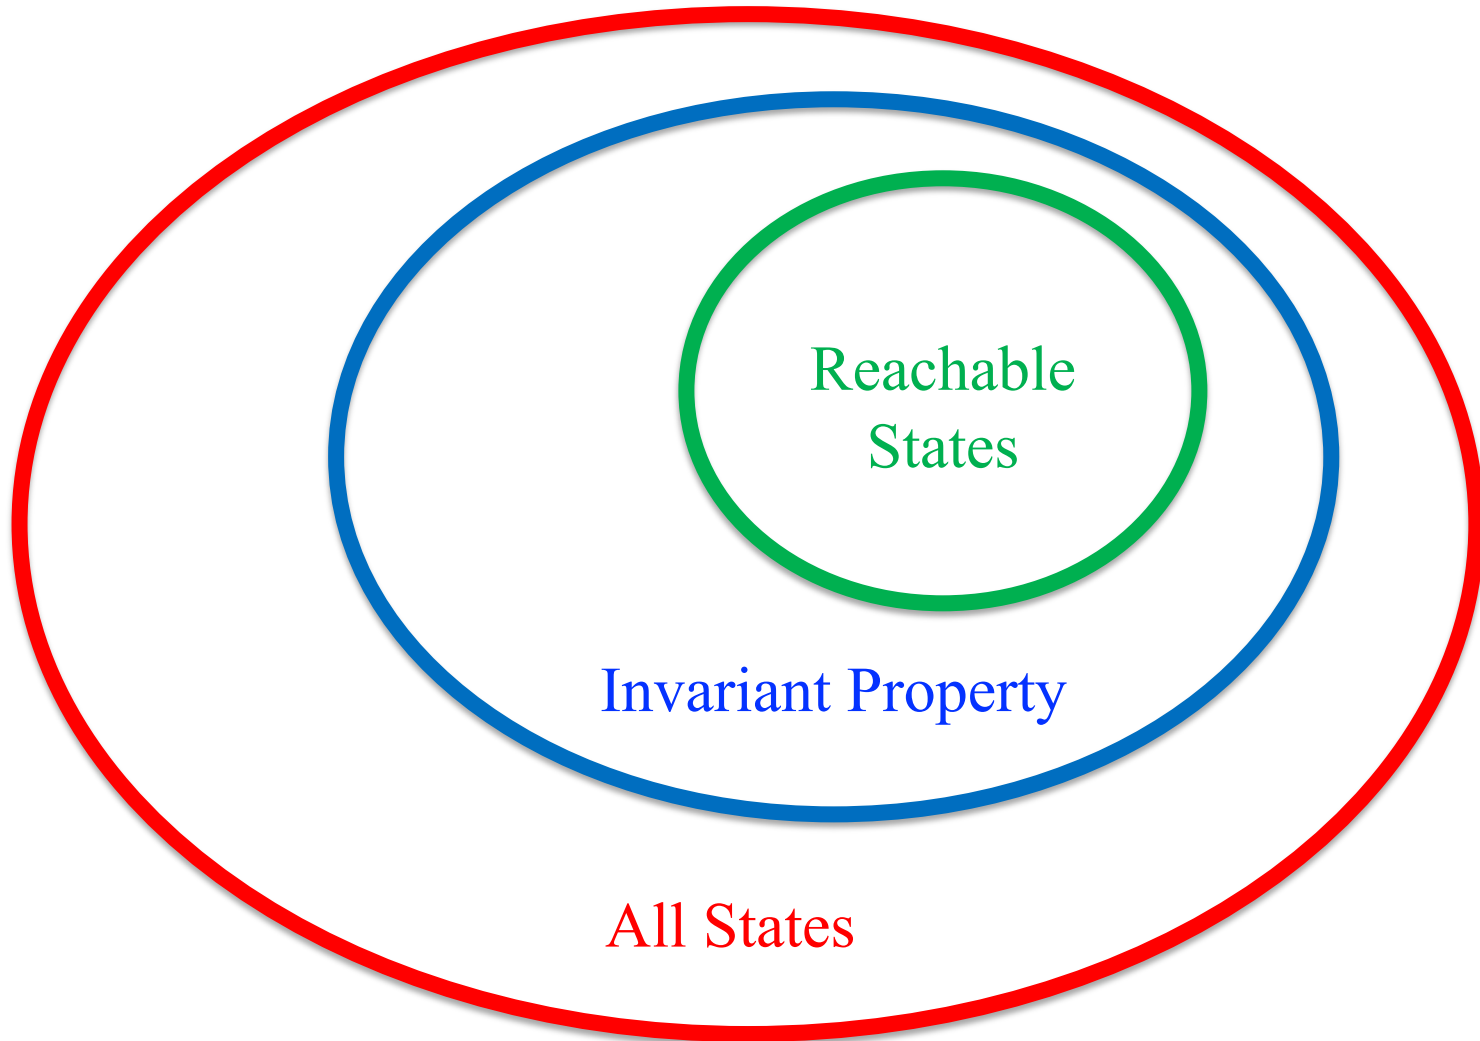
What is a property?

A property is a set of states

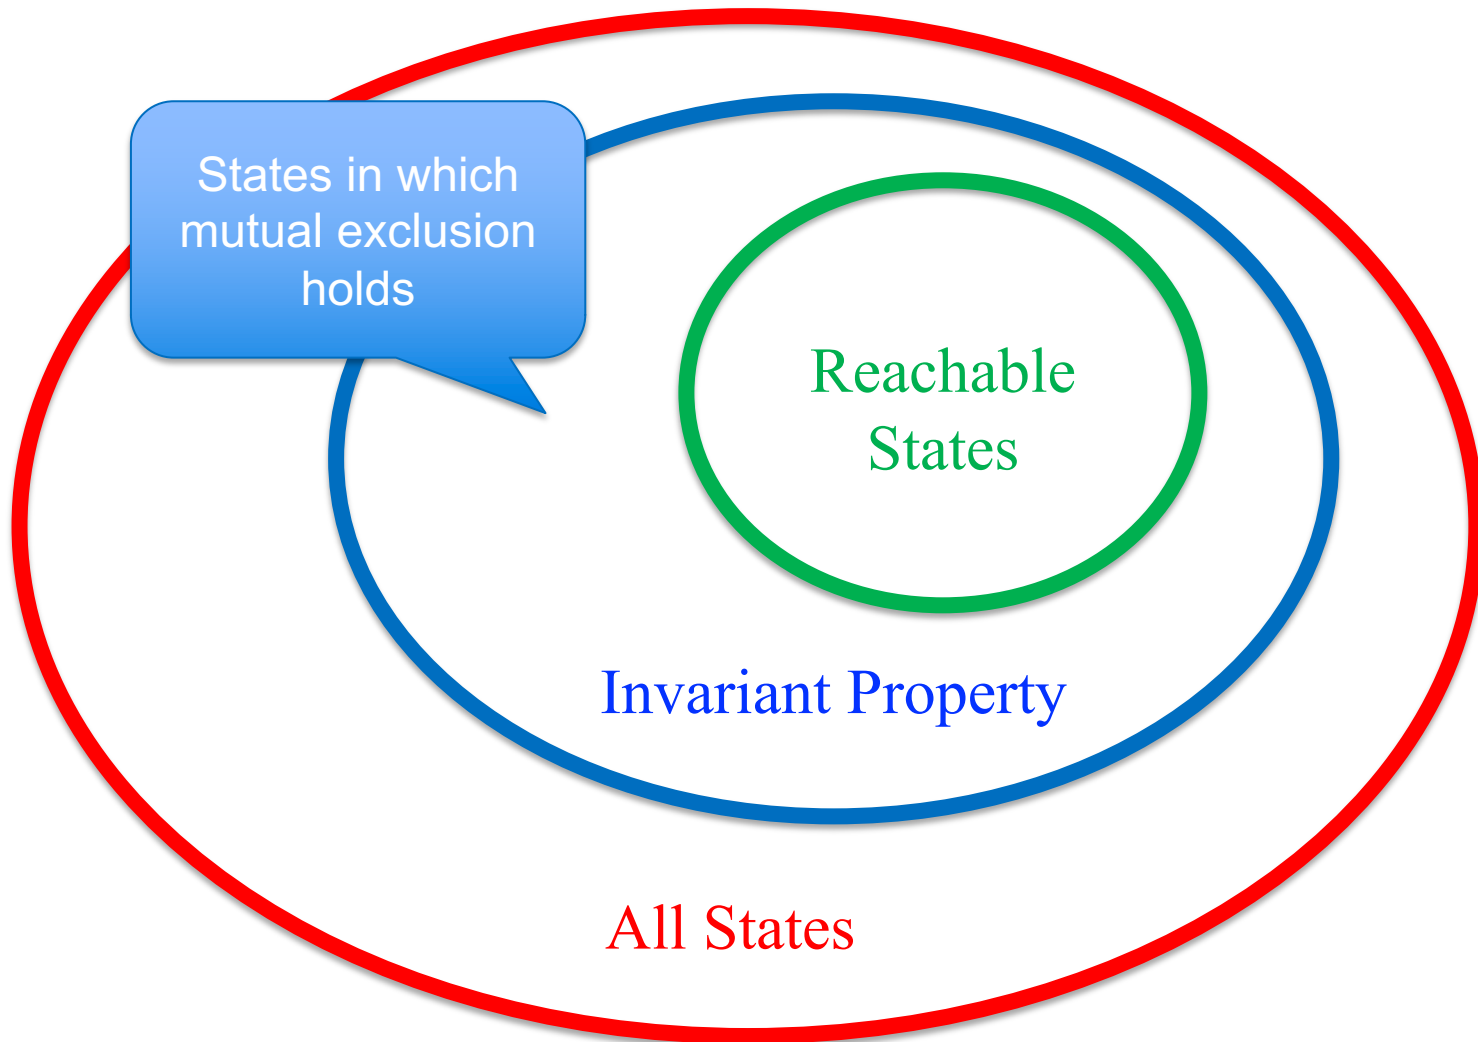
often succinctly described using a predicate

(all states that satisfy the predicate and no others)

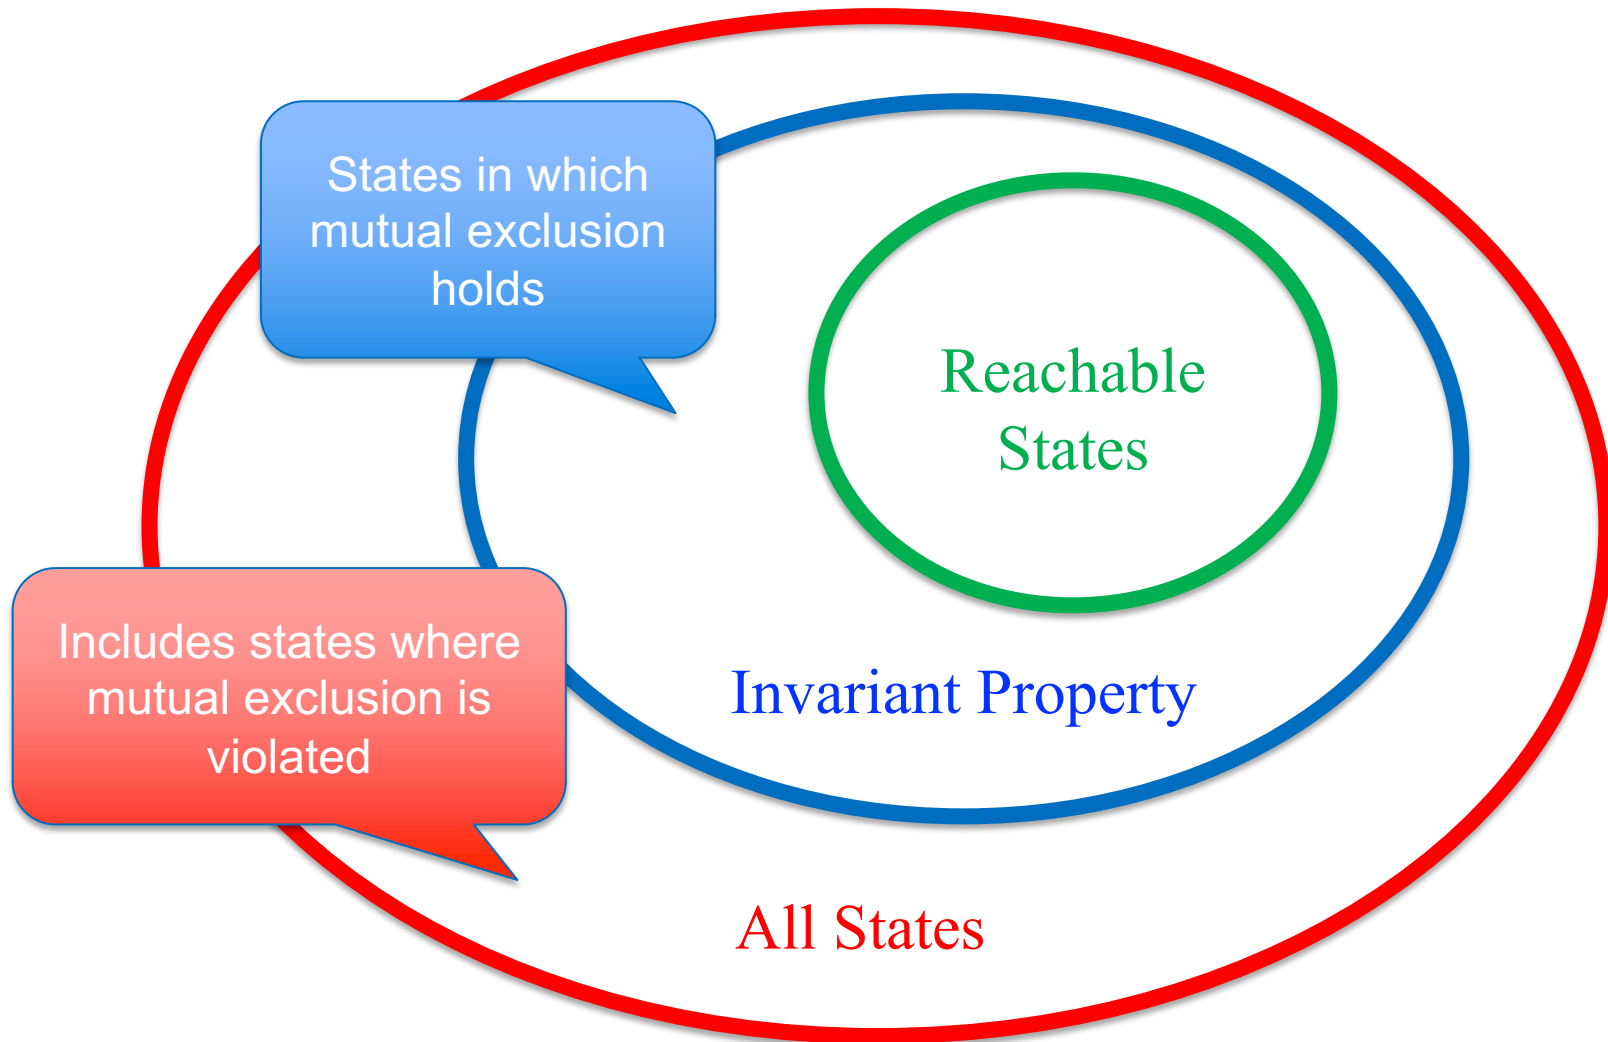
Invariant Property



Invariant Property



Invariant Property



How to prove an invariant?

- Need to show that, for any execution, all states reached satisfy the invariant
- Sounds similar to sorting:
 - Need to show that, for any list of numbers, the resulting list is ordered
- Let's try *proof by induction* on the length of an execution

Proof by induction

You want to prove that some *Induction Hypothesis* $IH(n)$ holds for any n :

- Base Case:

- show that $IH(0)$ holds

- Induction Step:

- show that if $IH(i)$ holds, then so does $IH(i+1)$

Proof by induction in our case

To show that some **IH** holds for an *execution* **E** of any number of *steps*:

- **Base Case:**

- show that **IH** holds in the initial state(s)

- **Induction Step:**

- show that if **IH** holds in a state produced by **E**, then for any possible next step **s**, **IH** also holds in the state produced by **E + [s]**

But there's a problem

- It turns out that mutual exclusion is hard to prove directly
 - it's hard to show that, if mutual exclusion holds in a state, it will also hold in the next state after making one execution step
 - not a good basis for induction
- Need a stronger invariant that implies mutual exclusion

Peterson's Algorithm: *flags & turn*

```
1  in_cs = 0
2  invariant in_cs in { 0, 1 }
3
4  sequential flags, turn
5  flags = [ False, False ]
6  turn = choose({0, 1})
7
8  def thread(self):
9      while choose({ False, True }):
10         # Enter critical section
11         flags[self] = True
12         turn = 1 - self
13         await (not flags[1 - self]) or (turn == self)
14
15         atomically in_cs += 1
16         # Critical section
17         atomically in_cs -= 1
18
19         # Leave critical section
20         flags[self] = False
21
22  spawn thread(0)
23  spawn thread(1)
```



Candidate invariant to prove

Peterson's Reconsidered

- Assumes that LOAD and STORE instructions are *atomic*
- Not guaranteed on a real processor
- Also not guaranteed by C, Java, Python,

...

```
1 in_cs = 0
2 invariant in_cs in { 0, 1 }
3
4 sequential flags, turn ← Loads and Stores are atomic
5 flags = [ False, False ]
6 turn = choose({0, 1})
7
8 def thread(self):
9     while choose({ False, True }):
10        # Enter critical section
11        flags[self] = True
12        turn = 1 - self
13        await (not flags[1 - self]) or (turn == self)
14
15        atomically in_cs += 1
16        # Critical section
17        atomically in_cs -= 1
18
19        # Leave critical section
20        flags[self] = False
21
22 spawn thread(0)
23 spawn thread(1)
```

Non-atomic load/store example

- Suppose x is a 64-bit integer
- Suppose you have a 32-bit CPU
- Then " $x = 0$ " requires 2 stores
 - because x occupies 2 words
- Similarly, reading x requires 2 loads
- Same is true if x is a 32-bit integer but x is not aligned on a word boundary
 - Writing to x would require two LOAD and two STORE operations on memory!

Concurrent writing

- Hardware may also cause problems
 - e.g., buffering of writes to memory for improved performance
- Because of all these issues, programming languages will typically leave the outcome of concurrent operations to a variable *undefined*
 - if at least one of those operations is a store

Data Race

- When two threads access the same variable
- And at least one is a STORE
- Then the semantics of the outcome is *undefined*

Harmony “sequential” statement

- **sequential** *turn, flags*
- ensures that loads/stores are atomic
- that is, concurrent operations appear to be executed sequentially
- This is called “*sequential consistency*”

For example

- Shared variable x contains 3
- Thread A stores 4 into x
- Thread B loads x
 - With atomic load/store operations, B will read either 3 or 4
 - With normal operations, the value that B reads is undefined

Sequential consistency

- Java has a similar notion:
 - **volatile int** *x*;
- Not to be confused with the same keyword in C and C++ though...
- Loading/storing volatile (sequentially consistent) variables is *more expensive* than loading/storing ordinary variables
 - because it restricts CPU and/or compiler optimizations

Peterson's Reconsidered Again

- Mutual Exclusion can be implemented with atomic LOAD and STORE instructions to access shared memory
 - hardware supports such instructions but they are very expensive
- Peterson's can be generalized to >2 processes
 - even more STOREs and LOADs

Too inefficient in practice

Enter *Interlock Instructions*

- Machine instructions that do multiple shared memory accesses atomically
- e.g., **TestAndSet s**
 - sets **s** to **True**
 - returns old value of **s**
- i.e., does the following:
 - LOAD r0, s # load variable s into register r0
 - STORE s, 1 # store TRUE in variable s
- Entire operation is *atomic*
 - other machine instructions cannot interleave

Harmony interlude: *pointers*

- If x is a shared variable, $?x$ is the address of x
- If p is a variable and p contains $?x$, then we say that p is a *pointer* to x
- Finally, $!p$ refers to the value of x

Harmony interlude: *pointers*

- If x is a shared variable, $?x$ is the address of x
- If p is a variable and p contains $?x$, then we say that p is a *pointer* to x
- Finally, $!p$ refers to the value of x



Where?
There!

Specifying a lock

```
1  def Lock() returns result:  
2      result = False  
3  
4  def acquire(lk):  
5      atomically when not !lk:  
6          !lk = True  
7  
8  def release(lk):  
9      atomically:  
10         assert !lk  
11         !lk = False
```

Specifying a lock

```
1 def Lock() returns result:  
2     result = False  
3  
4 def acquire(lk):  
5     atomically when not !lk:  
6         !lk = True  
7  
8 def release(lk):  
9     atomically:  
10        assert !lk  
11        !lk = False
```



returns initial value



acquires lock atomically once available



releases lock atomically

Critical Section using a *lock*

```
1  import lock
2
3  const NTHREADS = 5
4
5  in_cs = 0
6  invariant in_cs in { 0, 1 }
7
8  thelock = lock.Lock()
9
10 def thread():
11     while choose({ False, True }):
12         lock.acquire(?thelock)
13
14         atomically in_cs += 1
15         # Critical section
16         atomically in_cs -= 1
17
18         lock.release(?thelock)
19
20 for i in {1..NTHREADS}:
21     spawn thread()
```

“Ghost” state

- We say that a lock is *held* or *owned* by a thread
 - implicit “ghost” state (not an actual variable)
 - nonetheless can be used for reasoning
- Two important invariants:
 1. $T@CriticalSection \Rightarrow T$ holds the lock
 2. at most one thread can hold the lock

Many (most?) systems do not keep track of who holds a particular lock, if anybody

Lock implementation (“spinlock”)

```
1 def test_and_set(s) returns result:  
2     atomically:  
3         result = !s  
4         !s = True  
5  
6 def atomic_store(p, v):  
7     atomically !p = v  
8  
9 def Lock() returns result:  
10     result = False  
11  
12 def acquire(lk):  
13     while test_and_set(lk):  
14         pass  
15  
16 def release(lk):  
17     atomic_store(lk, False)
```

specification of the CPU's
test_and_set functionality

specification of the CPU's
atomic store functionality

lock implementation

Specification vs Implementation

```
1 def Lock() returns result:  
2     result = False  
3  
4 def acquire(lk):  
5     atomically when not !lk:  
6         !lk = True  
7  
8 def release(lk):  
9     atomically:  
10         assert !lk  
11         !lk = False
```

```
1 def test_and_set(s) returns result:  
2     atomically:  
3         result = !s  
4         !s = True  
5  
6 def atomic_store(p, v):  
7     atomically !p = v  
8  
9 def Lock() returns result:  
10     result = False  
11  
12 def acquire(lk):  
13     while test_and_set(lk):  
14         pass  
15  
16 def release(lk):  
17     atomic_store(lk, False)
```

Specification: describes *what an abstraction does*

Implementation: describes *how*



Spinlocks and Time Sharing

- Spinlocks work well when threads on different cores need to synchronize
- But how about when it involves two threads time-shared on the same core:
 - when there is no pre-emption?
 - when there is pre-emption?

Spinlocks and Time Sharing

- Spinlocks work well when threads on different cores need to synchronize
- But how about when it involves two threads time-shared on the same core:
 - when there is no pre-emption?
 - can cause all threads to get stuck while one is trying to obtain a spinlock
 - when there is pre-emption?

Spinlocks and Time Sharing

- Spinlocks work well when threads on different cores need to synchronize
- But how about when it involves two threads time-shared on the same core:
 - when there is no pre-emption?
 - can cause all threads to get stuck while one is trying to obtain a spinlock
 - when there is pre-emption?
 - can cause delays and waste of CPU cycles while a thread is trying to obtain a spinlock

Context switching in Harmony

- Harmony allows contexts to be saved and restored (i.e., **context switch**)

- ***r = stop p***

- stops the current thread and stores context in *!p*

- ***go (!p) r***

- adds a thread with the given context to the bag of threads. Thread resumes from **stop** expression, returning *r*

Locks using **stop** and **go**

```
1 def Lock() returns result:
2   result = { .acquired: False, .suspended: [] }
3
4 def acquire(lk):
5   atomically:
6     if lk->acquired:
7       stop ?lk->suspended[len lk->suspended]
8       assert lk->acquired
9     else:
10      lk->acquired = True
11
12 def release(lk):
13   atomically:
14     assert lk->acquired
15     if lk->suspended == []:
16       lk->acquired = False
17     else:
18       go (lk->suspended[0]) ()
19     del lk->suspended[0]
```

.acquired: boolean
.suspended: queue of contexts

Locks using **stop** and **go**

```
1 def Lock() returns result:
2   result = { .acquired: False, .suspended: [] }
3
4 def acquire(lk):
5   atomically:
6     if lk->acquired:
7       stop ?lk->suspended[ len lk->suspended ]
8       assert lk->acquired
9     else:
10      lk->acquired = True
11
12 def release(lk):
13   atomically:
14     assert lk->acquired
15     if lk->suspended == []:
16       lk->acquired = False
17     else:
18       go (lk->suspended[0]) ()
19     del lk->suspended[0]
```

.acquired: boolean
.suspended: queue of contexts

← put thread on wait queue

← resume first thread on wait queue

Locks using **stop** and **go**

```
1 def Lock() returns result:  
2     result = { .acquired: False, .suspended: [] }  
3  
4 def acquire(lk):  
5     atomically:
```

Similar to a Linux “futex”: if there is no contention (hopefully the common case) acquire() and release() are cheap. If there is contention, they involve a context switch.

```
13     atomically:  
14         assert lk->acquired  
15         if lk->suspended == []:  
16             lk->acquired = False  
17         else:  
18             go (lk->suspended[0]) ()  
19             del lk->suspended[0]
```

Choosing modules in Harmony

- “synch” is the (default) module that has the specification of a lock
- “synchS” is the module that has the **stop/go** version of lock
- you can select which one you want:

`harmony -m synch=synchS x.hny`

- “synch” tends to be faster than “synchS”
 - smaller state graph

Atomic section \neq Critical Section

Atomic Section	Critical Section
only one thread can execute	multiple threads can execute concurrently, just not within a critical section
rare programming language paradigm	ubiquitous: locks available in many mainstream programming languages
good for specifying interlock instructions	good for implementing concurrent data structures

Harmony demo:

Demo 1:
data race

```
x = 0

def f():
    x = x + 1

def g():
    x = x + 1

spawn f()
spawn g()
```

Demo 2: no data race

```
x = 0

def atomic_load(p) returns v:
    atomically v = !p

def atomic_store(p, v):
    atomically !p = v

def f():
    atomic_store(?x, atomic_load(?x) + 1)

def g():
    atomic_store(?x, atomic_load(?x) + 1)

spawn f()
spawn g()
```

Demo 3: same
semantics as
Demo 2:

```
sequential x

x = 0

def f():
    x = x + 1

def g():
    x = x + 1

spawn f()
spawn g()
```


Harmony demo

Demo 4: still a data race

```
x = 0

def atomic_load(p) returns v:
    atomically v = !p

def atomic_store(p, v):
    atomically !p = v

def f():
    atomic_store(?x, x + 1)

def g():
    atomic_store(?x, atomic_load(?x) + 1)

spawn f()
spawn g()
```

Demo 5: data race
freedom does not imply
no race conditions

```
sequential x
finally x == 2

x = 0

def f():
    x += 1

def g():
    x += 1

spawn f()
spawn g()
```

Harmony demo

Demo 6: spec of
what we want

```
finally x == 2

x = 0

def f():
    atomically x += 1

def g():
    atomically x += 1

spawn f()
spawn g()
```

Demo 7: implementation
using critical section

```
from synch import Lock, acquire, release

finally x == 2

x = 0
thelock = Lock()

def f():
    acquire(?thelock)
    x += 1
    release(?thelock)

def g():
    acquire(?thelock)
    x += 1
    release(?thelock)

spawn f()
spawn g()
```

Harmony demo

Demo 8: broken implementation using two critical sections

```
from synch import Lock, acquire, release

finally x == 2

x = 0
thelock1 = Lock()
thelock2 = Lock()

def f():
    acquire(?thelock1)
    x += 1
    release(?thelock1)

def g():
    acquire(?thelock2)
    x += 1
    release(?thelock2)

spawn f()
spawn g()
```

Summary

- A *Data Race* occurs when two threads try to access the same variable and at least one access is non-atomic and at least one access is an update.
 - The outcome of the operations may be undefined and almost always is a bug
- A *Race Condition* occurs when the correctness of the program depends on ordering of variable access
 - Race Condition does not imply Data Race

Summary, cont'd

- A *Critical Section* consists of one or more regions of code in which at most thread can execute at a time
 - usually protected by a *lock*
 - not the same as atomic because threads can continue to execute other regions of the code
- Beware of code with multiple critical sections
 - e.g., code that uses multiple locks

Data Structure *consistency*

- Each data structure maintains some *consistency property*
 - e.g., in a linked list, there is a head, a tail, a list of nodes such that head points to first node, tail points to the last node, and each node points to the next one except the last, which points to **None**. However, if the list is empty, head and tail are both **None**.

Consistency using locks

- Each data structure maintains some *consistency property*
 - e.g., in a linked list, there is a head, a tail, a list of nodes such that head points to first node, tail points to the last node, and each node points to the next one except the last, which points to **None**. However, if the list is empty, head and tail are both **None**.
- *You can assume the property holds right after obtaining the lock*
- *You must make sure the property holds again right before releasing the lock*

Consistency using locks

- Each data structure maintains some *consistency property*
- *Invariant:*
 - lock not held \implies data structure consistent
- *Or equivalently:*
 - data structure inconsistent \implies lock held

Building a Concurrent Queue

- `q = queue.Queue()`: initialize a new queue
- `queue.put(q, v)`: add `v` to the tail of queue `q`
- `v = queue.get(q)`: returns `None` if `q` is empty or `v` if `v` was at the head of the queue

Specifying a concurrent queue

```
1 def Queue() returns empty:  
2   empty = []  
3  
4 def put(q, v):  
5   !q += [v,]  
6  
7 def get(q) returns next:  
8   if !q == []:  
9     next = None  
10  else:  
11    next = (!q)[0]  
12    del (!q)[0]
```

(a) [[code/queue_nonatom.hny](#)] Sequential

```
1 def Queue() returns empty:  
2   empty = []  
3  
4 def put(q, v):  
5   atomically !q += [v,]  
6  
7 def get(q) returns next:  
8   atomically:  
9     if !q == []:  
10      next = None  
11   else:  
12     next = (!q)[0]  
13     del (!q)[0]
```

(b) [[code/queue.hny](#)] Concurrent

Example of using a queue

```
1  import queue
2
3  def sender(q, v):
4      queue.put(q, v)  enqueue v onto !q
5
6  def receiver(q):
7      let v = queue.get(q):  dequeue and check
8          assert v in { None, 1, 2 }
9
10 demoq = queue.Queue()  create queue
11 spawn sender(?demoq, 1)
12 spawn sender(?demoq, 2)
13 spawn receiver(?demoq)
14 spawn receiver(?demoq)
```


Specifying a concurrent queue

```
1 def Queue() returns empty:  
2     empty = []  
3  
4 def put(q, v):  
5     atomically !q += [v,]  
6  
7 def get(q) returns next:  
8     atomically:  
9         if !q == []:  
10            next = None  
11        else:  
12            next = (!q)[0]  
13            del (!q)[0]
```

Not a good implementation because


- *operations are $O(n)$*
- *code uses **atomically***

Queue implementation, v1



```
1 from synch import Lock, acquire, release
2 from alloc import malloc, free
3
4 def Queue() returns empty:
5     empty = { .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)
```


Queue implementation, v1



```
1 from synch import Lock, acquire, release
2 from alloc import malloc, free
3
4 def Queue() returns empty:
5     empty = { .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)
```

dynamic memory allocation


Queue implementation, v1



```
1 from synch import Lock, acquire, release
2 from alloc import malloc, free
3
4 def Queue() returns empty:
5     empty = { .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)
```

empty queue


Queue implementation, v1




```
1 from synch import Lock, acquire, release
2 from alloc import malloc, free
3
4 def Queue() returns empty:
5     empty = { .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)
```

allocate node


Queue implementation, v1



```
1 from synch import Lock, acquire, release
2 from alloc import malloc, free
3
4 def Queue() returns empty:
5     empty = { .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)
```




Queue implementation, v1



```
1 from synch import Lock, acquire, release
2 from alloc import malloc, free
3
4 def Queue() returns empty:
5     empty = { .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)
```

the hard stuff


Queue implementation, v1



```
1 from synch import Lock, acquire, release
2 from alloc import malloc, free
3
4 def Queue() returns empty:
5     empty = { .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)
```


← release lock

Queue implementation, v1



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


Queue implementation, v1



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

the hard stuff

Queue implementation, v1



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

malloc'd memory must be explicitly released (cf. C)

How important are concurrent queues?


- Answer: **all important**
 - any resource that needs scheduling
 - CPU run queue
 - disk, network, printer waiting queue
 - lock waiting queue
 - inter-process communication
 - Posix pipes:
 - **cat file | tr a-z A-Z | grep RVR**
 - actor-based concurrency
 - ...

How important are concurrent queues?

- Answer: **all important**
 - any resource that needs scheduling
 - CPU run queue
 - disk, network, printer waiting queue
 - lock waiting queue
 - inter-process communication
 - Posix pipes:
 - **cat file | tr a-z A-Z | grep RVR**
 - actor-based concurrency
 - ...


Good performance is critical!

Concurrent queue v2: 2 locks




```
1 from synch import Lock, acquire, release, atomic_load, atomic_store
2 from alloc import malloc, free
3
4 def Queue() returns empty:
5     let dummy = malloc({ .value: (), .next: None }):
6         empty = { .head: dummy, .tail: dummy, .hdlock: Lock(), .tllock: Lock() }
7
8 def put(q, v):
9     let node = malloc({ .value: v, .next: None }):
10         acquire(?q→tllock)
11         atomic_store(?q→tail→next, node)
12         q→tail = node
13         release(?q→tllock)
```

Concurrent queue v2: 2 locks




```
1 from synch import Lock, acquire, release, atomic_load, atomic_store
2 from alloc import malloc, free
3
4 def Queue() returns empty:
5     let dummy = malloc({ .value: (), .next: None }):
6         empty = { .head: dummy, .tail: dummy, .hdlock: Lock(), .tllock: Lock() }
7
8 def put(q, v):
9     let node = malloc({ .value: v, .next: None }):
10         acquire(?q->tllock)
11         atomic_store(?q->tail->next, node) ← atomically q->tail->next = node
12         q->tail = node
13         release(?q->tllock)
```

Concurrent queue v2: 2 locks



```
15 def get(q) returns next:
16     acquire(?q→hdlock)
17     let dummy = q→head
18     let node = atomic_load(?dummy→next):
19         if node == None:
20             next = None
21             release(?q→hdlock)
22         else:
23             next = node→value
24             q→head = node
25             release(?q→hdlock)
26             free(dummy)
```


Concurrent queue v2: 2 locks



```
15 def get(q) returns next:
16   acquire(?q→hdlock)
17   let dummy = q→head
18   let node = atomic_load(?dummy→next):
19     if node == None:
20       next = None
21       release(?q→hdlock)
22     else:
23       next = node→value
24       q→head = node
25       release(?q→hdlock)
26       free(dummy)
```

No contention for concurrent
enqueue and dequeue operations!
→ more concurrency → faster

Concurrent queue v2: 2 locks



```
15 def get(q) returns next:
16   acquire(?q→hdlock)
17   let dummy = q→head
18   let node = atomic_load(?dummy→next):
19     if node == None:
20       next = None
21       release(?q→hdlock)
22     else:
23       next = node→value
24       q→head = node
25       release(?q→hdlock)
26       free(dummy)
```


No contention for concurrent enqueue and dequeue operations!
→ more concurrency → faster

Needs to avoid data race on `dummy`→`next` when queue is empty

Global vs. Local Locks

- The two-lock queue is an example of a data structure with *finer-grained locking*
- A global lock is easy, but limits concurrency
- Fine-grained or local locking can improve concurrency, but tends to be trickier to get right

Sorted Linked List with Lock per Node




```
1 from synch import Lock, acquire, release
2 from alloc import malloc, free
3
4 def _node(v, n): # allocate and initialize a new list node
5     result = malloc({ .lock: Lock(), .value: v, .next: n })
6
7 def _find(lst, v):
8     var before = lst
9     acquire(?before→lock)
10    var after = before→next
11    acquire(?after→lock)
12    while after→value < (0, v):
13        release(?before→lock)
14        before = after
15        after = before→next
16        acquire(?after→lock)
17    result = (before, after)
18
19 def SetObject():
20    result = _node((-1, None), _node((1, None), None))
```

- $-\infty$ represented by $(-1, \text{None})$
 - v represented by $(0, v)$
 - ∞ represented by $(1, \text{None})$
- Note that $\forall v: (-1, \text{None}) < (0, v) < (1, \text{None})$
(lexicographical ordering)




Sorted Linked List with Lock per Node



```
1 from synch import Lock, acquire, release
2 from alloc import malloc, free
3
4 def _node(v, n): # allocate and initialize a new list node
5     result = malloc({ .lock: Lock(), .value: v, .next: n })
6
7 def _find(lst, v):
8     var before = lst
9     acquire(?before→lock)
10    var after = before→next
11    acquire(?after→lock)
12    while after→value < (0, v):
13        release(?before→lock)
14        before = after
15        after = before→next
16        acquire(?after→lock)
17    result = (before, after)
18
19 def SetObject():
20    result = _node((-1, None), _node((1, None), None))
```

Helper routine to find and lock two consecutive nodes *before* and *after* such that $before \rightarrow value < v \leq after \rightarrow value$

Sorted Linked List with Lock per Node



```
1 from synch import Lock, acquire, release
2 from alloc import malloc, free
3
4 def _node(v, n): # allocate and initialize a new list node
5     result = malloc({ .lock: Lock(), .value: v, .next: n })
6
7 def _find(lst, v):
8     var before = lst
9     acquire(?before→lock)
10    var after = before→next
11    acquire(?after→lock)
12    while after→value < (0, v):
13        release(?before→lock)
14        before = after
15        after = before→next
16        acquire(?after→lock)
17    result = (before, after)
18
19 def SetObject():
20    result = _node((-1, None), _node((1, None), None))
```

Helper routine to find and lock two consecutive nodes *before* and *after* such that $before \rightarrow value < v \leq after \rightarrow value$

Hand-over hand locking
(good for data structures without cycles)

Sorted Linked List with Lock per Node

```
22  def insert(lst, v):
23      let before, after = _find(lst, v):
24          if after→value != (0, v):
25              before→next = _node((0, v), after)
26              release(?after→lock)
27              release(?before→lock)
28
29  def remove(lst, v):
30      let before, after = _find(lst, v):
31          if after→value == (0, v):
32              before→next = after→next
33              release(?after→lock)
34              free(after)
35          else:
36              release(?after→lock)
37              release(?before→lock)
38
39  def contains(lst, v):
40      let before, after = _find(lst, v):
41          result = after→value == (0, v)
42          release(?after→lock)
43          release(?before→lock)
```

Sorted Linked List with Lock per Node

```
22 def insert(lst, v):
23     let before, after = _find(lst, v):
24         if after→value != (0, v):
25             before→next = _node((0, v), after)
26             release(?after→lock)
27             release(?before→lock)
28
29 def remove(lst, v):
30     let before, after = _find(lst, v):
31         if after→value == (0, v):
32             before→next = after→next
33             release(?after→lock)
34             free(after)
35         else:
36             release(?after→lock)
37         release(?before→lock)
38
39 def contains(lst, v):
40     let before, after = _find(lst, v):
41         result = after→value == (0, v)
42         release(?after→lock)
43         release(?before→lock)
```

Multiple threads can access the list simultaneously, but they can't *overtake* one another

Systematic Testing

Systematic Testing


- Sequential case
 - try all “sequences” of 1 operation
 - put or get (in case of queue)
 - try all sequences of 2 operations
 - put+put, put+get, get+put, get+get, ...
 - try all sequences of 3 operations
 - ...
- How do you know if a sequence is correct?
 - compare “behaviors” of running test against implementation with running test against the sequential specification

Systematic Testing

- Concurrent case
 - try all “interleavings” of 1 operation
 - try all interleavings of 2 operations
 - try all interleavings of 3 operations
 - ...
- How do you know if an interleaving is correct?
 - compare “behaviors” of running test against concurrent implementation with running test against the concurrent specification

Life of an atomic operation

process invokes operation operation happens atomically process resumes with result



Concurrency and Overlap

Is the following a possible scenario?

1. customer X orders a burger
2. customer Y orders a burger (afterwards)
3. customer Y is served a burger
4. customer X is served a burger (afterwards)

Concurrency and Overlap

Is the following a possible scenario?

1. customer X orders a burger
2. customer Y orders a burger (afterwards)
3. customer Y is served a burger
4. customer X is served a burger (afterwards)


We've all seen this happen. It's a matter of how things get scheduled!

Specification

- One operation: order a burger
 - result: a burger (at some later time)
- Semantics: the burger manifests itself atomically *sometime during the operation*
- ***Atomically: no two manifestations overlap***
- It's easier to specify something when you don't have to worry about overlap
 - i.e., you can simply give a sequential specification
- Allows many implementations


Implementation?

- Suppose the diner has one small hot plate and two cooks
- Cooks use a lock for access to the hot plate
- Possible scenario:
 1. customer X orders burger, order ends up with cook 1
 2. customer Y orders burger, order ends up with cook 2
 3. cook 1 was busy with something else, so cook 2 grabs the lock first
 4. cook 2 cooks burger for Y
 5. cook 2 releases lock
 6. cook 1 grabs lock
 7. cook 1 cooks burger for X
 8. cook 1 releases lock
 9. customer Y receives burger
 10. customer X receives burger



Implementation?

- Suppose the diner has one small hot plate and two cooks
- Cooks use a lock for access to the hot plate
- Possible scenario:
 1. customer X orders burger, order ends up with cook 1
 2. customer Y orders burger, order ends up with cook 2
 3. cook 1 was busy with something else, so cook 2 grabs the lock first
 4. cook 2 cooks burger for Y
 5. cook 2 releases lock
 6. cook 1 grabs lock
 7. cook 1 cooks burger for X
 8. cook 1 releases lock
 9. customer Y receives burger
 10. customer X receives burger



- *can't happen if Y orders burger after X receives burger*
- *but if operations overlap, any ordering can happen...*

Correct Behaviors

(1)

put(1)


get() → ?

TIME

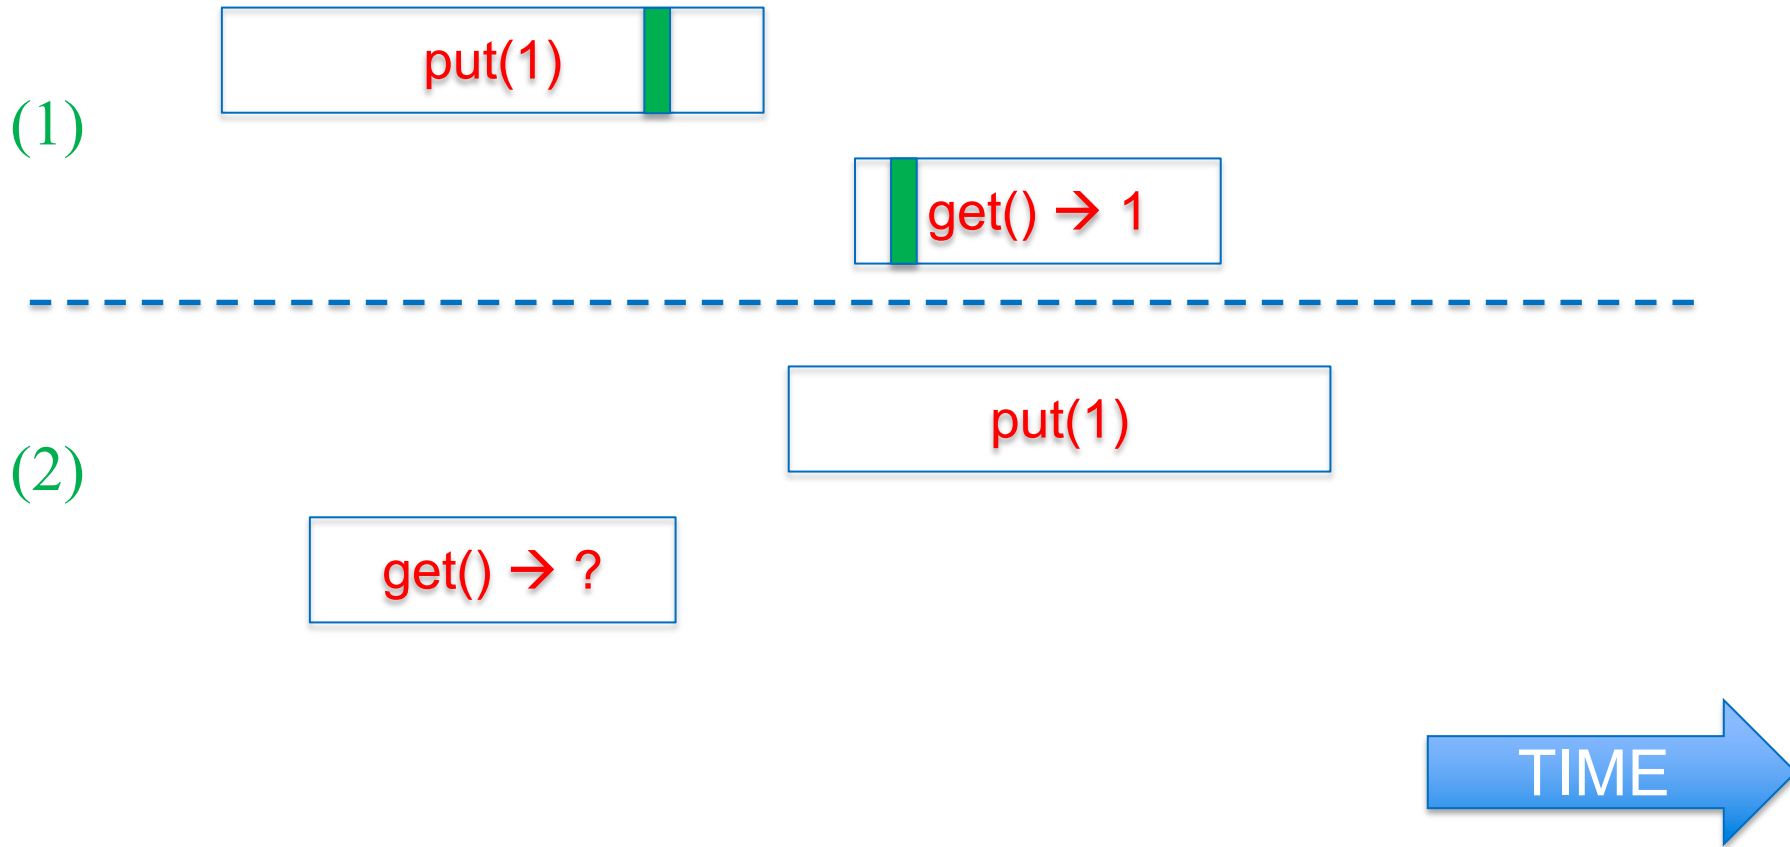


Correct Behaviors

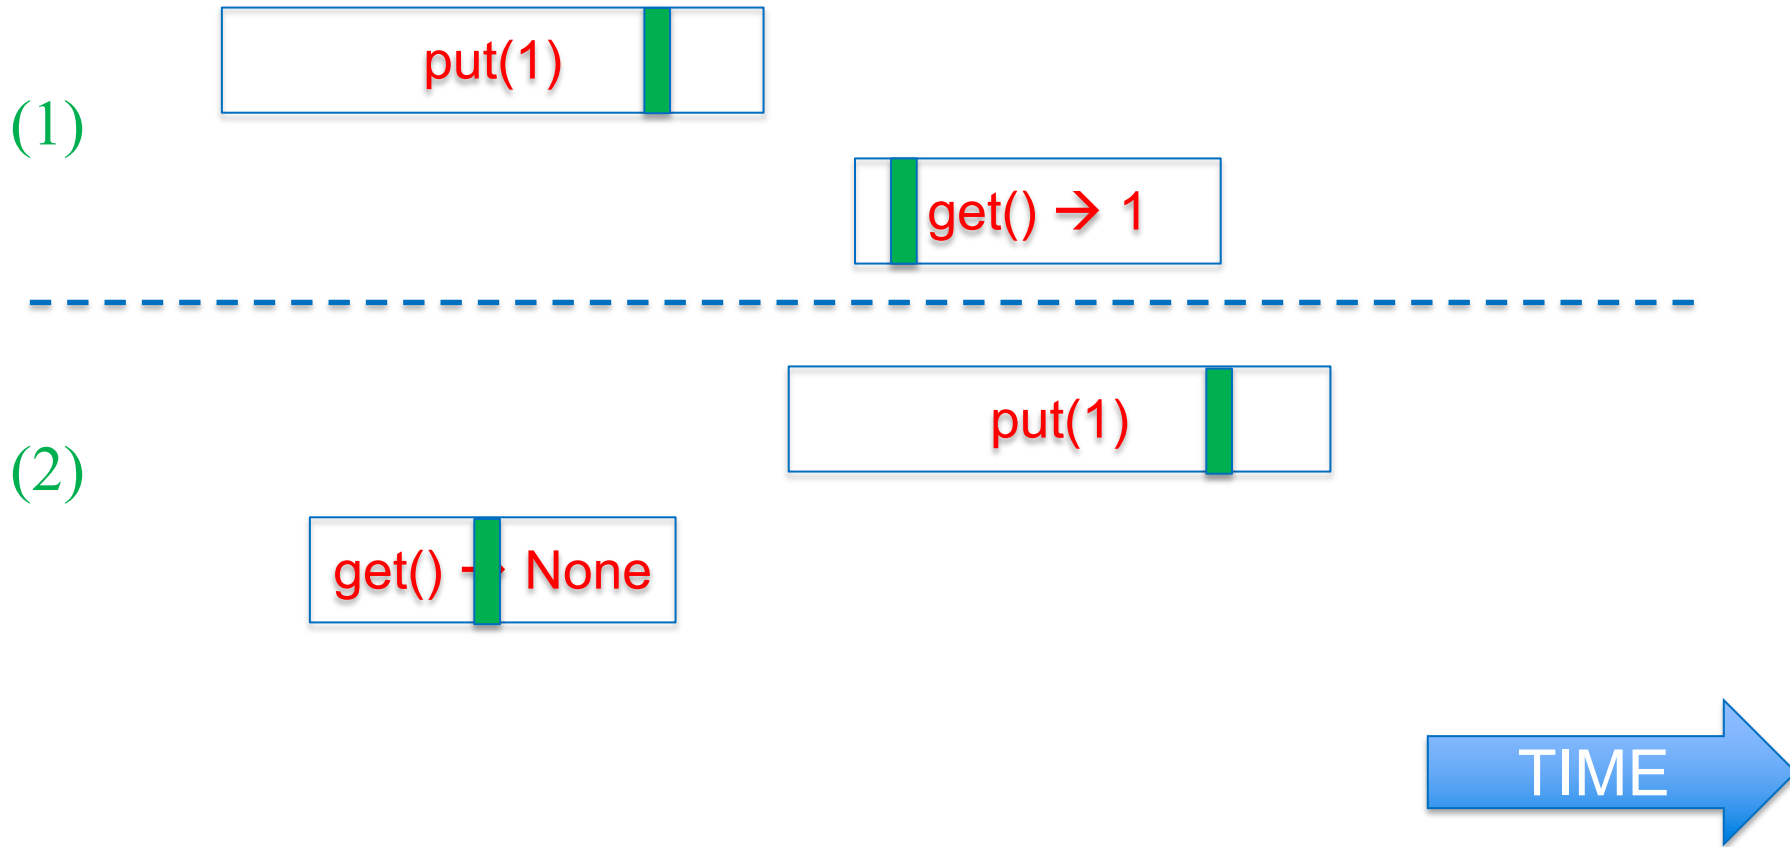
(1)



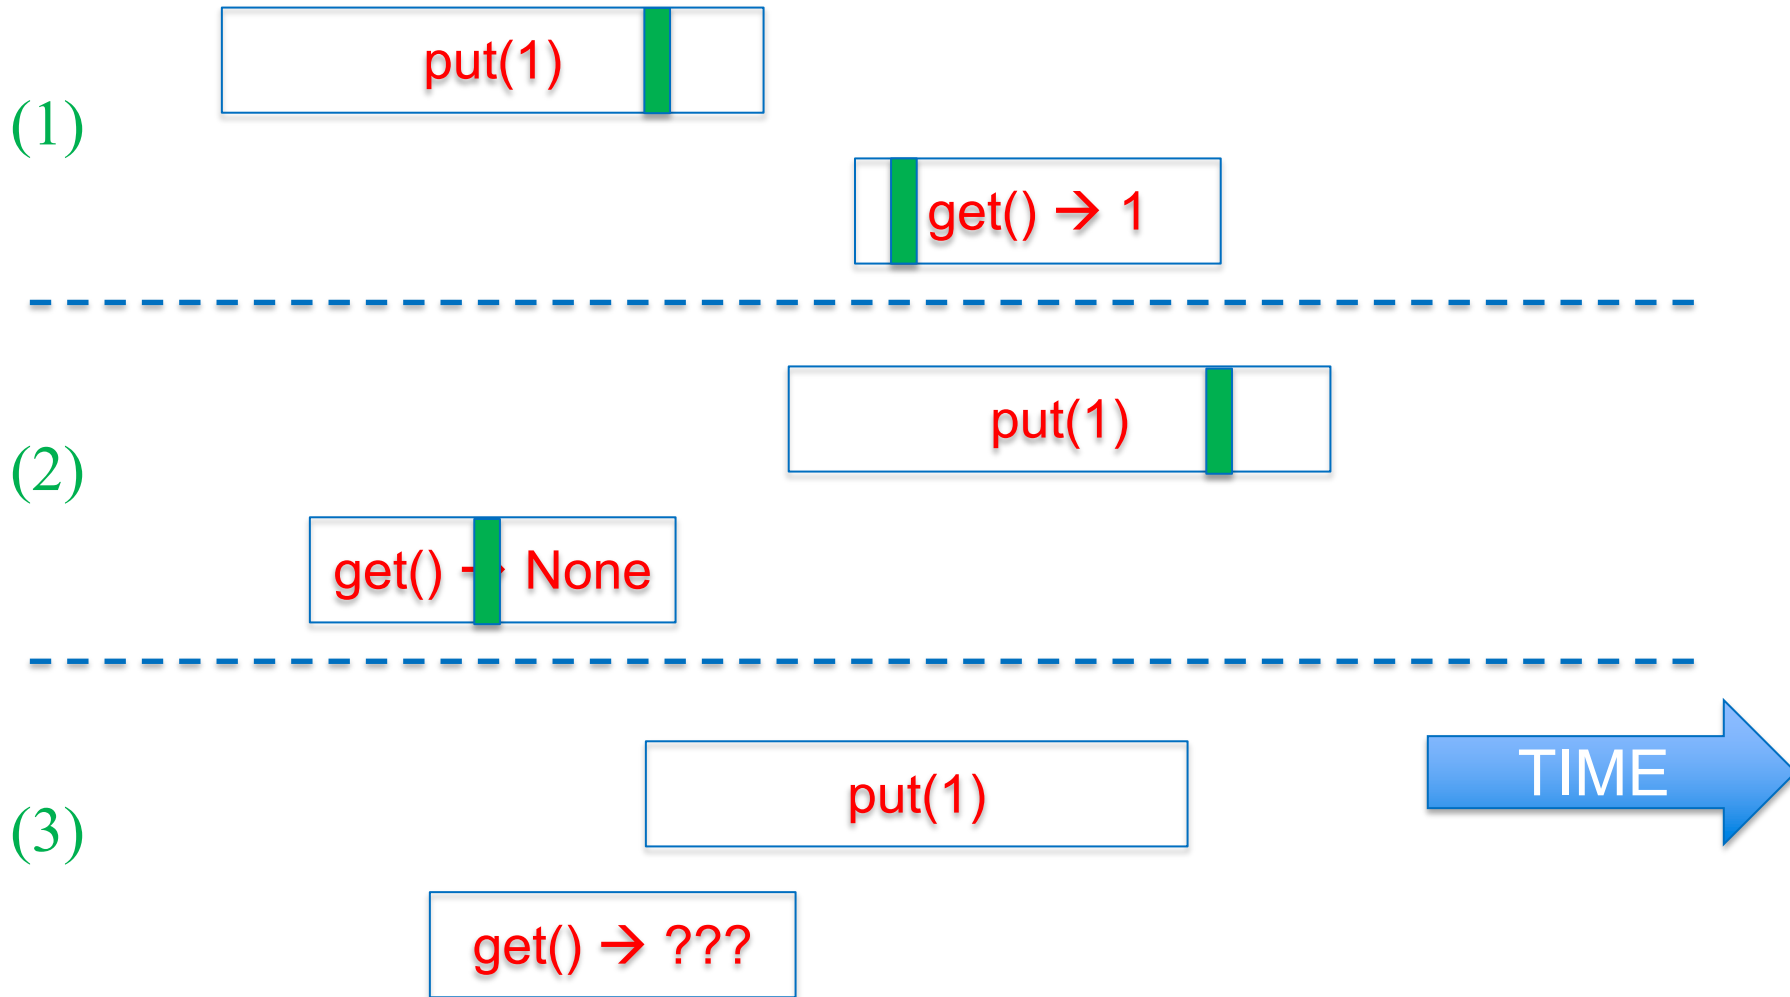
Correct Behaviors



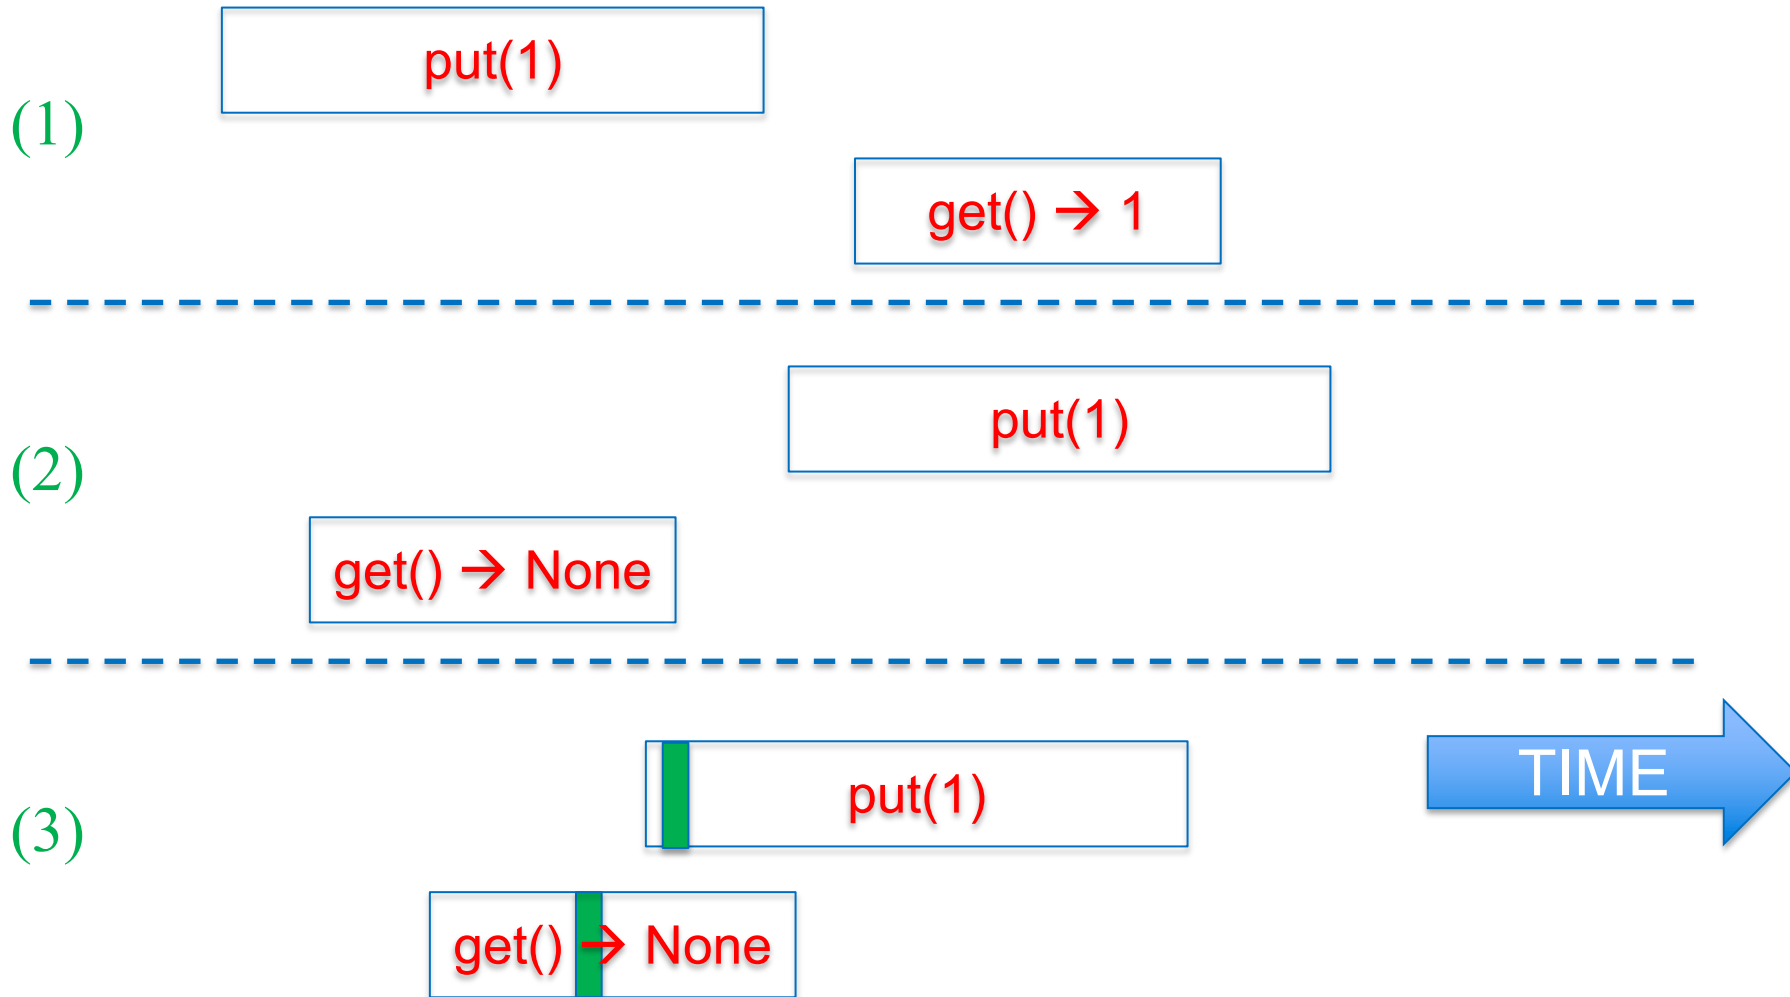
Correct Behaviors




Correct Behaviors



Correct Behaviors



Correct Behaviors



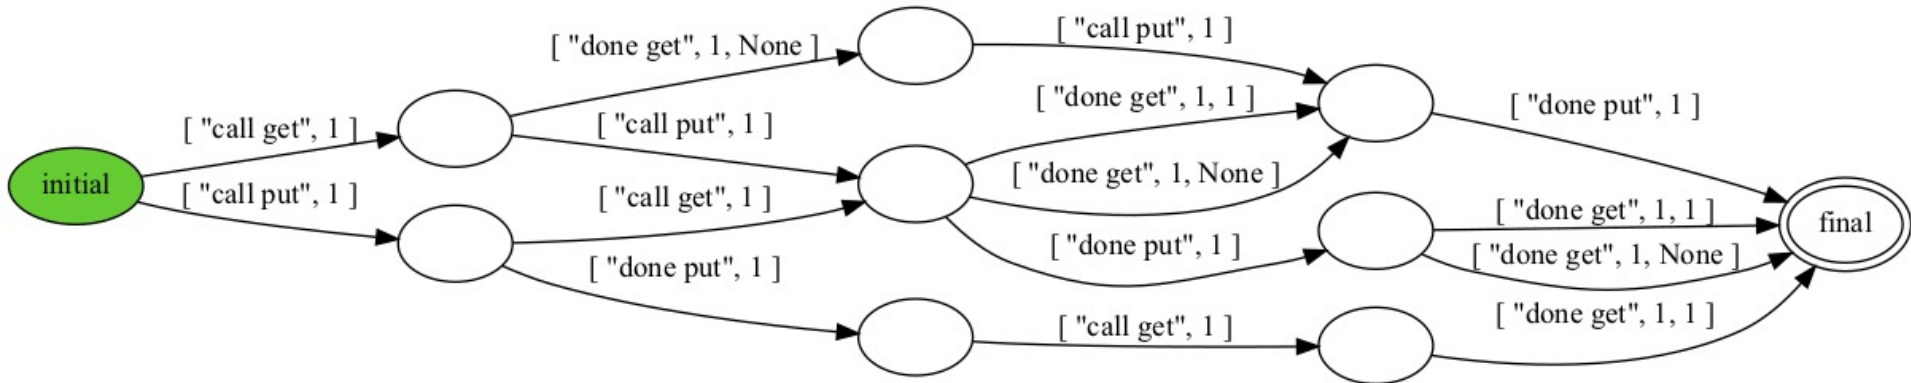
Testing Concurrent Objects

- Concurrent case
 - try all “interleavings” of 1 operation
 - try all interleavings of 2 operations
 - try all interleavings of 3 operations
 - ...
- How do you know if an interleaving is correct?
 - compare “behaviors” of running test against concurrent implementation with running test against the concurrent specification

Concurrent queue test program

```
1  import queue
2
3  const NOPS = 4
4  q = queue.Queue()
5
6  def put_test(self):
7      print("call put", self)
8      queue.put(?q, self)
9      print("done put", self)
10
11 def get_test(self):
12     print("call get", self)
13     let v = queue.get(?q):
14         print("done get", self, v)
15
16 nputs = choose {1..NOPS-1}
17 for i in {1..nputs}:
18     spawn put_test(i)
19 for i in {1..NOPS-nputs}:
20     spawn get_test(i)
```

Behavior (NOPS=2: 1 get, 1 put)



```
$ harmony -c NOPS=2 -o spec.png code/queue_btest1.hny
```

Testing: comparing behaviors

```
$ harmony -o queue4.hfa code/queue_btest1.hny  
$ harmony -B queue4.hfa -m queue=queue_lock code/queue_btest1.hny
```

- The first command outputs the behavior of running the test program against the specification in file queue4.hfa
- The second command runs the test program against the implementation and checks if its behavior matches that stored in queue4.hfa