


# Concurrent Programming in Harmony: 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
- Further complicated by interrupt handlers that also access those data structures

# So I talked with a recruiter last week...

- Not making this up...

# 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 reason about concurrent programs?
- How to construct correct concurrent programs?

# 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**

# Two Threads, One Variable

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?

# Two Threads, One Variable


Might execute like this:

T1

```
...  
r1 = load from amount  
r1 = r1 - 10,000  
store r1 to amount  
...
```

T2

```
...  
r2 = load from amount  
r2 = r2 / 2  
store r2 to amount  
...
```



Or vice versa (T1 then T2 → 45,000)...  
either way is fine...

# Two Threads, One Variable


Or it might execute like this:

T1

```
...  
r1 = load from amount  
r1 = r1 - 10,000  
store r1 to amount  
...
```

T2

```
...  
r2 = load from amount  
...  
r2 = r2 / 2  
store r2 to amount  
...
```




*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 (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

# My experience until now

1. Students develop their concurrent 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
5. Finds most submissions are broken
6. RVR unhappy, students unhappy

# It's not stupidity

- 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 about concurrency get it wrong sometimes

# My take on the problem

- Handwritten correctness proofs just as likely to have bugs as programs
  - or even more likely as you can't test handwritten proofs
- Lack of mainstream tools to check concurrent algorithms
- Tools that do exist have a steep learning curve

# Enter *Harmony*

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

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

# Example (same as before)

```
def T1():  
    amount -= 10000;  
    done1 = True;  
;  
def T2():  
    amount /= 2;  
    done2 = True;  
;
```

# Example (same as before)

```
def T1():  
    amount -= 10000;  
    done1 = True;  
;  
def T2():  
    amount /= 2;  
    done2 = True;  
;  
def main():  
    await done1 and done2;  
    assert (amount == 40000) or (amount == 45000), amount;  
;  
done1 = done2 = False;  
amount = 100000;  
spawn T1();  
spawn T2();  
spawn main();
```

# Example (same as before)

```
def T1():
    amount -= 10000;
    done1 = True;
;
def T2():
    amount /= 2;
    done2 = True;
;
def main():
    await done1 and done2;
    assert (amount == 40000) or (amount == 45000), amount;
;
done1 = done2 = False;
amount = 100000;
spawn T1();
spawn T2();
spawn main();
```

Equivalent to:

```
while not (done1 and done2):
    pass;
```

;

# Example (same as before)

```
def T1():
    amount -= 10000;
    done1 = True;
;
def T2():
    amount /= 2;
    done2 = True;
;
def main():
    await done1 and done2;
    assert (amount == 40000) or (amount == 45000), amount;
;
done1 = done2 = False;
amount = 100000;
spawn T1();
spawn T2();
spawn main();
```

Assertion: useful to check properties

# Example (same as before)

```
def T1():
    amount -= 10000;
    done1 = True;
;
def T2():
    amount /= 2;
    done2 = True;
;
def main():
    await done1 and done2;
    assert (amount == 40000) or (amount == 45000), amount;
;
done1 = done2 = False;
amount = 100000;
spawn T1();
spawn T2();
spawn main();
```

Output amount if assertion fails

# Example (same as before)

```
def T1():  
    amount -= 10000;  
    done1 = True;  
;  
def T2():  
    amount /= 2;  
    done2 = True;  
;  
def main():  
    await done1 and done2;  
    assert (amount == 40000) or (amount == 45000), amount;  
;  
done1 = done2 = False;  
amount = 100000;  
spawn T1();  
spawn T2();  
spawn main();
```

Initialize shared variables

# Example (same as before)

```
def T1():  
    amount -= 10000;  
    done1 = True;  
;  
def T2():  
    amount /= 2;  
    done2 = True;  
;  
def main():  
    await done1 and done2;  
    assert (amount == 40000) or (amount == 45000), amount;  
;  
done1 = done2 = False;  
amount = 100000;  
spawn T1();  
spawn T2();  
spawn main();
```

Spawn three processes (threads)

# Example (same as before)

```
def T1():
    amount -= 10000;
    done1 = True;
;
def T2():
    amount /= 2;
    done2 = True;
;
def main():
    await done1 and done2;
    assert (amount == 40000) or (amount == 45000), amount;
;
```

```
done1 = done2 = False;
amount = 100000;
spawn T1();
spawn T2();
spawn main();
```

```
#states = 100 diameter = 5
===== Safety violation =====
__init__ /() [0,40-58] 58 { amount: 100000, done1: False, done2: False }
T1/() [1-4] 5 { amount: 100000, done1: False, done2: False }
T2/() [10-17]. 17 { amount: 50000, done1: False, done2: True }
T1/() [5-8] 8 { amount: 90000, done1: True, done2: True }
main/() [19-23,25-34,36-37] 37 { amount: 90000, done1: True, done2: True }
>>> Harmony Assertion (file=test.hny, line=11) failed: 90000
```

# 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



# 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 Output

```
#states = 100 diameter = 5
===== Safety violation =====
__init__()/ [0,40-58] 58 { amount: 100000, done1: False, done2: False }
T1()/ [1-4]          5 { amount: 100000, done1: False, done2: False }
T2()/ [10-17].      17 { amount: 50000, done1: False, done2: True }
T1()/ [5-8]         8 { amount: 90000, done1: True, done2: True }
main()/ [19-23,25-34,36-37] 37 { amount: 90000, done1: True, done2: True }
>>> Harmony Assertion (file=test.hny, line=11) failed: 90000
```

# Output

#states in the state graph

#states = 100 diameter = 5

==== Safety violation =====

```
__init__ /() [0,40-58] 58 { amount: 100000, done1: False, done2: False }
T1/() [1-4]          5 { amount: 100000, done1: False, done2: False }
T2/() [10-17].      17 { amount: 50000, done1: False, done2: True }
T1/() [5-8]         8 { amount: 90000, done1: True, done2: True }
main/() [19-23,25-34,36-37] 37 { amount: 90000, done1: True, done2: True }
>>> Harmony Assertion (file=test.hny, line=11) failed: 90000
```

# Output

diameter of the state graph

```
#states = 100 diameter = 5
```

```
==== Safety violation =====
```

```
__init__ /() [0,40-58] 58 { amount: 100000, done1: False, done2: False }  
T1/() [1-4]          5 { amount: 100000, done1: False, done2: False }  
T2/() [10-17].      17 { amount: 50000, done1: False, done2: True }  
T1/() [5-8]         8 { amount: 90000, done1: True, done2: True }  
main/() [19-23,25-34,36-37] 37 { amount: 90000, done1: True, done2: True }  
>>> Harmony Assertion (file=test.hny, line=11) failed: 90000
```

# Output

something went wrong in (at least) one path in the graph  
(*assertion failure*)

```
#states = 100 diameter = 5
```

```
==== Safety violation ====
```

```
__init__ /() [0,40-58] 58 { amount: 100000, done1: False, done2: False }  
T1/() [1-4]          5 { amount: 100000, done1: False, done2: False }  
T2/() [10-17]       17 { amount: 50000, done1: False, done2: True }  
T1/() [5-8]         8 { amount: 90000, done1: True, done2: True }  
main/() [19-23,25-34,36-37] 37 { amount: 90000, done1: True, done2: True }  
>>> Harmony Assertion (file=test.hny, line=11) failed: 90000
```

# Output

shortest path to  
assertion failure

#states = 100 diameter = 5

==== Safety violation =====

path {

```
__init__ /() [0,40-58] 58 { amount: 100000, done1: False, done2: False }
T1 /() [1-4] 5 { amount: 100000, done1: False, done2: False }
T2 /() [10-17] 17 { amount: 50000, done1: False, done2: True }
T1 /() [5-8] 8 { amount: 90000, done1: True, done2: True }
main /() [19-23,25-34,36-37] 37 { amount: 90000, done1: True, done2: True }
```

>>> Harmony Assertion (file=test.hny, line=11) failed: 90000

# Output

```
#states = 100 diameter = 5
```

```
==== Safety violation ====
```

```
_init_ __init__/() [0,40-58] 58 { amount: 100000, done1: False, done2: False }  
T1/() [1-4] 5 { amount: 100000, done1: False, done2: False }  
T2/() [10-17] 17 { amount: 50000, done1: False, done2: True }  
T1/() [5-8] 8 { amount: 90000, done1: True, done2: True }  
main/() [19-23,25-34,36-37] 37 { amount: 90000, done1: True, done2: True }  
>>> Harmony Assertion (file=test.hny, line=11) failed: 90000
```

# Output

```
#states = 100 diameter = 5
```

```
===== Safety violation =====
```

```
__init__ __init__/() [0,40-58] 58 { amount: 100000, done1: False, done2: False }  
T1ab T1/() [1-4] 5 { amount: 100000, done1: False, done2: False }  
T2/() [10-17] 17 { amount: 50000, done1: False, done2: True }  
T1/() [5-8] 8 { amount: 90000, done1: True, done2: True }  
main/() [19-23,25-34,36-37] 37 { amount: 90000, done1: True, done2: True }  
>>> Harmony Assertion (file=test.hny, line=11) failed: 90000
```

# Output

```
#states = 100 diameter = 5
```

```
===== Safety violation =====
```

```
__init__ __init__/() [0,40-58] 58 { amount: 100000, done1: False, done2: False }  
T1abc T1/() [1-4] 5 { amount: 100000, done1: False, done2: False }  
T2abc T2/() [10-17] 17 { amount: 50000, done1: False, done2: True }  
T1/() [5-8] 8 { amount: 90000, done1: True, done2: True }  
main/() [19-23,25-34,36-37] 37 { amount: 90000, done1: True, done2: True }  
>>> Harmony Assertion (file=test.hny, line=11) failed: 90000
```

# Output

```
#states = 100 diameter = 5
```

```
===== Safety violation =====
```

```
_init_  __init__/() [0,40-58] 58 { amount: 100000, done1: False, done2: False }  
T1ab   T1/() [1-4]           5 { amount: 100000, done1: False, done2: False }  
T2abc  T2/() [10-17]          17 { amount: 50000,  done1: False, done2: True  }  
T1c   T1/() [5-8]           8 { amount: 90000,  done1: True,  done2: True  }  
main/() [19-23,25-34,36-37] 37 { amount: 90000, done1: True, done2: True }  
>>> Harmony Assertion (file=test.hny, line=11) failed: 90000
```

# Output

```
#states = 100 diameter = 5
```

```
===== Safety violation =====
```

```
_init_  __init__/() [0,40-58] 58 { amount: 100000, done1: False, done2: False }  
T1ab   T1/() [1-4]           5 { amount: 100000, done1: False, done2: False }  
T2abc  T2/() [10-17]          17 { amount: 50000, done1: False, done2: True }  
T1c    T1/() [5-8]           8 { amount: 90000, done1: True, done2: True }  
main   main/() [19-23,25-34,36-37] 37 { amount: 90000, done1: True, done2: True }  
>>> Harmony Assertion (file=test.hny, line=11) failed: 90000
```

# Output



```
#states = 100 diameter = 5
```

```
===== Safety violation =====
```

```
_init_  __init__ /() [0,40-58] 58 { amount: 100000, done1: False, done2: False }  
T1ab    T1/() [1-4]           5 { amount: 100000, done1: False, done2: False }  
T2abc   T2/() [10-17]        17 { amount: 50000, done1: False, done2: True }  
T1c     T1/() [5-8]          8 { amount: 90000, done1: True, done2: True }  
main    main/() [19-23,25-34,36-37] 37 { amount: 90000, done1: True, done2: True }  
>>> Harmony Assertion (file=test.hny, line=11) failed: 90000
```

# Output

“name tag” of a process

```
__init__ /() [0,40-58] 58 { amount: 100000, done1: False, done2: False }  
T1/() [1-4] 5 { amount: 100000, done1: False, done2: False }  
T2/() [10-17]. 17 { amount: 50000, done1: False, done2: True }  
T1/() [5-8] 8 { amount: 90000, done1: True, done2: True }  
main/() [19-23,25-34,36-37] 37 { amount: 90000, done1: True, done2: True }
```

# Output

“microsteps” =  
list of program counters  
of machine instructions  
executed

```
__init__ /() [0,40-58] 58 { amount: 100000, done1: False, done2: False }  
T1/() [1-4]          5 { amount: 100000, done1: False, done2: False }  
T2/() [10-17]       17 { amount: 50000, done1: False, done2: True }  
T1/() [5-8]         8 { amount: 90000, done1: True, done2: True }  
main/() [19-23,25-34,36-37] 37 { amount: 90000, done1: True, done2: True }
```

# Harmony Machine Code

0 Jump 40

1 Frame T1 ()

2 Load amount      T1a: LOAD amount

3 Push 10000

4 2-ary –            T1b: SUB 10000

5 Store amount      T1c: STORE amount

6 Push True

7 Store done1       T1d: done1 = True

8 Return

9 Jump 40

10 Frame T2 ()

11 Load amount      T2a: LOAD amount

12 Push 2

13 2-ary /            T2b: DIV 2

14 Store amount      T2c: STORE amount

15 Push True

16 Store done2       T2d: done2 = True

17 Return

18 ...

# Harmony Machine Code

0 Jump 40

PC := 40

1 Frame T1 ()

2 Load amount

3 Push 10000

4 2-ary —

5 Store amount

6 Push True

7 Store done1

8 Return

9 Jump 40

10 Frame T2 ()

11 Load amount

12 Push 2

13 2-ary /

14 Store amount

15 Push True

16 Store done2

17 Return

18 ...

# Harmony Machine Code

0 Jump 40

PC := 40

1 Frame T1 ()

2 Load amount

push amount onto the stack of process T1

3 Push 10000

4 2-ary —

5 Store amount

6 Push True

7 Store done1

8 Return

9 Jump 40

10 Frame T2 ()

11 Load amount

12 Push 2

13 2-ary /

14 Store amount

15 Push True

16 Store done2

17 Return

18 ...

# Harmony Machine Code

0 Jump 40

PC := 40

1 Frame T1 ()

2 Load amount

push amount onto the stack of process T1

3 Push 10000

push 10000 onto the stack of process T1

4 2-ary –

replace top two elements of stack with difference

5 Store amount

6 Push True

7 Store done1

8 Return

9 Jump 40

10 Frame T2 ()

11 Load amount

12 Push 2

13 2-ary /

14 Store amount

15 Push True

16 Store done2

17 Return

18 ...

# Harmony Machine Code

0 Jump 40

PC := 40

1 Frame T1 ()

2 Load amount

push amount onto the stack of process T1

3 Push 10000

push 10000 onto the stack of process T1

4 2-ary –

replace top two elements of stack with difference

5 Store amount

store top of the stack of T1 into amount

6 Push True

7 Store done1

8 Return

9 Jump 40

10 Frame T2 ()

11 Load amount

12 Push 2

13 2-ary /

14 Store amount

15 Push True

16 Store done2

17 Return

18 ...

# Harmony Machine Code

0 Jump 40

PC := 40

1 Frame T1 ()

2 Load amount

push amount onto the stack of process T1

3 Push 10000

push 10000 onto the stack of process T1

4 2-ary –

replace top two elements of stack with difference

5 Store amount

store top of the stack of T1 into amount

6 Push True

push True onto the stack of process T1

7 Store done1

store top of the stack of T1 into done1

8 Return

9 Jump 40

10 Frame T2 ()

11 Load amount

12 Push 2

13 2-ary /

14 Store amount

15 Push True

16 Store done2

17 Return

18 ...

# Harmony Machine Code

0 Jump 40

PC := 40

1 Frame T1 ()

2 Load amount

push amount onto the stack of process T1

3 Push 10000

push 10000 onto the stack of process T1

4 2-ary –

replace top two elements of stack with difference

5 Store amount

store top of the stack of T1 into amount

6 Push True

push True onto the stack of process T1

7 Store done1

store top of the stack of T1 into done1

8 Return

9 Jump 40

10 Frame T2 ()

11 Load amount

push amount onto the stack of process T2

12 Push 2

push 2 onto the stack of process T2

13 2-ary /

replace top two elements of stack with division

14 Store amount

store top of the stack of T2 into amount

15 Push True

push True onto the stack of process T2

16 Store done2

store top of the stack of T2 into done2

17 Return

18 ...

# Output

current program counter  
(after microsteps)

```
__init__()/() [0,40-58] 58 { amount: 100000, done1: False, done2: False }  
T1()/() [1-4]          5 { amount: 100000, done1: False, done2: False }  
T2()/() [10-17]       17 { amount: 50000, done1: False, done2: True }  
T1()/() [5-8]         8 { amount: 90000, done1: True, done2: True }  
main()/() [19-23,25-34,36-37] 37 { amount: 90000, done1: True, done2: True }
```

# Output

current state  
(after microsteps)

```
__init__()/() [0,40-58] 58 { amount: 100000, done1: False, done2: False }  
T1()/() [1-4]          5 { amount: 100000, done1: False, done2: False }  
T2()/() [10-17]       17 { amount: 50000, done1: False, done2: True }  
T1()/() [5-8]         8 { amount: 90000, done1: True, done2: True }  
main()/() [19-23,25-34,36-37] 37 { amount: 90000, done1: True, done2: True }
```

# Harmony Virtual Machine *State*

Three parts:

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

State represents one vertex in the graph model

# *Context* (state of a process)

- Name tag
- PC (program counter)
- stack (+ implicit stack pointer)
- local variables
  - parameters (aka arguments)
  - “result”
    - there is no **return** statement
- local variables
  - declared in **let** and **for** statements

# Harmony != Python

Harmony	Python
tries all possible executions	executes just one
every statement ends in ;	; at end of statement optional
indentation recommended	indentation required
( ... ) == [ ... ] == ...	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
{ } is empty set	set() != { }
few operator precedence rules --- use brackets often	many operator precedence rules
variables global unless declared otherwise	depends... Sometimes must be explicitly declared global
no <b>return</b> , <b>break</b> , <b>continue</b>	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:
  - **assert**  $x + y < 10$
  - **assert**  $x + y < 10, (x, y)$

# I/O in Harmony?

- Input:
  - **choose** expression
    - `x = choose({ 1, 2, 3 })`
    - allows Harmony to choose

- **const**

- `Harmony` with “-c x=4” flag to Harmony

- `assert x + y < 10`
- `assert x + y < 10, (x, y)`


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

# Non-determinism in Harmony

Two sources:

1. **choose** expressions
2. **process** interleavings

# Limitation: models must be finite!





- But models are allowed to have cycles.
- Executions are allowed to be unbounded!
- Harmony does check 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

100,000

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:** all threads make it into the CS if desired

**Fairness:** equal chances of getting into CS

... in practice, fairness rarely guaranteed

# Critical Section

Must be serialized due to shared memory access



## Goals

**Mutual Exclusion:** 1 thread in a critical section at time

**Progress:** all threads make it into the CS if desired

**Fairness:** equal chances of getting into CS

... in practice, fairness rarely guaranteed

# Critical Sections in Harmony

```
def process(self):  
    while True:  
        ... # code outside critical section  
        ... # code to enter the critical section  
        ... # critical section itself  
        ... # code to exit the critical section  
    ;  
;  
spawn process(1);  
spawn process(2);  
...
```

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

# Critical Sections in Harmony

```
def process(self):
    while True:
        ... # code outside critical section
        ... # code to enter the critical section
        @cs: assert atLabel.cs == dict{ nametag(): 1 };
        ... # code to exit the critical section
    ;
;
spawn process(1);
spawn process(2);
...
```

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

# Critical Sections in Harmony

```
def process(self):
    while choose( { False, True } ):
        ... # code outside critical section
        ... # code to enter the critical section
        @cs: assert atLabel.cs == dict{ nametag(): 1 };
        ... # code to exit the critical section
    ;
;
spawn process(1);
spawn process(2);
...
```

- How do we check mutual exclusion?
- How do we check progress?
  - *if code to enter/exit the critical section does not terminate, Harmony with balk*

# *Sounds like you need a lock...*

- True, but this is an O.S. class!
- The question is:

*How does one build a lock?*

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

*You have to program them!*

# First attempt: a naïve lock

```
1  def process(self):
2      while choose({ False, True }):
3          # Enter critical section
4          await not lockTaken;
5          lockTaken = True;
6
7          # Critical section
8          @cs: assert atLabel.cs == dict{ nametag(): 1 };
9
10         # Leave critical section
11         lockTaken = False;
12     ;
13 ;
14 lockTaken = False;
15 spawn process(0);
16 spawn process(1);
```

Figure 5.4: [[code/naiveLock.hny](#)] Naïve implementation of a shared lock.

# First attempt: a naïve lock

```
1  def process(self):
2      while choose({ False, True }):
3          # Enter critical section
4          await not lockTaken;
5          lockTaken = True;
6
7          # Critical section
8          @cs: assert atLabel.cs == dict{ nametag(): 1 };
9
10         # Leave critical section
11         lockTaken = False;
12     ;
13 ;
14 lockTaken = False;
15 spawn process(0);
16 spawn process(1);
```



Figure 5.4: [[code/naiveLock.hny](#)] Naïve implementation of a shared lock.

# First attempt: a naïve lock

```
1  def process(self):
2      while choose({ False, True }):
3          # Enter critical section
4          await not lockTaken;
5          lockTaken = True;
6
7          # Critical section
8          @cs: assert atLabel.cs == dict{ nametag(): 1 };
9
10         # Leave critical section
11         lockTaken = False;
12     ;
13 ;
14 lockTaken = False;
15 spawn process(0);
16 spawn process(1);
```



wait till lock is free, then take it

release the lock

Figure 5.4: [[code/naiveLock.hny](#)] Naïve implementation of a shared lock.

# First attempt: a naïve lock

```
1  def process(self):
2      while choose({ False, True }):
3          # Enter critical section
4          await not lockTaken;
5          lockTaken = True;
6
7          # Critical section
8          @cs: assert atLabel.cs == dict{ nametag(): 1 };
9
10         # Leave critical section
11         lockTaken = False;
12     ;
13 ;
14 lockTaken = False;
15 spawn process(0);
16 spawn process(1);
```

==== Safety violation ====

__init__ /() [0,26-36]	36 { lockTaken: False }
process/0 [1-2,3(choose True),4-7]	8 { lockTaken: False }
process/1 [1-2,3(choose True),4-8]	9 { lockTaken: True }
process/0 [8-19]	19 { lockTaken: True }

Figure 5.4: [code/naiv](#)

>>> Harmony Assertion (file=code/naiveLock.hny, line=8) failed

# Second attempt: *flags*

```
1  def process(self):
2      while choose({ False, True }):
3          # Enter critical section
4          flags[self] = True;
5          await not flags[1 - self];
6
7          # Critical section
8          @cs: assert atLabel.cs == dict{ nametag(): 1 };
9
10         # Leave critical section
11         flags[self] = False;
12     ;
13 ;
14 flags = [ False, False ];
15 spawn process(0);
16 spawn process(1);
```

Figure 5.6: [[code/naiveFlags.hny](#)] Naïve use of flags to solve mutual exclusion.

# Second attempt: *flags*

```
1  def process(self):
2      while choose({ False, True }):
3          # Enter critical section
4          flags[self] = True;
5          await not flags[1 - self];
6
7          # Critical section
8          @cs: assert atLabel.cs == dict{ nametag(): 1 };
9
10         # Leave critical section
11         flags[self] = False;
12     ;
13 ;
14 flags = [ False, False ];
15 spawn process(0);
16 spawn process(1);
```



Figure 5.6: [[code/naiveFlags.hny](#)] Naïve use of flags to solve mutual exclusion.

# Second attempt: *flags*

```
1  def process(self):
2      while choose({ False, True }):
3          # Enter critical section
4          flags[self] = True;
5          await not flags[1 - self];
6
7          # Critical section
8          @cs: assert atLabel.cs == dict{ nametag(): 1 };
9
10         # Leave critical section
11         flags[self] = False;
12     ;
13 ;
14 flags = [ False, False ];
15 spawn process(0);
16 spawn process(1);
```




Figure 5.6: [[code/naiveFlags.hny](#)] Naïve use of flags to solve mutual exclusion.

# Second attempt: *flags*

```
1  def process(self):
2      while choose({ False, True }):
3          # Enter critical section
4          flags[self] = True;
5          await not flags[1 - self];
6
7          # Critical section
8          @cs: assert atLabel.cs == dict{ nametag(): 1 };
9
10         # Leave critical section
11         flags[self] = False;
12     ;
13 ;
14 flags = [ False, False ]
15 spawn process(0);
16 spawn process(1);
```

==== Non-terminating State ====

\_\_init\_\_ /() [0,36-46] 46 { flags: [False, False] }  
process/0 [1-2,3(choose True),4-12] 13 { flags: [True, False] }  
process/1 [1-2,3(choose True),4-12] 13 { flags: [True, True] }  
blocked process: process/1 pc = 13  
blocked process: process/0 pc = 13

Figure 5.6: [code/naiv](#)

# Third attempt: *turn* variable

```
1  def process(self):
2      while choose({ False, True }):
3          # Enter critical section
4          await turn == self;
5
6          # Critical section
7          @cs: assert atLabel.cs == dict{ nametag(): 1 };
8
9          # Leave critical section
10         turn = 1 - self;
11     ;
12 ;
13 turn = 0;
14 spawn process(0);
15 spawn process(1);
```

Figure 5.8: [[code/naiveTurn.hny](#)] Naïve use of turn variable to solve mutual exclusion.

# Third attempt: *turn* variable


```
1  def process(self):
2      while choose({ False, True }):
3          # Enter critical section
4          await turn == self;
5
6          # Critical section
7          @cs: assert atLabel.cs == dict{ nametag(): 1 };
8
9          # Leave critical section
10         turn = 1 - self;
11     ;
12 ;
13 turn = 0;
14 spawn process(0);
15 spawn process(1);
```



Figure 5.8: [[code/naiveTurn.hny](#)] Naïve use of turn variable to solve mutual exclusion.

# Third attempt: *turn* variable

```
1  def process(self):
2      while choose({ False, True }):
3          # Enter critical section
4          await turn == self;
5
6          # Critical section
7          @cs: assert atLabel.cs == dict{ nametag(): 1 };
8
9          # Leave critical section
10         turn = 1 - self;
11     ;
12 ;
13 turn = 0;
14 spawn process(0);
15 spawn process(1);
```



wait for your turn

let the other process take a turn

Figure 5.8: [[code/naiveTurn.hny](#)] Naïve use of turn variable to solve mutual exclusion.

# Third attempt: *turn* variable

```
1  def process(self):
2      while choose({ False, True }):
3          # Enter critical section
4          await turn == self;
5
6          # Critical section
7          @cs: assert atLabel.cs == dict{ nametag(): 1 };
8
9          # Leave critical section
10         turn = 1 - self;
11     ;
12 ;
13 turn = 0;
14 spawn process(0);
15 spawn process(1);
```

==== Non-terminating State ====

```
__init__ /() [0,28-38]          38 { turn: 0 }
process/0 [1-2,3(choose True),4-26,2,3(choose True),4] 5 { turn: 1 }
process/1 [1-2,3(choose False),4,27]          27 { turn: 1 }
blocked process: process/0 pc = 5
```


Figure 5.8: [\[code/naiveTu](#)

# Peterson's Algorithm: *flags & turn*

```
1  def process(self):
2      while choose({ False, True }):
3          # Enter critical section
4          flags[self] = True;
5          turn = 1 - self;
6          await (not flags[1 - self]) or (turn == self);
7
8          # critical section is here
9          @cs: assert atLabel.cs == dict{ nametag(): 1 };
10
11         # Leave critical section
12         flags[self] = False;
13     ;
14 ;
15 flags = [ False, False ];
16 turn = choose({0, 1});
17 spawn process(0);
18 spawn process(1);
```


# Peterson's Algorithm: *flags & turn*

```
1  def process(self):
2      while choose({ False, True }):
3          # Enter critical section
4          flags[self] = True;
5          turn = 1 - self;
6          await (not flags[1 - self]) or (turn == self);
7
8          # critical section is here
9          @cs: assert atLabel.cs == dict{ nametag(): 1 };
10
11         # Leave critical section
12         flags[self] = False;
13     ;
14 ;
15 flags = [ False, False ];
16 turn = choose({0, 1});
17 spawn process(0);
18 spawn process(1);
```



# Peterson's Algorithm: *flags & turn*

```
1  def process(self):
2      while choose({ False, True }):
3          # Enter critical section
4          flags[self] = True;
5          turn = 1 - self;
6          await (not flags[1 - self]) or (turn == self);
7
8          # critical section is here
9          @cs: assert atLabel.cs == dict{ nametag(): 1 };
10
11         # Leave critical section
12         flags[self] = False;
13     ;
14 ;
15 flags = [ False, False ];
16 turn = choose({0, 1});
17 spawn process(0);
18 spawn process(1);
```




"you go first"



wait until alone or  
it's my turn

# Peterson's Algorithm: *flags & turn*

```
1  def process(self):
2      while choose({ False, True }):
3          # Enter critical section
4          flags[self] = True;
5          turn = 1 - self;
6          await (not flags[1 - self]) or (turn == self);
7
8          # critical section is here
9          @cs: assert atLabel.cs == dict{ nametag(): 1 };
10
11         # Leave critical section
12         flags[self] = False;
13     ;
14 ;
15 flags = [ False, False ];
16 turn = choose({0, 1});
17 spawn process(0);
18 spawn process(1);
```



“you go first”



wait until alone or  
it's my turn



leave

# Peterson's Algorithm: *flags & turn*

```
1  def process(self):
2      while choose({ False, True }):
3          # Enter critical section
4          flags[self] = True;
5          turn = 1 - self;
6          await (not flags[1 - self]) or (turn == self);
7
8          # critical section is here
9          @cs: assert atLabel.cs == dict{ nametag(): 1 };
10
11         # Leave critical section
12         flags[self] = False;
13     ;
14 ;
15 flags = [ False, False ];
16 turn = choose({0, 1});
17 spawn process(0);
18 spawn process(1);
```

**#states = 104 diameter = 5**  
**#components: 37**  
**no issues found**

So, we proved Peterson's Algorithm correct by brute force, enumerating all possible executions.

*But how does one prove it by deduction?  
so one might understand **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*
- 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]**

# First question: what should IH be?

- Obvious answer: mutual exclusion itself
  - if  $P0$  is in the critical section, then  $P1$  is not
    - without loss of generality...
  - Formally:  $P0@cs \implies \neg P1@cs$
- Unfortunately, this won't work...

# State before P1 takes a step:

```
1  def process(self):
2      while choose({ False, True }):
3          # Enter critical section
4          flags[self] = True;
5          turn = 1 - self;
6  P1 → await (not flags[1 - self]) or (turn == self);
7
8          # critical section is here
9  P0 → @cs: assert atLabel.cs == dict{ nametag(): 1 };
10
11         # Leave critical section
12         flags[self] = False;
13     ;
14 ;
15 flags = [ False, False ];
16 turn = choose({0, 1});
17 spawn process(0);
18 spawn process(1);
```

flags == [ True, True ]  
turn == 1

# State after P1 takes a step:

```
1  def process(self):
2      while choose({ False, True }):
3          # Enter critical section
4          flags[self] = True;
5          turn = 1 - self;
6          await (not flags[1 - self]) or (turn == self);
7
8          # critical section is here
9  P0 → @cs: assert atLabel.cs == dict{ nametag(): 1 }; ← P1
10
11         # Leave critical section
12         flags[self] = False;
13     ;
14 ;
15 flags = [ False, False ];
16 turn = choose({0, 1});
17 spawn process(0);
18 spawn process(1);
```

flags == [ True, True ]  
turn == 1

So, is Peterson's Algorithm broken?

# No, it'll turn out this prior state cannot be reached from the initial state

```
1  def process(self):
2      while choose({ False, True }):
3          # Enter critical section
4          flags[self] = True;
5          turn = 1 - self;
6  P1 → await (not flags[1 - self]) or (turn == self);
7
8          # critical section is here
9  P0 → @cs: assert atLabel.cs == dict{ nametag(): 1 };
10
11         # Leave critical section
12         flags[self] = False;
13     ;
14 ;
15 flags = [ False, False ];
16 turn = choose({0, 1});
17 spawn process(0);
18 spawn process(1);
```

flags == [ True, True ]  
turn == 1

# Let's try another obvious one

- Based on the **await** condition:

$$P0@cs \Rightarrow \neg flags[1] \vee turn == 0$$

- Promising because if  $P0@cs \wedge P1@cs$  then

$$\left. \begin{array}{l} P0@cs \Rightarrow \neg flags[1] \vee turn == 0 \wedge \\ P1@cs \Rightarrow \neg flags[0] \vee turn == 1 \end{array} \right\} \Rightarrow \left\{ \begin{array}{l} turn == 0 \wedge \\ turn == 1 \end{array} \right. \\ \Rightarrow \text{False (therefore mutual exclusion)}$$

- Unfortunately, this is not an invariant...

# State before P1 takes a step:

```
1  def process(self):
2      while choose({ False, True }):
3           # Enter critical section
4          flags[self] = True;
5          turn = 1 - self;
6          await (not flags[1 - self]) or (turn == self);
7
8          # critical section is here
9           @cs: assert atLabel.cs == dict{ nametag(): 1 };
10
11         # Leave critical section
12         flags[self] = False;
13     ;
14 ;
15 flags = [ False, False ];
16 turn = choose({0, 1});
17 spawn process(0);
18 spawn process(1);
```

*flags == [ True, False ]*  
*turn == 1*

*P0@cs  $\Rightarrow \neg flags[1] \vee turn == 0$  holds*

note: this is a reachable state

# State after P1 takes a step:

```
1  def process(self):
2      while choose({ False, True }):
3          # Enter critical section
4          flags[self] = True;
5          P1 → turn = 1 - self;
6          await (not flags[1 - self]) or (turn == self);
7
8          # critical section is here
9          P0 → @cs: assert atLabel.cs == dict{ nametag(): 1 };
10
11         # Leave critical section
12         flags[self] = False;
13     ;
14 ;
15 flags = [ False, False ];
16 turn = choose({0, 1});
17 spawn process(0);
18 spawn process(1);
```

flags == [ True, True ]  
turn == 1

$P0@cs \Rightarrow \neg flags[1] \vee turn == 0$  violated

note: this is also a reachable state

# But suggests an improved hypothesis

```
1  def process(self):
2      while choose({ False, True }):
3          # Enter critical section
4          flags[self] = True;
5          P1 → @gate: turn = 1 - self;
6              await (not flags[1 - self]) or (turn == self);
7
8          # Critical section
9          P0 → @cs: assert (not flags[1 - self]) or (turn == self) or
10              (atLabel.gate == dict{nametags[1 - self] : 1})
11               $P0@cs \Rightarrow \neg flags[1] \vee turn == 0 \vee P1@gate$ 
12
13          # Leave critical section
14          flags[self] = False;
15      ;
16  ;
17  flags = [ False, False ];
18  turn = choose({0, 1});
19  nametags = [ dict{ .name: .process, .tag: tag } for tag in {0, 1} ];
20  spawn process(0);
21  spawn process(1);
```

Figure 6.3: [[code/PetersonInductive.hny](#)] Peterson's Algorithm with Inductive Invariant

# Invariance proof

To prove:  $P0@cs \Rightarrow \neg flags[1] \vee turn == 0 \vee P1@gate$

By induction:

Base case:

- In initial state  $\neg P0@cs$ 
  - false implies anything

Induction Step: assume  $P0@cs$  and  $P1$  takes a step when

Case 1:  $\neg flags[1]$

then after step either  $\neg flags[1]$  or  $P1@gate$

Case 2:  $turn == 0$

then after step still  $turn == 0$  ( $P1$  never sets turn to 1)

Case 3:  $P1@gate$

then after step  $turn == 0$

# Finally, prove mutual exclusion

$$P0@cs \wedge P1@cs \implies$$

$$\left\{ \begin{array}{l} \neg flags[1] \vee turn == 0 \vee P1@gate \\ \neg flags[0] \vee turn == 1 \vee P0@gate \end{array} \right. \wedge$$

$$\implies turn == 0 \wedge turn == 1$$

$$\implies \textit{False}$$



# Finally, prove mutual exclusion

$$P0@cs \wedge P1@cs \implies$$


$$\left\{ \begin{array}{l} \neg flags[1] \vee turn == 0 \vee P1@gate \\ \neg flags[0] \vee turn == 1 \vee P0@gate \end{array} \right. \wedge$$

$$\implies turn == 0 \wedge turn == 1$$

$$\implies \textit{False}$$




# Review in Pictures: State Space




Mutual Exclusion Holds

Mutual  
Exclusion  
Violated


# Review in Pictures: State Space



# Review in Pictures: State Space



# Review in Pictures: State Space



# Inductive Invariant

It is an *inductive invariant* if for *any* state  $S$  (including unreachable ones!):

- Base case: It holds if  $S$  is an initial state
- Induction step: if It holds in  $S$ , then It also holds in any states reachable from  $S$  in one step

Note, an ordinary invariant only needs to hold in all reachable states

It is useful if it implies an invariant that we are interested in (mutual exclusion in this case)

# Peterson's Reconsidered

- Mutual Exclusion can be implemented with LOAD and STORE instructions to access shared memory
  - 3 STOREs and 1 or more LOADs
- 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  $s, p$** 
  - sets  $p$  to the (old) value of  $s$
  - sets  $s$  to **True**
    - i.e., **LOAD  $s$ , STORE  $p$ , STORE  $s$**
- Entire operation is *atomic*
  - other machine instructions cannot interleave

# Enter *Interlock Instructions*

- Machine instructions that do multiple shared memory accesses atomically

- E.g., **TestAndSet**  $s, p$ 
  - sets  $p$  to the (old) value of  $s$
  - sets  $s$  to **True**
    - i.e., **LOAD**  $s$ , **STORE**  $p$ , **STORE**  $s$

```
def tas( $s, p$ ):  
    atomic:  
        ! $p = !s$ ;  
        ! $s = \text{True}$ ;  
    ;  
;
```

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

```
def tas( $s, p$ ):  
    atomic:  
         $!p = !s$ ;  
         $!s = \text{True}$ ;  
    ;  
;
```

$s$  and  $p$  are pointers, thus  $\text{tas}(s, p)$  can be used with any two shared variables:  
 $\text{tas}(?x, ?y)$  or  $\text{tas}(?q, ?r)$

# Critical Sections with TAS

```
1  const N = 3;
2
9  def process(self):
10     while choose({ False, True }):
11         # Enter critical section
12         while private[self]:
13             tas(?shared, ?private[self]);
14         ;
15
16         # Critical section
17         @cs: assert (not private[self]) and
18                 (atLabel.cs == dict{ nametag(): 1 })
19         ;
20
21         # Leave critical section
22         private[self] = True;
23         shared = False;
24     ;
25 ;
26 shared = False;
27 private = [ True, ] * N;
28 for i in {0..N-1}:
29     spawn process(i);
30 ;
```

*number of processes*

**“spinlock”**

$process(self)@cs \Rightarrow \neg private[self]$

*private[ i ] belongs to process( i )*

Figure 8.1: [[code/spinlock.hny](#)] Mutual Exclusion using a “spinlock” based on test-and-set.

# Two essential invariants

1.  $\forall i: process(i)@cs \Rightarrow \neg private[i]$
2. at most 1 of *shared* and *private[i]* is False

1. Obvious
2. Easy proof by induction

both can also be checked by Harmony (see book)

If at most one *private[i]* can be False, then at most one *process(i)* can be @cs

# Checking the second invariant

```
1  import spinlock;
2
3  def checkInvariant():
4      let sum = 0;
5          if not shared:
6              sum = 1;
7          ;
8          for i in {0..N-1} such that not private[i]:
9              sum += 1;
10         ;
11         result = sum <= 1;
12     ;
13 ;
14 def invariantChecker():
15     assert checkInvariant();
16 ;
17 spawn invariantChecker();
```

Check that at most one of *shared* and *private*[*i*] is False

← check it here, atomically

Figure 8.2: [[code/spinlockInv.hny](#)] Checking invariants.

**assert** statements are evaluated atomically

# Checking the second invariant

```
1  import spinlock;
2
3  def checkInvariant():
4      let sum = 0;
5          if not shared:
6              sum = 1;
7          ;
8          for i in {0..N-1} such that not private[i]:
9              sum += 1;
10         ;
11         result = sum <= 1;
12     ;
13 ;
14 def invariantChecker():
15     assert checkInvariant();
16 ;
17 spawn invariantChecker();
```

*Riddle: this code checks the invariant only once, and yet it checks the invariant at every state.*

*How can that be?*

Figure 8.2: [[code/spinlockInv.hny](#)] Checking invariants.

# “Locks”

Best understood as “baton passing”

- At most one process, or *shared*, can “hold” **False**



# Locks in the “synch” module

```
1  def tas(lk):
2      atomic:
3          result = !lk;
4          !lk = True;
5      ;
6  ;
7  def Lock():
8      result = False;
9  ;
10 def lock(lk):
11     await not tas(lk);
12 ;
13 def unlock(lk):
14     !lk = False;
15 ;
```

Observation: *private[i]* does not need to be a shared variable. Just return the old value

Figure 9.2: [[modules/synch.hny](#)] The Lock interface and implementation in the synch module.

# “Ghost” state

- No longer have *private*[*i*]
- Instead:
  - We say that a lock is *held* or *owned* by a process
- The invariants become:
  1.  $P@cs \Rightarrow P$  holds the lock
  2. at most one process can hold the lock

# Using locks from the sync module

```
1  import sync;
2
3  def process(self):
4      lock(?countlock);
5      count = count + 1;
6      unlock(?countlock);
7      done[self] = True;
8
9  def main(self):
10     await all(done);
11     assert count == 2, count;
12
13     count = 0;
14     countlock = Lock();
15     done = [ False, False ];
16     spawn process(0);
17     spawn process(1);
18     spawn main();
```



*import the sync module*

Figure 9.3: [[code/UpLock.hny](#)] Program of Figure 3.1 fixed with a lock.

# Using locks from the sync module

```
1  import sync;
2
3  def process(self):
4      lock(?countlock);
5      count = count + 1;
6      unlock(?countlock);
7      done[self] = True;
8
9  def main(self):
10     await all(done);
11     assert count == 2, count;
12
13     count = 0;
14     countlock = Lock();
15     done = [ False, False ];
16     spawn process(0);
17     spawn process(1);
18     spawn main();
```



*import the sync module*



*initialize lock*

Figure 9.3: [[code/UpLock.hny](#)] Program of Figure 3.1 fixed with a lock.

# Using locks from the sync module

```
1  import sync;
2
3  def process(self):
4      lock(?countlock);
5      count = count + 1;
6      unlock(?countlock);
7      done[self] = True;
8  ;
9  def main(self):
10     await all(done);
11     assert count == 2, count;
12 ;
13 count = 0;
14 countlock = Lock();
15 done = [ False, False ];
16 spawn process(0);
17 spawn process(1);
18 spawn main();
```

*import the sync module*

*enter critical section*

*initialize lock*

Figure 9.3: [[code/UpLock.hny](#)] Program of Figure 3.1 fixed with a lock.

# Using locks from the sync module

```
1  import synch;
2
3  def process(self):
4      lock(?countlock);
5      count = count + 1;
6      unlock(?countlock);
7      done[self] = True;
8  ;
9  def main(self):
10     await all(done);
11     assert count == 2, count;
12 ;
13 count = 0;
14 countlock = Lock();
15 done = [ False, False ];
16 spawn process(0);
17 spawn process(1);
18 spawn main();
```

*import the sync module*

*enter critical section*

*?countlock is the address of countlock  
process self holds countlock*

*initialize lock*

Figure 9.3: [[code/UpLock.hny](#)] Program of Figure 3.1 fixed with a lock.

# Using locks from the sync module

```
1  import sync;
2
3  def process(self):
4      lock(?countlock);
5      count = count + 1;
6      unlock(?countlock);
7      done[self] = True;
8
9  def main(self):
10     await all(done);
11     assert count == 2, count;
12
13     count = 0;
14     countlock = Lock();
15     done = [ False, False ];
16     spawn process(0);
17     spawn process(1);
18     spawn main();
```

*import the sync module*

*enter critical section*

*exit critical section*

*initialize lock*

Figure 9.3: [[code/UpLock.hny](#)] Program of Figure 3.1 fixed with a lock.

# Spinlocks and Time Sharing

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

# Context switching in Harmony

- Harmony allows contexts to be saved and restored
- $r = \mathbf{stop}$  *list*
  - stops the current process and places its context at the end of the given list
- $\mathbf{go}$  *context*  $r$ 
  - adds a process with the given context to the bag of processes. Process resumes from **stop** expression, returning  $r$


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

```
1  import list;
2
3  def Lock():
4      result = dict{ .locked: False, .suspended: [ ] };
5
6  ;
7  def lock(lk):
8      atomic:
9          if lk→locked:
10             stop lk→suspended;
11             assert lk→locked;
12         else:
13             lk→locked = True;
14
15     ;
16
17 def unlock(lk):
18     atomic:
19         if lk→suspended == [ ]:
20             lk→locked = False;
21         else:
22             go (head(lk→suspended)) ();
23             lk→suspended = tail(lk→suspended);
24
25     ;
```

Figure 9.4: [[modules/syncS.hny](#)] The Lock interface in the `syncS` module uses suspension.

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

```
1  import list;
2
3  def Lock():
4      result = dict{ .locked: False, .suspended: [ ] };
5
6  ;
7  def lock(lk):
8      atomic:
9          if lk->locked:
10             stop lk->suspended;
11             assert lk->locked;
12         else:
13             lk->locked = True;
14
15     ;
16
17 def unlock(lk):
18     atomic:
19         if lk->suspended == [ ]:
20             lk->locked = False;
21         else:
22             go (head(lk->suspended)) ();
23             lk->suspended = tail(lk->suspended);
24
25     ;
```



*lk*->locked is short for (!*lk*).locked (cf. C, C++)

Figure 9.4: [[modules/syncS.hny](#)] The Lock interface in the `syncS` module uses suspension.

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

```
1  import list;
2
3  def Lock():
4      result = dict{ .locked: False, .suspended: [ ] };
5
6  ;
7  def lock(lk):
8      atomic:
9          if lk→locked:
10             stop lk→suspended;
11             assert lk→locked;
12         else:
13             lk→locked = True;
14
15     ;
16
17 def unlock(lk):
18     atomic:
19         if lk→suspended == [ ]:
20             lk→locked = False;
21         else:
22             go (head(lk→suspended)) ();
23             lk→suspended = tail(lk→suspended);
24
25     ;
```

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

# Choosing modules in Harmony

- “synch” is the (default) module that has the TAS version of 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`

- “sync” tends to be faster than “syncS”
  - smaller state graph

# Atomic section $\neq$ Critical Section

<b>Atomic Section</b>	<b>Critical Section</b>
only one process can execute	multiple process can execute concurrently, just not within a critical section
rare programming paradigm	ubiquitous: locks available in many mainstream programming languages
good for implementing interlock instructions	good for building concurrent data structures

# Building a Concurrent Queue

- $q = \text{Qnew}()$ : allocate a new queue
- $\text{Qenqueue}(q, v)$ : add  $v$  to the tail of queue  $q$
- $r = \text{Qdequeue}(q)$ : returns  $r = ()$  if  $q$  is empty or  $r = (v,)$  (a singleton tuple) if  $v$  was at the head of the queue

# Queue Test Program Example


```
1  import queue;
2
3  def sender(q, v):
4      Qenqueue(q, v);
5  ;
6  def receiver(q):
7      let done = False:
8          while not done:
9              let v = Qdequeue(q):
10                 done = v == ();
11                 assert done or (v[0] in { 1, 2 });
12             ;
13         ;
14     ;
15 ;
16
17 queue = Qnew();
18 spawn sender(?queue, 1);
19 spawn sender(?queue, 2);
20 spawn receiver(?queue);
21 spawn receiver(?queue);
```

enqueue *v* onto *q*

dequeue until queue *q* is empty


Figure 10.1: [\[code/queuetest.hny\]](#) Test program for a concurrent queue.

# Queue implementation, v1




```
1  import alloc;
2
3  def Qnew():
4      result = dict{ .head: None, .tail: None, .lock: Lock() };
5  ;
6  def Qenqueue(q, v):
7      let node = malloc(dict{ .value: v, .next: None }):
8          lock(?q->lock);
9          if q->head == None:
10             q->head = q->tail = node;
11         else:
12             q->tail->next = node;
13             q->tail = node;
14         ;
15         unlock(?q->lock);
16     ;
17 ;
```

# Queue implementation, v1



```
1 import alloc;
2
3 def Qnew():
4     result = dict{ .head: None, .tail: None, .lock: Lock() };
5 ;
6 def Qenqueue(q, v):
7     let node = malloc(dict{ .value: v, .next: None }):
8         lock(?q->lock);
9         if q->head == None:
10            q->head = q->tail = node;
11        else:
12            q->tail->next = node;
13            q->tail = node;
14        ;
15        unlock(?q->lock);
16    ;
17 ;
```


# Queue implementation, v1



```
18 def Qdequeue(q):
19     lock(?q→lock);
20     let node = q→head:
21         if node == None:
22             result = ();
23         else:
24             result = (node→value,);
25             q→head = node→next;
26             if q→head == None:
27                 q→tail = None;
28             ;
29             free(node);
30         ;
31     ;
32     unlock(?q→lock);
33 ;
```

Figure 10.2: [[code/queue.hny](#)] A basic concurrent queue data structure.

# Queue implementation, v1



```
18 def Qdequeue(q):
19     lock(?q→lock);
20     let node = q→head:
21         if node == None:
22             result = ();
23         else:
24             result = (node→value,);
25             q→head = node→next;
26             free(node);
27     ;
28 ;
29     unlock(?q→lock);
30 ;
```


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

Figure 10.2: [\[code/queue.hny\]](#) A basic concurrent queue data structure.

# 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
  - ...

# Better concurrent queue: 2 locks



```
1  import alloc;
2
3  def Qnew():
4      let dummy = malloc(dict{ .value: (), .next: None }):
5          result = dict{ .head: dummy, .tail: dummy, .hdlock: Lock(), .tllock: Lock() };
6      ;
7  ;
8  def Qenqueue(q, v):
9      let node = malloc(dict{ .value: v, .next: None }):
10         lock(?q->tllock);
11         q->tail->next = node;
12         q->tail = node;
13         unlock(?q->tllock);
14     ;
15 ;
```

# Better concurrent queue: 2 locks



```
16 def Qdequeue(q):
17     lock(?q→hdlock);
18     let dummy = q→head
19     let node = dummy→next:
20         if node == None:
21             result = ();
22         else:
23             free(dummy);
24             result = (node→value,);
25             q→head = node;
26     ;
27 ;
28 unlock(?q→hdlock);
29 ;
```

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

Figure 10.3: [[code/queueMS.hny](#)] A queue with separate locks

# How to get more concurrency?

**Idea:** allow multiple read-only operations to execute concurrently

- In many cases, reads are much more frequent than writes

→ reader/writer lock

Either:


- multiple readers, or
- a single writer

*thus not:*

- *a reader and a writer, nor*
- *multiple writers*

# Reader/writer lock state diagram

- (1) process not in c.s. terminates
- (2) process enters r.c.s.
- (3) process leaves r.c.s.
- (4) process enters w.c.s.
- (5) process leaves w.c.s.



# Reader/writer lock interface:

- **acquire\_rlock()**
  - get a read lock. Multiple processes can have the read lock simultaneously, but no process can have a write lock simultaneously
- **release\_rlock()**
  - release a read lock. Other processes may still have the read lock. When the last read lock is released, a write lock may be acquired
- **acquire\_wlock()**
  - acquire the write lock. Only one process can have a write lock, and if so no process can have a read lock
- **release\_wlock()**
  - release the write lock. Allows other processes to either get a read or write lock

# R/W lock, Implementation #1

- Uses a single ordinary lock and two integers to count #readers and #writers

```
37     rwlock = Lock();  
38     nreaders = 0;  
39     nwriters = 0;
```

Figure 11.1: [[code/RW.hny](#)] Busy-Waiting Reader/Writer Lock

## Invariants:

- if  $n$  readers in the critical section, then  $nreaders \geq n$
- if  $n$  writers in the critical section, then  $nwriters \geq n$
- $(nreaders \geq 0 \wedge nwriters = 0) \vee (nreaders = 0 \wedge 0 \leq nwriters \leq 1)$

*rwlock* protects the *nreaders/nwriters* variables, not the critical section!

# R/W lock, Implementation #1

```
1  import synch;
2
3  def acquire_rlock():
4      let blocked = True:
5          while blocked:
6              lock(?rwlock);
7              if nwriters == 0:
8                  nreaders += 1;
9                  blocked = False;
10             ;
11             unlock(?rwlock);
12         ;
13     ;
14 ;
15 def release_rlock():
16     lock(?rwlock);
17     nreaders -= 1;
18     unlock(?rwlock);
19 ;
```

# R/W lock, Implementation #1

```
1  import synch;
2
3  def acquire_rlock():
4      let blocked = True:
5          while blocked:
6              lock(?rwlock);
7              if nwriters == 0:
8                  nreaders += 1;
9                  blocked = False;
10             ;
11             unlock(?rwlock);
12         ;
13     ;
14 ;
15 def release_rlock():
16     lock(?rwlock);
17     nreaders -= 1;
18     unlock(?rwlock);
19 ;
```

“busy wait” (i.e., spin) until no writer in the critical section

# R/W lock, Implementation #1

```
20     def acquire_wlock():
21         let blocked = True:
22             while blocked:
23                 lock(?rwlock);
24                 if (nreaders + nwriters) == 0:
25                     nwriters = 1;
26                     blocked = False;
27                 ;
28                 unlock(?rwlock);
29             ;
30         ;
31     ;
32     def release_wlock():
33         lock(?rwlock);
34         nwriters = 0;
35         unlock(?rwlock);
```

“busy wait”  
until no other  
process in the  
critical section

# R/W Locks: test for mutual exclusion

```
1  import RW;
2
3  def process():
4      while choose({ False, True }):
5          if choose({ .read, .write }) == .read:
6              acquire_rlock();
7              @rcs: assert atLabel.wcs == dict{}; ← no writer
8              release_rlock();
9          else:                                # .write
10             acquire_wlock();
11             @wcs: assert (atLabel.wcs == dict{ nametag(): 1 }) and
12                 (atLabel.rcs == dict{})
13             ;
14             release_wlock();
15         ;
16     ;
17 ;
18 for i in {1..4}:
19     spawn process();
20 ;
```





Figure 11.2: [[code/RWtest.hny](#)] Test code for Figure 11.1.

# About *busy waiting*

- ok for multi-core (true) parallelism
- bad for time-sharing (virtual) parallelism

# R/W Lock, implementation #2

- Uses two ordinary locks and an integer that counts the number of readers

```
25     rwlock = Lock();
26     rlock = Lock();
27     nreaders = 0;
```

Figure 12.1: [[code/RWlock.hny](#)]

## Invariants:

- if  $n$  readers in the critical section, then  $nreaders \geq n$
- if a writer in the critical section, then  $nreaders = 0$
- if writer  $W$  in the critical section, then  $W$  holds *rwlock*  
(if some reader in the critical section, the readers collectively hold *rwlock*)

*rlock* protects the *nreaders* variable

# R/W Lock, implementation #2

```
3     def acquire_rlock():
4         lock(?rlock);
5         if nreaders == 0:
6             lock(?rwlock);
7         ;
8         nreaders += 1;
9         unlock(?rlock);
10    ;
11    def release_rlock():
12        lock(?rlock);
13        nreaders -= 1;
14        if nreaders == 0:
15            unlock(?rwlock);
16        ;
17        unlock(?rlock);
18    ;
19    def acquire_wlock():
20        lock(?rwlock);
21    ;
22    def release_wlock():
23        unlock(?rwlock);
24    ;
```

*rlock* protects the *nreaders* variable

first reader acquires *rwlock*

last reader releases *rwlock*

writer acquires *rwlock*

writer releases *rwlock*

# R/W Lock, implementation #2

*nreaders == 0*



Reader R1 holds *rlock* and is waiting for *rwlock*

Reader R2 is waiting for *rlock*

Reader R1 holds *rwlock* and released *rlock*

Reader R2 released *rlock*

Reader R1 leaves but *rwlock* is still “held”



Writer W left the critical section and holds *rwlock*  
Writer W left the critical section

```
3 def acquire_rlock():
4     lock(?rlock);
5     if nreaders == 0:
6         lock(?rwlock);
7     ;
8     nreaders += 1;
9     unlock(?rlock);
10 ;
11 def release_rlock():
12     lock(?rlock);
13     nreaders -= 1;
14     if nreaders == 0:
15         unlock(?rwlock);
16     ;
17     unlock(?rlock);
18 ;
19 def acquire_wlock():
20     lock(?rwlock);
21 ;
22 def release_wlock():
23     unlock(?rwlock);
24 ;
```

# R/W Lock, implementation #2

```
3     def acquire_rlock():
4         lock(?rlock);
5         if nreaders == 0:
6             lock(?rwlock);
7         ;
8         nreaders += 1;
9         unlock(?rlock);
10        ;
11    def release_rlock():
12        lock(?rlock);
13        nreaders -= 1;
14        if nreaders == 0:
15            unlock(?rwlock);
16        ;
17        unlock(?rlock);
18    ;
19    def acquire_wlock():
20        lock(?rwlock);
21    ;
22    def release_wlock():
23        unlock(?rwlock);
24    ;
```

no busy waiting!

both readers and writers  
“block” when they can’t  
enter the critical section

# More testing of reader/writer lock implementations

- Prior test only checks mutual exclusion
- How do you test if the implementation allows multiple readers?
- How do you test if the implementation uses busy waiting or not?

For both, see book