

# Concurrent Programming with Harmony

#### **Robbert van Renesse**





# **Concurrency Lectures Outline**

- What are the problems?
  o no determinism, no atomicity
- What is the solution?
   o some form of mutual exclusion
- How to specify concurrent problems?
   atomic operations
- How to construct correct concurrent code?
   behaviors
- How to test concurrent programs?
  - comparing behaviors

# **Concurrency Lectures Outline**

- How to build Concurrent Data Structures?
   o using locks
- How to wait for some condition?
   using condition variables
- How to deal with deadlock?
   prevention, avoidance, detection
- How to use barrier synchronization?
   improve scalability
- How to make code interrupt-safe?
   o enabling/disabling interrupts



# The problems





# 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
  - with concurrency, this leads to *non-deterministic* interleavings

# Harmony

- A new concurrent programming language

   heavily based on Python syntax to reduce
   learning curve for many
- A new underlying virtual machine

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

### The problem with non-determinism



#### What will happen if you run each?

### The problem with non-determinism

#### sequential concurrent shared = Trueshared = True1 1 2 2 3 def f(): assert shared 3 def f(): assert shared 4 def g(): shared = False 4 def q(): shared = False 5 5 6 6 f() spawn f()7 **g**() 7 spawn g() #states: 2 Schedule thread TO: init() Line 1: Initialize shared to True No issues 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

# The problem with non-atomicity



What will happen if you run each?

# The problem with non-atomicity



### **Race Conditions**

#### = timing dependent error involving shared state

- A schedule is an interleaving of (i.e., total order on) the machine instructions executed by each thread
- Usually, many interleavings are possible
- A race condition occurs when at least one interleaving gives an undesirable result

#### **Race Conditions are Hard to Debug**

- Number of possible interleavings is usually huge
- Bad interleavings, if they exist, may happen only rarely
  - Works 1000x ≠ no race condition
- Timing dependent: small changes hide bugs
   o add print statement → bug no longer seems to happen
- Harmony is designed to help identify such bugs
   model checking!



# State Space and Model Checking





## Harmony Machine Code



# Harmony Virtual Machine State

#### Three parts:

- 1. code (never changes)
- 2. values of the shared variables
- 3. state of each of the running threads
  - PC and stack (aka *context*)

HVM state represents one vertex in a graph of states





initial state













All possible states after one "step"









### Harmony





# 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)
no return, break, continue	various flow control escapes
pointers	object-oriented

# I/O in Harmony?

- Input:
  - o choose expression
    - x = choose({ 1, 2, 3 })
    - allows Harmony to know all possible inputs
  - o 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  $-x = choose(\{1, 2, 3\})$ puts - allows Harm  $\circ$  cons i.en with "-c x=4" flag to harmony ca  $\circ$  Out
  - print x + y
  - **assert** x + y < 10, (x, y)

# Non-determinism in Harmony

#### Three sources:

- 1. choose expressions
- 2. thread interleavings
- 3. interrupts

### Limitation: models must be finite!



28

# Limitation: models must be finite!



- That is, there must be a finite number of states and edges.
- But models are allowed to have cycles.
- Executions are allowed to be unbounded!
- Harmony checks for *possibility* of termination.



# Critical Sections





## Back to our problem...

2 threads updating a shared variable



## Back to our problem...

2 threads updating a shared variable



# Back to our problem...

2 threads updating a shared variable



#### Goals

Mutual Exclusion: 1 thread in a critical section at time Progress: a thread can get in when there is no other thread Fairness: equal chances of getting into CS

... in practice, fairness rarely guaranteed or needed

# Mutual Exclusion and Progress

#### Need both:

 $\circ$  either one is trivial to achieve by itself

#### Specifying Critical Sections in Harmony

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

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

### Specifying Critical Sections in Harmony

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

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



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



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



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



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



Progress: Harmony checks that all thread *can* terminate



## Building a lock is hard





#### Specification vs implementation

- Spec is fine, but we'll need an implementation too
- Sounds like we need a *lock*
- The question is:

#### How does one build a lock?

### First attempt: a naïve lock

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

5

#### First attempt: a naïve lock

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

### First attempt: a naïve lock

```
in_cs = 0
 1
 2
      invariant in_cs in \{0, 1\}
 3
       lockTaken = False
 4
 5
       def thread(self):
 6
 7
             while choose({ False, True }):

    Schedule thread T0: init()

                   # Enter critical section
 8

    Line 1: Initialize in cs to 0

    Line 4: Initialize lockTaken to False

 9
                   await not lockTaken

    Thread terminated

                   lockTaken = True
10

    Schedule thread T3: thread(1)

    Line 7: Choose True

11
                                                                    • Preempted in thread(1) about to store True into lockTaken in line 10
12
                   atomically in_cs += 1

    Schedule thread T2: thread(0)

                   # Critical section

    Line 7: Choose True

13

    Line 10: Set lockTaken to True (was False)

14
                   atomically in_cs -= 1

    Line 12: Set in cs to 1 (was 0)

    Preempted in thread(0) about to execute atomic section in line 14

15

    Schedule thread T3: thread(1)

                   # Leave critical section
16

    Line 10: Set lockTaken to True (unchanged)

17
                   lockTaken = False

    Line 12: Set in_cs to 2 (was 1)

    Preempted in thread(1) about to execute atomic section in line 14

18

    Schedule thread T1: invariant()

       spawn thread(0)

    Line 2: Harmony assertion failed

19
20
       spawn thread(1)
```

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

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

```
in_cs = 0
1
2
    invariant in_cs in { 0, 1 }
 3
    flags = [ False, False ]
4
5
6
    def thread(self):
        while choose({ False, True }):
 7
             # Enter critical section
 8
                                              show intent to enter critical section
9
             flags[self] = True
             await not flags [1 - self]
10
                                                wait until there's no one else
11
12
             atomically in_cs += 1
             # Critical section
13
14
             atomically in_cs -= 1
15
16
             # Leave critical section
             flags[self] = False
17
18
19
    spawn thread(0)
20
    spawn thread(1)
```

```
in_cs = 0
 1
                                                         Summary: some execution cannot terminate
     invariant in_cs in { 0, 1 }
 2
                                                         Here is a summary of an execution that exhibits the issue:
 3
 4
     flags = [ False, False ]

    Schedule thread T0: init()

    Line 1: Initialize in_cs to 0

 5

    Line 4: Initialize flags to [False, False]

     def thread(self):
 6

    Thread terminated

           while choose({ False, True }):
 7

    Schedule thread T1: thread(0)

                 # Enter critical section
 8

    Line 7: Choose True

                 flags[self] = True

    Line 9: Set flags[0] to True (was False)

 9

    Preempted in thread(0) about to load variable flags[1] in line 10

                 await not flags [1 - self]
10

    Schedule thread T2: thread(1)

11

    Line 7: Choose True

12
                 atomically in_cs += 1

    Line 9: Set flags[1] to True (was False)

                 # Critical section
13

    Preempted in thread(1) about to load variable flags[0] in line 10

                 atomically in_cs -= 1
14
                                                         Final state (all threads have terminated or are blocked):
15
                 # Leave critical section
16

    Threads:

17
                 flags[self] = False

    T1: (blocked) thread(0)

                                                                      about to load variable flags[1] in line 10
18

    T2: (blocked) thread(1)

19
      spawn thread(0)
                                                                      about to load variable flags[0] in line 10
20
      spawn thread(1)
```

```
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 }):
             # Enter critical section
 8
             turn = 1 - self
 9
10
             await turn == self
11
             atomically in_cs += 1
12
             # Critical section
13
14
             atomically in_cs -= 1
15
             # Leave critical section
16
17
18
    spawn thread(0)
19
    spawn thread(1)
```

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

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

```
Summary: some execution cannot terminate
      in cs = 0
 1
 2
      invariant in_cs in { 0, 1 }
                                                      Here is a summary of an execution that exhibits the issue:
 3

    Schedule thread T0: init()

 4
     turn = 0

    Line 1: Initialize in_cs to 0

 5

    Line 4: Initialize turn to 0

 6
      def thread(self):

    Thread terminated

 7
           while choose({ False, True }):

    Schedule thread T2: thread(1)

                # Enter critical section

    Line 7: Choose False

 8

    Thread terminated

 9
                turn = 1 - self

    Schedule thread T1: thread(0)

                 await turn == self
10

    Line 7: Choose True

11

    Line 9: Set turn to 1 (was 0)

12
                atomically in_cs += 1

    Preempted in thread(0) about to load variable turn in line 10

                # Critical section
13
14
                atomically in_cs -= 1
                                                      Final state (all threads have terminated or are blocked):
15

    Threads:

16
                # Leave critical section

    T1: (blocked) thread(0)

17
                                                                   about to load variable turn in line 10
18
      spawn thread(0)

    T2: (terminated) thread(1)

      spawn thread(1)
19
```

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

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

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











Proving a concurrent program correct





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)







#### 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]
# Example

- Theorem: if T is in the critical section, then flags[T] = True
- Base case: true because initially T is not in the critical section and False implies anything
- Induction: easy to show (using Hoare logic) because flags[T] can only be changed by T itself



#### Data Races





#### Peterson's Reconsidered

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

```
in cs = 0
 2
    invariant in_cs in { 0, 1 }
                                     Loads and Stores are atomic
    sequential flags, turn
    flags = [ False, False ]
    turn = choose(\{0, 1\})
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)
```

#### For example

- CPU with 16-bit architecture
- 32-bit integer variable *x* stored in memory in two adjacent locations (aligned on word boundary)
- Initial value is 0
- Thread 1 writes FFFFFF to x (requires 2 STOREs)
- Thread 2 reads *x* (requires 2 LOADs)
- What are the possible values that thread 2 will read?

#### For example

- CPU with 16-bit architecture
- 32-bit integer variable *x* stored in memory in two adjacent locations (aligned on word boundary)
- Initial value is 0
- Thread 1 writes FFFFFF to x (requires 2 STOREs)
- Thread 2 reads *x* (requires 2 LOADs)
- What are the possible values that thread 2 will read?
   FFFFFFF
  - o 0000000
  - FFFF0000
  - 0000FFFF

#### Other examples

- In Python, integers are arbitrary precision

   that is, each integer variable is a complex data structure, and an operation may require multiple loads and stores
- Suppose your C compiler decides to pack multiple bits in a single word
  - E.g., flags[0], flags[1], and turn
  - $\,\circ\,$  Then setting a bit involves a load and a store

#### Concurrent memory access

- Hardware may also cause problems in efforts to improve performance

  buffering of writes
  caching of reads
  out-of-order execution
- Because of all these issues, programming languages will typically leave the outcome of concurrent operations to a variable *undefined*
  - $\circ$  if at least one of those operations is a store

#### Data Race

- When two or more threads wish to access the same variable at the same time
- And at least one access is a STORE
- Then the semantics of the outcome is *undefined* That is:
  - The value stored in the variable is undefined
  - The value loaded (if any) is undefined
  - Undefined means random (or worse, like a crash)

# Harmony "sequential" statement

- sequential turn, flags in Peterson's
- 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:
  - $\circ$  volatile int *x*;
  - All accesses to volatile variables are sequentially consistent (but not whole program)
- 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
  - $\circ$  e.g., rules out caching

#### Peterson's Reconsidered Again

- Peterson's algorithm is correct with atomic LOAD and STORE instructions

   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*



# Specifying a lock





# Back to basics: specifying a lock

- What does a lock *do* exactly?
- What if we want more than one?

#### 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* and it *points to* x
- Finally, **!p** refers to the value of **x**



# Specifying a lock

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

# Specifying a lock



#### Critical Section using a lock

```
from synch import Lock, acquire, release
1
 2
3
    shared = 0
    thelock = Lock()
4
5
    def f():
6
 7
        acquire(?thelock)
        shared += 1
8
9
        release(?thelock)
10
    spawn f()
11
    spawn f()
12
13
    finally shared == 2
14
```

## Critical Section using a lock



#### "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 ⇒ T holds the lock
  2. at most one thread can hold the lock
  Together guarantee mutual exclusion

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



# Implementing a lock





#### Implementing a lock

We saw that it is hard and inefficient to implement a lock with just LOAD and STORE instructions

#### Enter Interlock Instructions

Machine instructions that do multiple shared memory accesses atomically

- e.g., test\_and\_set s o 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

# Lock implementation ("spinlock")

```
def test_and_set(s) returns result:
1
                                                specification of the CPU's
test_and_set functionality
 2
         atomically:
              result = !s
 3
4
              !s = True
 5
                                                specification of the CPU's
6
     def atomic_store(p, v):
                                                atomic store functionality
 7
         atomically |p| = v
8
9
     def Lock() returns result:
         result = False
10
11
     def acquire(lk):
12
                                                lock implementation
13
         while test_and_set(lk):
14
              pass
15
16
     def release(lk):
         atomic_store(lk, False)
17
```

# **Specification vs Implementation**

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

```
def test_and_set(s) returns result:
 1
         atomically:
 2
             result = !s
 3
             !s = True
 4
 5
    def atomic_store(p, v):
 6
         atomically |p| = v
 7
 8
    def Lock() returns result:
9
         result = False
10
11
12
    def acquire(lk):
         while test_and_set(lk):
13
14
             pass
15
    def release(lk):
16
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:
  o when there is no pre-emption?

o 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:
   o when there is no pre-emption?
  - can cause all threads to get stuck while one is trying to obtain a spinlock
  - o 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:
   o when there is no pre-emption?
  - can cause all threads to get stuck while one is trying to obtain a spinlock
  - o 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

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

## Locks using stop and go

```
def Lock() returns result:
1
 2
        result = { .acquired: False, .suspended: ] }
 3
    def acquire(lk):
 4
        atomically:
 5
 6
  Similar to a Linux "futex": if there is no contention
9
10
  (hopefully the common case) acquire() and release() are
   cheap. If there is contention, they involve a context switch.
11
12
13
        aconicculty.
            assert lk->acquired
14
            if lk->suspended == []:
15
                lk->acquired = False
16
17
            else:
18
                go (lk->suspended[0]) ()
                del lk->suspended[0]
19
```

# 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 ≠ Critical Section

Atomic Section	<b>Critical Section</b>
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

#### Data Race ≠ Race Condition

- 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 is undefined
- A *Race Condition* occurs when the correctness of the program depends on ordering of variable access
  - Race Condition can happen without a Data Race

#### Data Race ≠ Race Condition

- Data Race: Harmony can automatically detect these because Harmony enumerates all behaviors and fails if there is undefined behavior
- Race Condition: Harmony can only detect these if you tell Harmony what it is that you want using assert, invariant, or finally
  - or by explicitly enumerating the correct behaviors, as we'll see later



#### Demo Time




= 0

Demo 1: data race

Demo 2: no data race

Demo 3: same semantics as Demo 2:

x = 0 def f(): x = x + 1 def g(): x = x + 1 spawn f() spawn g()

```
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()
```

sequential x
x = 0
def f():
 x = x + 1
def g():
 x = x + 1
spawn f()
spawn g()

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()
```

## Demo 6: spec of what we want



## Demo 7: implementation using critical section

from synch import Lock,	acquire,	release
finally x == 2		
x = 0 thelock = Lock()		
<pre>def f():     acquire(?thelock)     x += 1     release(?thelock)</pre>		
<pre>def g():     acquire(?thelock)     x += 1     release(?thelock)</pre>		
spawn f() spawn g()		

Demo 8: broken implementation using two critical sections

```
from synch import Lock, acquire, release
finally x == 2
\mathbf{x} = \mathbf{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()
```



# Concurrent Data Structure Consistency





### 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:
  - $\circ$  lock not held  $\Rightarrow$  data structure consistent
- Or equivalently:
  - $\circ$  data structure inconsistent  $\Rightarrow$  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

#### How important are concurrent queues?

- Answer: all important
  - any resource that needs scheduling
    - CPU run queue
    - disk, network, printer waiting queue
    - lock waiting queue
  - $\circ$  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
  - $\circ$  inter-process communication
    - Posix pipes:
      - cat file | tr a-z A-Z | grep RVR
  - actor-based concurrency
  - 0...

## Specifying a concurrent queue

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

## Specifying a concurrent queue

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

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

```
Sequential
```

Concurrent

## Example of using a queue



## Specifying a concurrent queue

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

#### Not a good implementation because

- operations are O(n)
- code uses **atomically** compiler cannot generate code



# Implementing a concurrent queue







```
from synch import Lock, acquire, release
 1
     from alloc import malloc, free
 2
 3
 4
    def Queue() returns empty:
 5
         empty = { .head: None, .tail: None, .lock: Lock() }
 6
    def put(q, v):
 7
 8
         let node = malloc({ .value: v, .next: None }):
 9
              acquire(?q->lock)
10
              if a->tail == None:
                  q \rightarrow tail = q \rightarrow head = node
11
              else:
12
13
                  q->tail->next = node
                  q \rightarrow tail = node
14
              release(?q->lock)
15
```







```
from synch import Lock, acquire, release
 1
     from alloc import malloc, free
 2
 3
 4
    def Queue() returns empty:
 5
         empty = { .head: None, .tail: None, .lock: Lock() }
 6
    def put(q, v):
 7
                                                                   allocate node
         let node = malloc({ .value: v, .next: None }):
 8
              acquire(?q->lock)
 9
10
              if a->tail == None:
                  q \rightarrow tail = q \rightarrow head = node
11
              else:
12
13
                  q->tail->next = node
                  q \rightarrow tail = node
14
              release(?q->lock)
15
```



```
from synch import Lock, acquire, release
1
    from alloc import malloc, free
2
 3
4
    def Queue() returns empty:
5
         empty = { .head: None, .tail: None, .lock: Lock() }
6
    def put(q, v):
7
         let node = malloc({ .value v next. None }):
8
                                           grab lock
             acquire(?q->lock)
9
10
             if a->tail == None:
                  q \rightarrow tail = q \rightarrow head = node
11
              else:
12
13
                  q->tail->next = node
                  q \rightarrow tail = node
14
              release(?q->lock)
15
```



```
from synch import Lock, acquire, release
 1
     from alloc import malloc, free
 2
 3
 4
    def Queue() returns empty:
 5
         empty = { .head: None, .tail: None, .lock: Lock() }
 6
    def put(q, v):
 7
 8
         let node = malloc({ .value: v, .next: None }):
 9
              acquire(?q->lock)
10
              if a->tail == None:
                  q->tail = q->head = node
11
                                                            the hard stuff
              else:
12
13
                  q \rightarrow tail \rightarrow next = node
                  q \rightarrow tail = node
14
              release(?q->lock)
15
```



```
from synch import Lock, acquire, release
 1
     from alloc import malloc, free
 2
 3
 4
    def Queue() returns empty:
 5
         empty = { .head: None, .tail: None, .lock: Lock() }
 6
    def put(q, v):
 7
 8
         let node = malloc({ .value: v, .next: None }):
              acquire(?q->lock)
 9
              if q->tail == None:
10
                  q \rightarrow tail = q \rightarrow head = node
11
              else:
12
13
                  q->tail->next = node
                  q \rightarrow tail = node
14
                                          release lock
              release(?q->lock)
15
```



17	<pre>def get(q) returns next:</pre>
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	<pre>free(node)</pre>
28	release(?q->lock)











```
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,
 7
                              .hdlock: Lock(), .tllock: Lock() }
8
9
    def put(q, v):
10
         let node = malloc({ .value: v, .next: None }):
11
             acquire(?q->tllock)
             atomic_store(?q->tail->next, node)
12
13
             q \rightarrow tail = node
14
             release(?q->tllock)
```



```
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 }):
             empty = { .head: dummy, .tail: dummy,
 6
 7
                              .hdlock: Lock(), .tllock: Lock() }
8
9
    def put(q, v):
         let node = malloc({ .value: v, .next: None }):
10
11
             acquire(?q->tllock)
             atomic_store(?q->tail->next, node)
12
                                                       atomically q->tail->next = node
13
             q \rightarrow tail = node
14
             release(?q->tllock)
```



16	<pre>def get(q) returns next:</pre>	
17	acquire(?q->hdlock)	
18	let dummy = $q$ ->head	
19	<pre>let node = atomic_load(?dummy-&gt;next):</pre>	
20	if node == None:	
21	next = None	
22	release(?q->hdlock)	
23	else:	
24	next = node->value	
25	q->head = node	
26	release(?q->hdlock)	
27	<pre>free(dummy)</pre>	



16	<pre>def get(q) returns next:</pre>		
17	acquire(?q->hdlock)		
18	<pre>let dummy = q-&gt;head</pre>		
19	<pre>let node = atomic_load(?dumm)</pre>	y->next):	
20	if node == None:		
21	next = None	No contention for concurrent	
22	release(?q->hdlock)	enqueue and dequeue operations!	
23	else:	$\rightarrow$ more concurrency $\rightarrow$ faster	
24	next = node->value		
25	q->head = node		
26	release(?q->hdlock)		
27	<pre>free(dummy)</pre>		



16	<pre>def get(q) returns next:</pre>	
17	acquire(?q->hdlock)	
18	<pre>let dummy = q-&gt;head</pre>	
19	<pre>let node = atomic_load(?dumm</pre>	ny->next):
20	<pre>if node == None:</pre>	No contention for concurrent
21	next = None	No contention for concurrent
22	release(?q->hdlock)	enqueue and dequeue operations!
23	else:	$\rightarrow$ more concurrency $\rightarrow$ faster
24	next = node->value	
25	$q \rightarrow head = node$	
26	release(?q->hdlock)	
27	free(dummy)	
Needs to avoid data race on		

 $dummy \rightarrow next$  when queue is empty



# Fine-Grained Locking





## 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



#### Sorted Linked List with Lock per Node

```
.value
                                                     .value
                   -\infty
                                                                          \infty
                                                                                     None
                                   .next
                                                                        .next -
                 .next
                                                     .next
    from synch import Lock, acquire, release
 1
    from alloc import malloc, free
 2
 3
 4
    def _node(v, n) returns node: # allocate and initialize
 5
         node = malloc({ .lock: Lock(), .value: v, .next: n })
 6
 7
    def _find(lst, v) returns pair:
                                             Helper routine to find and lock two
         var before = 1st
 8
                                             consecutive nodes before and after such that
         acquire(?before->lock)
 9
                                             before \rightarrow value < v \leq after \rightarrow value
         var after = before->next
10
11
         acquire(?after->lock)
12
         while after->value < (0, v):
13
             release(?before->lock)
             before = after
14
             after = before -> next
15
16
             acquire(?after->lock)
17
         pair = (before, after)
18
19
    def SetObject() returns object:
         object = _node((-1, None), _node((1, None), None))
20
```

#### Sorted Linked List with Lock per Node

```
.value
                                                    .value
                  -\infty
                                                                       \infty
                                                    .next
                                                                     .next -
                                                                                  None
                                  .next
                 .next
    from synch import Lock, acquire, release
 1
    from alloc import malloc, free
 2
 3
 4
    def _node(v, n) returns node: # allocate and initialize
 5
        node = malloc({ .lock: Lock(), .value: v, .next: n })
 6
 7
    def _find(lst, v) returns pair:
                                            Helper routine to find and lock two
        var before = 1st
 8
                                            consecutive nodes before and after such that
        acquire(?before->lock)
 9
                                            before \rightarrow value < v \leq after \rightarrow value
        var after = before->next
10
11
        acquire(?after->lock)
12
        while after->value < (0, v):
                                                  Hand-over hand locking
13
            release(?before->lock)
            before = after
14
                                                  (good for data structures
            after = before -> next
15
16
            acquire(?after->lock)
                                                  without cycles)
17
        pair = (before, after)
18
19
    def SetObject() returns object:
        object = _node((-1, None), _node((1, None), None))
20
```
#### Sorted Linked List with Lock per Node

```
def insert(lst, v):
    let before, after = _find(lst, v):
        if after->value != (0, v):
            before->next = _node((0, v), after)
        release(?after->lock)

def remove(lst, v):
    let before, after = _find(lst, v):
        if after->value == (0, v):
            before->next = after->next
            free(after)
        else:
            release(?after->lock)
```

```
release(?before->lock)
```

```
def contains(lst, v) returns present:
    let before, after = _find(lst, v):
        present = after->value == (0, v)
        release(?after->lock)
        release(?before->lock)
```

#### Sorted Linked List with Lock per Node

```
def insert(lst, v):
    let before, after = _find(lst, v):
        if after->value != (0, v):
            before->next = _node((0, v), after)
        release(?after->lock)
        release(?before->lock)
```

```
def remove(lst, v):
    let before, after = _find(lst, v):
        if after->value == (0, v):
            before->next = after->next
            free(after)
        else:
            release(?after->lock)
        release(?before->lock)
```

```
def contains(lst, v) returns present:
    let before, after = _find(lst, v):
        present = after->value == (0, v)
        release(?after->lock)
        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
- 0...
- 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

#### How do we capture behaviors?

• And what is a behavior?

#### Life of an atomic operation

process invokes operation happens process resumes operation atomically with result



# Concurrency and Overlap

Is the following a possible scenario?

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

# Concurrency and Overlap

Is the following a possible scenario?

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

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

## 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. 3. customer Y orders burger, order ends up with cook 2
- cook 1 was busy with something else, so cook 2 grabs the lock first
- cook 2 cooks burger
   cook 2 releases lock cook 2 cooks burger for Y
- 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. 3. customer Y orders burger, order ends up with cook 2
- cook 1 was busy with something else, so cook 2 grabs the lock first
- cook 2 cooks burger for Y cook 2 cooks burger
   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**





#### **Correct Behaviors**



















#### Concurrent queue test program

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

# Correct behaviors (1 get, 1 put)



\$ harmony -c N\_GET=1 -c N\_PUT=1 code/queue\_btest2.hny

# Testing: comparing behaviors

- \$ harmony -o queue.hfa code/queue\_btest2.hny
- \$ harmony -B queue.hfa -m queue=queue\_lock code/queue\_btest2.hny
- The first command outputs the behavior of running the test program against the specification in file queue.hfa
- The second command runs the test program against the implementation and checks if its behavior matches that stored in queue.hfa

# **Black Box Testing**

- Not allowed to look under the covers

   can't use *rw*->nreaders, etc.
- Only allowed to invoke the interface methods and observe behaviors
- Your job: try to find bad behaviors
  - o compare against a specification
  - o how would you test a clock? An ATM machine?
    - without looking inside
- In general testing cannot ensure correctness
  - $\circ~$  only a correctness proof can
  - o testing may or may not expose a bug
  - $\circ$  model checking helps expose bugs



# Conditional Waiting





# **Conditional Waiting**

- Thus far we've shown how threads can wait for one another to avoid multiple threads in the critical section
- Sometimes there are other reasons:

   Wait until queue is non-empty
   Wait until there are no readers (or writer) in a reader/writer section

0 ...

## Reader/writer lock

Idea: allow multiple read-only operations to execute concurrently

- Still no data races
- In many cases, reads are much more frequent than writes

#### →Either:

- multiple readers, or
- a single writer

thus not:

- a reader and a writer, nor
- multiple writers

### **Reader/Writer Lock Specification**

```
def RWlock() returns lock:
1
2
        lock = { .nreaders: 0, .nwriters: 0 }
 3
    def read_acquire(rw):
4
        atomically when rw->nwriters == 0:
 5
             rw->nreaders += 1
 6
 7
8
    def read_release(rw):
9
        atomically rw->nreaders -= 1
10
11
    def write_acquire(rw):
        atomically when (rw->nreaders == 0) and (rw->nwriters == 0):
12
             rw->nwriters = 1
13
14
15
    def write_release(rw):
        atomically rw->nwriters = 0
16
```

## **Reader/Writer Lock Specification**

```
def RWlock() returns lock:
 1
 2
         lock = { .nreaders: 0, .nwriters: 0 }
 3
 4
    def read_acquire(rw):
         atomically when rw->nwriters == 0:
 5
             rw->nreaders += 1
 6
 7
    def read_release(rw):
 8
 9
         atomically rw->nreaders -= 1
10
11
    def write_acquire(rw):
12
         atomically when (rw - > nreaders == 0) and (rw - > nwriters == 0):
             rw \rightarrow nwriters = 1
13
14
15
    def write_release(rw):
16
         atomically rw->nwriters = 0
```

Invariants:

- if *n* readers in the R/W critical section, then  $nreaders \ge n$
- if *n* writers in the R/W critical section, then *nwriters*  $\geq n$
- $(nreaders \ge 0 \land nwriters = 0) \lor (nreaders = 0 \land 0 \le nwriters \le 1)$

#### R/W Locks: test for mutual exclusion

```
import rwlock
 2
 3
    nreaders = nwriters = 0
    invariant ((nreaders >= 0) and (nwriters == 0)) or
 4
 5
                ((nreaders == 0) and (nwriters == 1))
 6
 7
    const NOPS = 4
 8
9
    rw = rwlock.RWlock()
10
11
    def thread():
12
        while choose({ False, True }):
            if choose({ "read", "write" }) == "read":
13
14
                rwlock.read_acquire(?rw)
                atomically nreaders += 1
15
                                                         no writer, one or more readers
16
                atomically nreaders -= 1
                rwlock.read_release(?rw)
17
                                        # write
18
            else:
19
                rwlock.write_acquire(?rw)
20
                atomically nwriters += 1
                                                               one writer, no readers
                atomically nwriters -= 1
21
22
                rwlock.write_release(?rw)
23
24
    for i in {1..NOPS}:
25
        spawn thread()
```

#### Cheating R/W lock implementation

```
1
    import synch
 2
 3
    def RWlock() returns lock:
 4
         lock = synch.Lock()
 5
 6
    def read_acquire(rw):
 7
         synch.acquire(rw)
 8
9
    def read_release(rw):
         synch.release(rw)
10
11
12
    def write_acquire(rw):
         synch.acquire(rw)
13
14
15
    def write_release(rw):
         synch.release(rw)
16
```

The *lock* protects the application's critical section

#### Cheating R/W lock implementation

```
1
    import synch
 2
 3
    def RWlock() returns lock:
 4
         lock = synch.Lock()
 5
 6
    def read_acquire(rw):
 7
         synch.acquire(rw)
 8
9
    def read_release(rw):
10
         synch.release(rw)
11
12
    def write_acquire(rw):
13
         synch.acquire(rw)
14
15
    def write_release(rw):
16
         synch.release(rw)
```

The *lock* protects the application's critical section

Allows only one reader to get the lock at a time

Does *not* have the same behavior as the specification

- it is missing behaviors
- no bad behaviors though

```
from synch import Lock, acquire, release
1
 2
 3
    def RWlock() returns lock:
        lock = { .lock: Lock(), .nreaders: 0, .nwriters: 0 }
 4
 5
    def read_acquire(rw):
 6
 7
        acquire(?rw->lock)
        while rw->nwriters > 0:
 8
 9
             release(?rw->lock)
             acquire(?rw->lock)
10
        rw->nreaders += 1
11
        release(?rw->lock)
12
13
    def read_release(rw):
14
15
        acquire(?rw->lock)
        rw->nreaders -= 1
16
        release(?rw->lock)
17
18
    def write_acquire(rw):
19
        acquire(?rw->lock)
20
21
        while rw->nreaders > 0 or rw->nwriters > 0:
22
             release(?rw->lock)
             acquire(?rw->lock)
23
24
        rw->nwriters = 1
        release(?rw->lock)
25
26
    def write_release(rw):
27
28
        acquire(?rw->lock)
        rw->nwriters = 0
29
        release(?rw->lock)
30
```

```
from synch import Lock, acquire, release
 1
 2
 3
    def RWlock() returns lock:
 4
        lock = { .lock: Lock(), .nreaders: 0, .nwriters: 0
 5
    def read_acquire(rw):
 6
 7
        acquire(?rw->lock)
 8
        while rw->nwriters > 0:
 9
             release(?rw->lock)
             acquire(?rw->lock)
10
        rw->nreaders += 1
11
        release(?rw->lock)
12
13
14
    def read_release(rw):
15
        acquire(?rw->lock)
        rw->nreaders -= 1
16
        release(?rw->lock)
17
18
    def write_acquire(rw):
19
        acquire(?rw->lock)
20
21
        while rw->nreaders > 0 or rw->nwriters > 0:
22
             release(?rw->lock)
23
             acquire(?rw->lock)
24
        rw->nwriters = 1
        release(?rw->lock)
25
26
    def write_release(rw):
27
        acquire(?rw->lock)
28
29
        rw->nwriters = 0
        release(?rw->lock)
30
```

The *lock* protects *nreaders* and *nwriters*, not the critical section of the application

```
from synch import Lock, acquire, release
 1
 2
 3
    def RWlock() returns lock:
 4
        lock = { .lock: Lock(), .nreaders: 0, .nwriters: 0 }
 5
    def read_acquire(rw):
 6
 7
        acquire(?rw->lock)
        while rw->nwriters > 0:
 8
             release(?rw->lock)
 9
            acquire(?rw->lock)
10
        rw->nreaders += 1
11
        release(?rw->lock)
12
13
    def read_release(rw):
14
15
        acquire(?rw->lock)
        rw->nreaders -= 1
16
                                                                              waiting conditions
        release(?rw->lock)
17
18
    def write_acquire(rw):
19
        acquire(?rw->lock)
20
        while rw->nreaders > 0 or rw->nwriters > 0:
21
22
            release(?rw->lock)
23
            acquire(?rw->lock)
24
        rw->nwriters = 1
        release(?rw->lock)
25
26
    def write_release(rw):
27
28
        acquire(?rw->lock)
29
        rw->nwriters = 0
        release(?rw->lock)
30
```

180

```
from synch import Lock, acquire, release
 1
 2
 3
    def RWlock() returns lock:
 4
         lock = { .lock: Lock(), .nreaders: 0, .nwriters: 0 }
 5
    def read_acquire(rw):
 6
 7
         acquire(?rw->lock)
         while rw->nwriters > 0:
 8
 9
             release(?rw->lock)
             acquire(?rw->lock)
10
         rw->nreaders += 1
11
         release(?rw->lock)
12
13
    def read_release(rw):
14
15
         acquire(?rw->lock)
         rw->nreaders -= 1
16
         release(?rw->lock)
17
18
    def write_acquire(rw):
19
         acquire(?rw->lock)
20
        while rw->nreaders > 0 or rw->nwriters > 0:
21
22
             release(?rw->lock)
23
             acquire(?rw->lock)
24
         rw->nwriters = 1
         release(?rw->lock)
25
26
    def write_release(rw):
27
28
         acauire(?rw->lock)
         rw->nwriters = 0
29
        release(?rw->lock)
30
```

Good: has the same behaviors as the implemention

Bad: process is continuously scheduled to try to get the lock even if it's not available

(*Harmony complains about this as well*)

#### Mesa Condition Variables

- A lock can have one or more *condition variables*
- A thread that holds the lock but wants to wait for some condition to hold can *temporarily* release the lock by *waiting* on some condition variable
- Associate a condition variable with each "waiting condition"
  - reader: no writer in the critical section
  - writer: no readers nor writers in the c.s.
## Mesa Condition Variables, cont'd

 When a thread that holds the lock notices that some waiting condition is satisfied it should *notify* the corresponding condition variable

#### R/W lock with Mesa condition variables





```
def read_acquire(rw):
9
        acquire(?rw->mutex)
10
        while rw->nwriters > 0:
11
12
             wait(?rw->r_cond, ?rw->mutex)
13
        rw->nreaders += 1
14
        release(?rw->mutex)
15
    def read_release(rw):
16
17
        acquire(?rw->mutex)
18
        rw->nreaders -= 1
        if rw->nreaders == 0:
19
             notify(?rw->w_cond)
20
        release(?rw->mutex)
21
```







compare with busy waiting

```
def read_acquire(rw):
    acquire(?rw->lock)
    while rw->nwriters > 0:
        release(?rw->lock)
        acquire(?rw->lock)
    rw->nreaders += 1
    release(?rw->lock)
```

```
def read_release(rw):
    acquire(?rw->lock)
    rw->nreaders -= 1
    release(?rw->lock)
```

def read\_acquire(rw):
 acquire(?rw->mutex)
 while rw->nwriters > 0:
 wait(?rw->r\_cond, ?rw->mutex)
 rw->nreaders += 1
 release(?rw->mutex)

```
def read_release(rw):
    acquire(?rw->mutex)
    rw->nreaders -= 1
    if rw->nreaders == 0:
        notify(?rw->w_cond)
    release(?rw->mutex)
```

#### compare with busy waiting



def read\_acquire(rw):
 acquire(?rw->mutex)
 while rw->nwriters > 0:
 Wait(?rw->r\_cond, ?rw->mutex)
 rw->nreaders += 1
 release(?rw->mutex)

def read\_release(rw):
 acquire(?rw->mutex)
 rw->nreaders -= 1
 if rw->nreaders == 0:

notify(?rw->w\_cond)
release(?rw->mutex)

### R/W Lock, writer part



# **Condition Variable interface**

- wait(cv, lock)
  - o may only be called while holding *lock*
  - o temporarily releases *lock* 
    - but re-acquires it before resuming
  - o if cv not notified, may block indefinitely
    - but wait() may resume "on its own"
- notify(cv)
  - no-op if nobody is waiting on *cv*
  - o otherwise wakes up at least one thread waiting on *cv*
- notify\_all(cv)
  - wakes up all threads currently waiting on *cv*

```
def test_and_set(s) returns oldvalue:
    atomically:
        oldvalue = !s
        !s = True
```

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

```
def Lock() returns initvalue:
    initvalue = False
```

```
def acquire(lk):
    while test_and_set(lk):
        pass
```

```
def release(lk):
    atomic_store(lk, False)
```

def read\_acquire(rw):
 acquire(?rw->lock)
 while rw->nwriters > 0:
 release(?rw->lock)
 acquire(?rw->lock)
 rw->nreaders += 1
 release(?rw->lock)

```
def read_release(rw):
    acquire(?rw->lock)
    rw->nreaders -= 1
    release(?rw->lock)
```

```
def test_and_set(s) returns oldvalue:
    atomically:
        oldvalue = !s
        !s = True
def atomic_store(p, v):
    atomically !p = v
def Lock() returns initvalue:
    initvalue = False
def acquire(lk):
    while test_and_set(lk)
        pass
def release(lk):
```

```
atomic_store(lk, False)
```

def read\_acquire(rw):
 acquire(?rw->lock)
 while rw->nwriters > 0:
 release(?rw->lock)
 acquire(?rw->lock)
 rw->nreaders += 1
 release(?rw->lock)

def read\_release(rw):
 acquire(?rw->lock)
 rw->nreaders -= 1

release(?rw->lock)

```
def test_and_set(s) returns oldvalue:
    atomically:
        oldvalue = !s
        !s = True
def atomic_store(p, v):
    atomically !p = v
def Lock() returns initvalue:
    initvalue = False
def acquire(lk):
    while test_and_set(lk)
        pass
def release(lk):
```

```
atomic_store(lk, False)
```

State unchanged while condition does not hold. This thread only "observes" the state until condition holds

${\rm def}$	<pre>read_acquire(rw):</pre>
	<pre>acquire(?rw-&gt;lock)</pre>
(	<pre>while rw-&gt;nwriters &gt; 0:</pre>
	release(?rw->lock)
	acquire(?rw <mark>-&gt;</mark> lock)
	rw->nreaders += 1
	release(?rw->lock)
def	<pre>read_release(rw):</pre>
	acquire(?rw->lock)
	rw->nreaders -= 1
	release(?rw->lock)

State conditionally changes while condition does not hold. This thread actively changes the state until the condition hold

```
def test_and_set(s) returns oldvalue:
    atomically:
        oldvalue = !s
        !s = True
def atomic_store(p, v):
    atomically !p = v
def Lock() returns initvalue:
    initvalue = False
def acquire(lk):
    while test_and_set(lk)
        pass
def release(lk):
```

atomic\_store(lk, False)

State unchanged while condition does not hold. This thread only "observes" the state until condition holds



State conditionally changes while condition does not hold. This thread actively changes the state until the condition hold

# Why is busy waiting bad?

- Consider a timesharing setting
- Threads T1 and T2 take turns on the CPU
   o switch every 100 milliseconds
- Suppose T1 has a write lock and is running
- Now suppose a clock interrupt occurs, T2 starts running and tries to acquire a read lock
- Non-busy-waiting acquisition:
  - T2 is put on a waiting queue and T1 resumes immediately and runs until T1 releases the write lock
    - which puts T2 back on the run queue
- Busy-waiting acquisition:
  - T2 keeps running (wasting CPU) until the next clock interrupt
  - T1 and T2 switch back and forth every 100 ms until T1 releases the write lock

#### **Busy Waiting vs Condition Variables**

Busy Waiting	<b>Condition Variables</b>
Use a lock and a loop	Use a lock and a collection of condition variables and a loop
Easy to write the code	Notifying is tricky
Easy to understand the code	Easy to understand the code
Progress property is easy	Progress requires careful consideration (both for correctness and efficiency)
Ok-ish for true multi-core, but bad for virtual threads	Good for both multi-core and virtual threading

#### Just Say No to Busy Waiting

# Why no naked waits? (reason 1)

A naked wait is a wait() without while around it

- By the time waiter gets the lock back, condition may no longer hold
  - E.g., given three threads: W1, R2, W3
  - W1 enters as a writer
  - R2 waits as a reader
  - W1 leaves, notifying R2
  - W3 enters as a writer
  - R2 wakes up
    - If R2 doesn't check condition again, R2 and W3 would both be in the critical section

# Why no naked waits? (reason 2)

- When notifying, be safe rather than sorry

   it's better to notify too many threads than
   too few
  - in case of doubt, use notify\_all() instead of just notify()
- Over-notifying can lead to some threads waking up when their condition is no longer satisfied

# Why no naked waits? (reason 3)

- Because you should use while around wait, many condition variable implementations allow "spurious wakeups"
  - wait() resumes even though condition variable was not notified
  - simplifies implementation of wait()

#### Just Say No to Naked Waits

### Hints for reducing unneeded wakeups

- Use separate condition variables for each waiting condition
- Don't use notify\_all when notify suffices
   but be safe rather than sorry
- You can use N calls to **notify** if you know at most N nodes can continue after a waiting condition holds



#### Deadlock





### Deadlock example

```
1
     from synch import Lock, acquire, release
 2
 3
    def Account(balance) returns account:
         account = \{ .lock: Lock(), .balance: balance \}
 4
 5
    accounts = \begin{bmatrix} Account(3), Account(7), Account(0) \end{bmatrix}
 6
 7
 8
    def transfer(a1, a2, amount):
 9
         acquire(?accounts[a1].lock)
         if amount \leq accounts [a1].balance:
10
11
             accounts[a1].balance -= amount
             acquire(?accounts[a2].lock)
12
13
             accounts[a2].balance += amount
14
             release(?accounts[a2].lock)
15
         release(?accounts[a1].lock)
16
17
     spawn transfer(0, 1, 1)
     spawn transfer(1, 0, 2)
18
```

#### What could go wrong?

### Harmony output

#### **Summary: some execution cannot terminate**

- Schedule thread T0: init()
   Line 6: Set accounts to [ { "balance": 3, "lock": False }, { "balance": 7, "lock": False } ]
- Schedule thread T1: transfer(0, 1, 1)
   Line synch/36: Set accounts[0]["lock"] to True (was False)
   Line 11: Set accounts[0]["balance"] to 2 (was 3)
   Preempted in transfer(0, 1, 1) --> acquire(?accounts[1]["lock"])
- Schedule thread T2: transfer(1, 0, 2)
   Line synch/36: Set accounts[1]["lock"] to True (was False)
   Line 11: Set accounts[1]["balance"] to 5 (was 7)
   Preempted in transfer(1, 0, 2) --> acquire(?accounts[0]["lock"])

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

#### Threads:

T1: (blocked) transfer(0, 1, 1) --> acquire(?accounts[1]["lock"])

T2: (blocked) transfer(1, 0, 2) --> acquire(?accounts[0]["lock"])

#### Variables:

accounts: [ { "balance": 2, "lock": True }, { "balance": 5, "lock": True } ]

## Harmony HTML Output

Issue: Non-terminating state				Shared Variables		
Turn	Thread	Instructions Executed	PC	accounts		
1	T0:init()	terminated	1309	[ { "balance": 3, "lock": False }, { "balance": 7, "lock": False } ]		
2	T1: transfer(0, 1, 1)	about to execute in <u>synch</u> :35: atomically when not !binsema:	949	[ { "balance": 2, "lock": True }, { "balance": 7, "lock": False } ]		
3	T2: transfer(1, 0, 2)	about to execute in <u>synch</u> :35: atomically when not !binsema:	949	[ { "balance": 2, "lock": True }, { "balance": 5, "lock": True } ]		

#### synch:34 def acquire(binsema):

934	934 LoadVar binsema		Threads						
935	5 DelVar binsema		Status	Stack Trace	Stack Top				
936	Load	T0	T0 terminatedinit()						
937	StoreVar result	т1	blocked	transfer(0, 1, 1) a1: 0, a2: 1, amount: 1					
938	ReturnOp(result)		DIOCKCU	acquire(?accounts[1]["lock"])binsema: ?accounts[1]["lock"]					
939	Jump 1214	тэ	blocked	transfer(1, 0, 2) a1: 1, a2: 0, amount: 2					
940	Frame held(binsema)			<pre>acquire(?accounts[0]["lock"])binsema: ?accounts[0]["lock"]</pre>					

#### **Deadlock vs Starvation**

- Starvation: some processes can run in theory, but the scheduler continually selects other processes to run first. Tied to fairness in scheduling.
- Deadlock: no process can run because all are waiting for another process to change the state. The scheduler can't help you now.

### Deadlock vs Livelock

- Livelock: some processes continually change their state but don't make progress (like polite people trying to pass one another in a narrow hallway). The scheduler could fix this in theory.
- Deadlock: no process can run because all are waiting for another process to change the state. The scheduler can't help you now.

# System Model

- Collection of resources and threads
  - Examples of resources: I/O devices, GPUs, locks, buffers, slots in a buffer, ...
- Exclusive access
  - Only one thread can use a resource at a time
    Protocol:
    - 1. Thread acquires resource
      - thread is blocked until resource is free
    - 2. Thread holds the resource
      - resource is allocated (not free) at this time
    - 3. Thread releases the resource

#### Necessary Conditions for Deadlock Edward Coffman 1971

#### **1. Mutual Exclusion**

• acquire() can block invoker until resource is free

#### 2. Hold & wait

• A thread can be blocked while holding resources

#### 3. No preemption

• Allocated resources cannot be reclaimed

#### 4. Circular wait

- Let  $T_i \rightarrow T_j$  denote " $T_i$  waits for  $T_j$  to release a resource".
- Then  $\exists T_1, \dots, T_n : T_1 \rightarrow T_2 \rightarrow \dots \rightarrow T_n \rightarrow T_1$

### **Example: Mutual Exclusion**

```
1
    from synch import Lock, acquire, release
 2
 3
    def Account(balance) returns account:
         account = \{ .lock: Lock(), .balance: balance \}
 4
 5
    accounts = \begin{bmatrix} Account(3), Account(7), Account(0) \end{bmatrix}
 6
 7
 8
    def transfer(a1, a2, amount):
                                                   Mutual exclusion
         acquire(?accounts[a1].lock)
 9
         if amount <= accounts [a1].balance:
10
11
             accounts[a1].balance -= amount
                                                      Mutual exclusion
12
             acquire(?accounts[a2].lock)
13
             accounts[a2].balance += amount
14
             release(?accounts[a2].lock)
15
         release(?accounts[a1].lock)
16
17
    spawn transfer(0, 1, 1)
    spawn transfer(1, 0, 2)
18
```

### Example: Hold & Wait

```
1
    from synch import Lock, acquire, release
 2
 3
    def Account(balance) returns account:
         account = \{ .lock: Lock(), .balance: balance \}
 4
 5
    accounts = \begin{bmatrix} Account(3), Account(7), Account(0) \end{bmatrix}
 6
 7
 8
    def transfer(a1, a2, amount):
 9
         acquire(?accounts[a1].lock)
                                                     Thread holds a1.lock
         if amount <= accounts[a1].balance:
10
11
             accounts[a1].balance -= amount
                                                   Thread wants a2.lock
12
             acquire(?accounts[a2].lock)
13
             accounts[a2].balance += amount
14
             release(?accounts[a2].lock)
15
         release(?accounts[a1].lock)
16
17
    spawn transfer(0, 1, 1)
    spawn transfer(1, 0, 2)
18
```

#### **Example: No Preemption**

```
1
     from synch import Lock, acquire, release
 2
 3
    def Account(balance) returns account:
         account = \{ .lock: Lock(), .balance: balance \}
 4
 5
    accounts = \begin{bmatrix} Account(3), Account(7), Account(0) \end{bmatrix}
 6
 7
 8
    def transfer(a1, a2, amount):
 9
         acquire(?accounts[a1].lock)
         if amount \leq accounts all balance:
10
11
             accounts[a1].balance -= amount
12
             acquire(?accounts[a2].lock)
13
             accounts[a2].balance += amount
14
             release(?accounts[a2].lock)
                                                Only holder can release lock
15
         release(?accounts[a1].lock)
16
17
     spawn transfer(0, 1, 1)
     spawn transfer(1, 0, 2)
18
```

## Example: Circular Wait

```
1
     from synch import Lock, acquire, release
 2
 3
    def Account(balance) returns account:
         account = \{ .lock: Lock(), .balance: balance \}
 4
 5
    accounts = \begin{bmatrix} Account(3), Account(7), Account(0) \end{bmatrix}
 6
 7
 8
    def transfer(a1, a2, amount):
 9
         acquire(?accounts[a1].lock)
         if amount <= accounts[a1].balance:
10
11
             accounts[a1].balance -= amount
12
             acquire(?accounts[a2].lock)
13
             accounts[a2].balance += amount
14
             release(?accounts[a2].lock)
15
         release(?accounts[a1].lock)
16
17
     spawn transfer(0, 1, 1)
                                      Circular wait conditions
     spawn transfer(1, 0, 2)
18
```

Three ways to deal with deadlock <u>Prevention</u>: Programmer ensures that at least one of the necessary conditions cannot hold

<u>Avoidance</u>: Scheduler avoids deadlock scenarios (e.g., by executing each thread to completion)

Detect and Recover: Allow deadlocks to happen. Detect them and recover in some way



# Deadlock Prevention




## Negate one of the following:

- **1. Mutual Exclusion**
- 2. Hold & wait
- 3. No preemption
- 4. Circular wait

#### 1. Negate Mutual Exclusion

- Make resources sharable without locks
  - Non-blocking concurrent data structures
    - See Harmony book for examples
- Have sufficient resources available so acquire() never blocks
  - bounded buffer: make sure it is large enough
    - Doesn't work for locks, as there is only one per critical section

#### 2. Negate Hold & Wait

```
1
    from synch import Lock, acquire, release
 2
 3
    def Account(balance) returns account:
        account = \{ .lock: Lock(), .balance: balance \}
4
 5
6
    accounts = \lceil Account(3), Account(7) \rceil
 7
8
    def transfer(a1, a2, amount):
9
        acquire(?accounts[a1].lock)
        if amount <= accounts[a1].balance:
10
11
             accounts[a1].balance -= amount
                                                   Release resource
12
             release(?accounts[a1].lock)
                                                before acquiring another
13
             acquire(?accounts[a2].lock)
14
             accounts[a2].balance += amount
15
             release(?accounts[a2].lock)
        else:
16
             release(?accounts[a1].lock)
17
18
19
    spawn transfer(0, 1, 1)
    spawn transfer(1, 0, 2)
20
```

## 2: Negate Hold & Wait, badly

```
from synch import Lock, acquire, release
1
2
    def Account(balance) returns account:
3
4
        account = { .lock: Lock(), .balance: balance }
5
6
    accounts = [Account(3), Account(7)]
7
    invariant all(a.balance \geq 0 for a in accounts)
8
9
10
    def transfer(a1, a2, amount):
        acquire(?accounts[a1].lock)
11
                                                                 check if funds are available
        var funds_available = amount <= accounts[a1].balance</pre>
12
        release(?accounts[a1].lock)
13
        if funds_available:
14
            acquire(?accounts[a1].lock)
15
                                                                 withdraw funds from a1
            accounts[a1].balance -= amount
16
            release(?accounts[a1].lock)
17
            acquire(?accounts[a2].lock)
18
                                                                 deposit funds for a2
            accounts[a2].balance += amount
19
            release(?accounts[a2].lock)
20
21
    spawn transfer(0, 1, 2)
22
                                                    What could go wrong?
    spawn transfer(0, 1, 2)
23
```

#### 2. Negate Hold & Wait, alternate

```
def Lock() returns lock:
 1
 2
        lock = False
 3
    def acquire2(lk1, lk2):
 4
                                                     Spec: Acquire two locks
 5
        atomically when not (!lk1 or !lk2):
            !!k1 = !!k2 = True
 6
 7
 8
    def release(lk):
 9
        atomically !lk = False
10
11
    def Account(balance) returns account:
12
        account = { .lock: Lock(), .balance: balance }
13
14
    accounts = [Account(3), Account(7)]
15
16
    def transfer(a1, a2, amount):
                                                                 Acquire resources at
        acquire2(?accounts[a1].lock, ?accounts[a2].lock)
17
                                                                     the same time
18
        if amount <= accounts[a1].balance:</pre>
19
            accounts all.balance -= amount
20
            accounts[a2].balance += amount
21
        release(?accounts[a1].lock)
22
        release(?accounts[a2].lock)
23
24
    spawn transfer(0, 1, 1)
    spawn transfer(1, 0, 2)
25
```

#### 3. Allow Preemption

- Time-multiplexing of resources
  - threads: context switching
  - memory: paging
- Database transactions
  - 2-phase locking + transaction abort and retry
- Not available for locks

#### 4: Negate circular wait

- Define a total order on resources
- Rule: a thread cannot acquire a resource that is "lower" than a resource already held
- Either:
  - a thread is careful to acquire resources that it needs in order, or
  - a thread that wants to acquire a resource R must first release all resources that are lower than R

## Why does resource ordering work?

# **Theorem**: Resource ordering prevents circular wait **Proof by contradiction**:

- Assume circular wait exists
- $\exists T_1, \dots, T_n : T_1 \rightarrow T_2 \rightarrow \dots \rightarrow T_n \rightarrow T_1$
- T<sub>i</sub> holds R<sub>i</sub>
- $T_i$  requests  $R_j$  held by  $T_j$   $(j = (i + 1) \mod n)$
- Resource ordering:  $R_1 < R_2$ , ...,  $R_{n-1} < R_n$ ,  $R_n < R_1$
- R<sub>1</sub> < R<sub>1</sub> (by *transitivity* of total order)
- Violates *irreflexivity* of total order

#### 4: Negate circular wait

```
from synch import Lock, acquire, release
1
2
3
    def Account(balance) returns account:
        account = \{ lock: Lock(), balance: balance \}
4
 5
6
    accounts = \lceil Account(3), Account(7) \rceil
 7
8
    def transfer(a1, a2, amount):
                                                      Acquire resources
9
        acquire(?accounts[min(a1, a2)].lock)
                                                           in order
10
        acquire(?accounts[max(a1, a2)].lock)
        if amount <= accounts[a1].balance:
11
12
             accounts al .balance -= amount
13
             accounts[a2].balance += amount
14
        release(?accounts[a1].lock)
15
         release(?accounts[a2].lock)
16
17
    spawn transfer(0, 1, 1)
18
    spawn transfer(1, 0, 2)
```







#### Deadlock in traffic







#### How can these be avoided?

- Scheduler carefully schedules threads so deadlock cannot occur
- For example, it might allow only one thread to run at a time, to completion

  This is extreme: no concurrency
  Doesn't work with conditional waiting
- Better solutions typically require that the scheduler has some abstract knowledge of what the threads are trying to accomplish

#### Safe States

- A *state* is an allocation of resources to threads
- The state changes each time a thread allocates or releases a resource
- A *safe state* is a state from which an execution exists that does not cause deadlock
- Notes:
  - the initial state is safe: threads can be scheduled one at a time and run to completion
  - an unsafe state is not necessarily deadlocked, but deadlock is unavoidable eventually
  - deadlock may be possible from a safe state, but it is avoidable through careful scheduling

Scheduler should only allow safe states to happen in an execution

When a thread tries to acquire() a resource, the scheduler should block the thread, if acquiring the resource leads to an unsafe state, until this is no longer the case
release() is always ok

```
1
     from synch import Lock, acquire, release
 2
 3
    def Account(balance) returns account:
         account = \{ .lock: Lock(), .balance: balance \}
 4
 5
    accounts = \begin{bmatrix} Account(3), Account(7), Account(0) \end{bmatrix}
 6
 7
 8
    def transfer(a1, a2, amount):
 9
         acquire(?accounts[a1].lock)
         if amount <= accounts[a1].balance:
10
11
             accounts[a1].balance -= amount
12
             acquire(?accounts[a2].lock)
13
             accounts[a2].balance += amount
14
             release(?accounts[a2].lock)
15
         release(?accounts[a1].lock)
16
17
     spawn transfer(0, 1, 1)
    spawn transfer(1, 0, 2)
18
```

How?

```
1
     from synch import Lock, acquire, release
 2
    def Account(balance) returns account:
 3
         account = \{ .lock: Lock(), .balance: balance \}
 4
 5
    accounts = \begin{bmatrix} Account(3), Account(7), Account(0) \end{bmatrix}
 6
 7
 8
    def transfer(a1, a2, amount):
         acquire(?accounts[a1].lock)
 9
         if amount \leq accounts all balance:
10
11
             accounts[a1].balance -= amount
             acquire(?accounts[a2].lock)
12
13
             accounts[a2].balance += amount
             release(?accounts[a2].lock)
14
15
         release(?accounts[a1].lock)
16
17
     spawn transfer(0, 1, 1)
     spawn transfer(1, 0, 2)
18
```

For example, don't schedule two threads transfer(a1, a2) and transfer(a3, a4) at the same time unless  $\{a1, a2\} \cap \{a3, a4\} = \emptyset$ 

## Avoidance specified in Harmony

```
from synch import Lock, acquire, release
def Account(balance) returns account:
    account = { .lock: Lock(), .balance: balance }
active = {}
accounts = \lceil Account(3), Account(7) \rceil
def transfer(a1, a2, amount):
    atomically when ({ a1, a2 } & active) == {}:
        active |= \{a1, a2\}
    acquire(?accounts[a1].lock)
    if amount <= accounts[a1].balance:
        accounts[a1].balance -= amount
        acquire(?accounts[a2].lock)
        accounts[a2].balance += amount
        release(?accounts[a2].lock)
    release(?accounts[a1].lock)
    atomically:
        active -= { a1, a2 }
```

1 2 3

4

5

6

7 8 9

10

11 12 13

14

15

16

17

18

19

20

21

22

keep track of which accounts are active

enforce no intersection with active transfers



# Deadlock Detection and Recovery





 Keep track of allocation of resources to threads



- Keep track of allocation of resources to threads
- Keep track of which threads are trying to acquire which resource



- Known as the Resource Allocation Graph
- Deadlock  $\equiv$  cycle in the graph



- Known as the Resource Allocation Graph
- Deadlock  $\equiv$  cycle in the graph



## Finding Cycles

- Graph Reduction Algorithm:
  - While there are nodes with no outgoing edges
    - select one such node
    - remove node and its incoming edges
  - If the resulting graph empty (no nodes), then no cycles
  - $\circ$  No cycles  $\Longrightarrow$  No deadlock











No more nodes can be removed, but graph is non-empty  $\rightarrow$  *cycle is present* 



- Deadlock detection is expensive
- When to run graph reduction?
  - $\circ$  When a resource request cannot be granted?
  - When a thread has been blocked for a certain amount of time?
  - Periodically?

#### **Deadlock Recovery Strategies**

- Blue screen and reboot
   Can lose data / results of long computations
- Deny a request to remove cycle
   Programmer responsible for exception
- Kill processes until cycle is gone
  - Can lose data / results of long computations
  - Select processes that have been running shortest amount of time
- Use transactions to access resources
  - Abort and retry transaction if deadlock exists
  - Requires roll-back or versioning of state



#### Actors



[Robbert van Renesse]

#### Actor Model

- An actor is a type of process
- Each actor has an incoming message queue
- No other shared state
- Actors communicate by "message passing"
   placing messages on message queues
- Supports modular concurrent programs
- Actors and message queues are abstractions

#### Mutual Exclusion with Actors

- Data structure owned by a "server actor" •
- Client actors can send request messages to the server and receive response messages if necessary
- Server actor awaits requests on its queue and executes one request at a time
- →
  - Mutual Exclusion (one request at a time) Ο
  - Progress (requests eventually get to the head of the queue) Fairness (requests are handled in FCFS order) 0
  - Ο



#### Conditional Critical Sections with Actors

- An actor can "wait" for a condition by waiting for a specific message
- An actor can "notify" another actor by sending it a message

#### Parallel processing with Actors

- Organize program with a Manager Actor and a collection of Worker Actors
- Manager Actor sends work requests to the Worker Actors
- Worker Actors send completion requests to the Manager Actor


## Parallel processing example

```
from synch import *
1
2
3
    ranges = { (2,10), (11,20), (21,30) }
   queues = { r:Queue() for r in ranges }
4
5
    maing = Queue()
6
7
    def isPrime(v) returns prime:
8
        prime = True
9
        var d = 2
        while prime and (d < v):
10
            if (v % d) == 0:
11
                prime = False
12
13
            d += 1
14
    def worker(q):
15
16
        while True:
17
            let rq, (start, finish) = get(q):
                for p in { start .. finish }:
18
                     if isPrime(p):
19
20
                         put(rq, p)
21
22
    def main(rq, workers):
        for r:q in workers:
23
            put(q, (rq, r))
24
        while True:
25
26
            print get(rq)
27
28
    for r in ranges:
        spawn eternal worker(?queues[r])
29
30
    spawn eternal main(?mainq, { r:?queues[r] for r in ranges })
```

# Pipeline Parallelism with Actors

- Organize program as a chain of actors
- For example, REST/HTTP server

   Network receive actor → HTTP parser actor
   → REST request actor → Application actor
   → REST response actor → HTTP response actor → Network send actor

automatic flow control (when actors run at different rates)

• with bounded buffer queues

# **Pipelining Example**

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

```
from synch import *
1
2
3
    const MAX = 10
4
5
    def isPrime(v) returns prime:
6
        prime = True
        var d = 2
7
8
        while prime and (d < v):
            if (v % d) == 0:
9
                 prime = False
10
11
            d += 1
12
13
    q1 = q2 = q3 = Queue
14
15
    def actor0():
16
        for v in \{2, MAX\}:
17
             put(?q1, v)
```

```
def actor1():
    while True:
        let v = get(?q1):
            put(?q2, (2 ** v) - 1)
def actor2():
    while True:
        let v = qet(?q2):
            if isPrime(v):
                put(?q3, v)
def actor3():
    while True:
        let v = get(?q3):
            print(v)
spawn actor0()
spawn eternal actor1()
spawn eternal actor2()
spawn eternal actor3()
```

#### Find Mersenne primes

Support for actors in programming languages

- Native support in languages such as Scala and Erlang
- "blocking queues" in Python, Harmony, Java
- Actor support libraries for Java, C, ...

Actors also nicely generalize to distributed systems!

# Actor disadvantages?

- Doesn't work well for "fine-grained" synchronization
  - overhead of message passing much higher than lock/unlock
- Sending/receiving messages just to access a data structure leads to significant extra code



# Barrier Synchronization





# Barrier Synchronization: the opposite of mutual exclusion...

- Set of processes run in rounds
- Must all complete a round before starting the next
- Popular in simulation, HPC, graph processing, model checking...
  - Lock-based synchronization reduces opportunities for parallelism
  - Barrier Synchronization supports scalable parallelism

## **Barrier abstraction**

- Barrier(N): barrier for N threads
- bwait(): start the next round



#### Example: dot product

```
import barrier
1
2
 3
    const NWORKERS = 2
4
    vec1 = [ 1, 2, 3, 4 ]
 5
    vec2 = [ 5, 6, 7, 8 ]
 6
    barr = barrier.Barrier(NWORKERS)
 7
    output = [0,] * NWORKERS
8
9
    def split(self, v) returns x:
10
        x = (self * len(v)) / NWORKERS
11
12
    def dotproduct(self, v1, v2):
13
        assert len(v1) == len(v2)
14
15
        var total = 0
16
        for i in { split(self, v1) .. split(self + 1, v1) - 1}:
            total += v1[i] * v2[i]
17
        output[self] = total
18
19
        barrier.bwait(?barr)
        print sum(output)
20
21
22
    for i in \{0 \dots \text{NWORKERS} - 1\}:
23
        spawn dotproduct(i, vec1, vec2)
```

# Test program for barriers

```
import barrier
1
2
3
    const NTHREADS = 3
    const NROUNDS = 4
4
5
    barr = barrier.Barrier(NTHREADS)
6
    before = after = [0,] * NTHREADS
7
8
    invariant min(before) >= max(after)
9
10
11
    def thread(self):
        for _ in { 1 .. NROUNDS }:
12
                                       work done before barrier
            before self += 1
13
            barrier.bwait(?barr)
14
                                        work done after barrier
            after self += 1
15
16
    for i in \{0 ... NTHREADS - 1 \}:
17
        spawn thread(i)
18
```

# Test program for barriers

```
import barrier
1
2
    const NTHREADS = 3
3
    const NROUNDS = 4
4
5
    barr = barrier.Barrier(NTHREADS)
6
    before = after = [0,] * NTHREADS
7
                                                 no one can pass
8
                                                  barrier until all
    invariant min(before) >= max(after)
9
10
                                               reached the barrier
11
    def thread(self):
        for _ in { 1 .. NROUNDS }:
12
                                      work done before barrier
13
            before self += 1
            barrier.bwait(?barr)
14
                                       work done after barrier
            after self += 1
15
16
    for i in \{0 ... NTHREADS - 1 \}:
17
        spawn thread(i)
18
```

```
def Barrier(required) returns barrier:
1
2
        barrier = { .required: required, .n: 0 }
                                                        State:
3
                                                           required: #threads
                                                           n: #threads that have
   def bwait(b):
4
                                                        _
5
        atomically b \rightarrow n += 1
                                                           reached the barrier
        atomically await b->n == b->required
6
```

266







```
def Barrier(required) returns barrier:
1
2
        barrier = { .required: required, .n: 0 }
3
   def bwait(b):
4
                                                       _
5
        atomically b \rightarrow n += 1
        atomically await b->n == b->required
6
```

```
State:
```

- *required*: #threads
- *n*: #threads that have reached the barrier

```
Only works one round
```

```
def Barrier(required) returns barrier:
1
2
        barrier = { .required: required, .n: 0 }
3
   def bwait(b):
4
        atomically:
5
            b->n += 1
6
7
            if b->n == b->required:
                b -> n = 0
8
        atomically await b->n == 0
9
```

```
1
    def Barrier(required) returns barrier:
2
        barrier = { .required: required, .n: 0 }
3
    def bwait(b):
4
        atomically:
5
            b->n += 1
6
7
             if b->n == b->required:
8
                 b -> n = 0
9
        atomically await b \rightarrow n = 0
```

Broken! (if used more than once)

271

```
def Barrier(required) returns barrier:
1
2
        barrier = { required: required, n: [0, 0] }
 3
    def turnstile(b, i):
4
5
        atomically:
6
            b->n[i] += 1
7
            if b->n[i] == b->required:
8
                 b \to n[1 - i] = 0
        atomically await b->n[i] == b->required
9
10
11
    def bwait(b):
12
        turnstile(b, 0)
        turnstile(b, 1)
13
```



```
def Barrier(required) returns barrier:
 1
 2
         barrier = { required: required, n: [0, 0] }
 3
4
    def turnstile(b, i):
 5
         atomically:
 6
             b->n[i] += 1
 7
             if b->n[i] == b->required:
 8
                  b \to n[1 - i] = 0
9
         atomically await b \rightarrow n[i] == b \rightarrow required
10
11
    def bwait(b):
12
         turnstile(b, 0)
                                        Works, but double
waiting is inefficient
         turnstile(b, 1)
13
```



273

## Barrier Specification, final version

```
def Barrier(required) returns barrier:
1
2
         barrier = { .required: required, .n: 0, .color: 0 }
 3
    def bwait(b):
4
                                                   State:
 5
         var color = None
                                                      required: #threads
         atomically:
 6
                                                      n: #threads that have
             color = b -> color
 7
             b -> n += 1
                                                      reached the barrier
 8
             if b->n == b->required:
 9
                                                    color: allows re-use of
                  b \rightarrow color = 1
10
                                                      barrier. Flipped each round
11
                  b -> n = 0
12
         atomically await b->color != color
```



#### **Barrier Implementation**

```
from synch import *
1
 2
3
    def Barrier(required) returns barrier:
         barrier = {
4
             .mutex: Lock(), .cond: Condition(),
 5
 6
             required: required, n: 0, color: 0
 7
         }
 8
9
    def bwait(b):
10
         acquire(?b->mutex)
11
         b -> n += 1
         if b->n == b->required:
12
             b \rightarrow color = 1
13
14
             b -> n = 0
15
             notify_all(?b->cond)
16
         else:
             let color = b->color:
17
18
                 while b->color == color:
                      wait(?b->cond, ?b->mutex)
19
         release(?b->mutex)
20
```

# **Advanced Barrier Synchronization**

- Given is a resource of finite capacity
   Bus with N seats, say
- Resource must be used at full capacity
   O Bus won't go until it is full
- Resource must be completed emptied before it can be re-used

Everybody must get off at destination
 before anybody can get back on the bus

# **Advanced Barrier Synchronization**

- Given is a resource of finite capacity
   Bus with N seats, say
- Resource must be used at full
   Bus won't go until it is full
- Resource must before it ased

• **Explanation** Leanybody can get back on the bus

# Interface

#### • enter(*resource*)

 must wait if resource is in use or if resource has not yet been fully unloaded
 after that, must wait until resource is full

• exit(*resource*)

 $\circ$  any time

## **Rounds and Phases**

- Round: each time the resource gets used
- Three phases in each round:
  - 1. Resource is loaded
  - 2. Resource is used
  - 3. Resource is unloaded
- Two waiting conditions:
   Wait until resource is fully unloaded

– Before starting to load the resource

- Wait until resource is fully loaded
  - Before starting to use the resource

# Rollercoaster

```
from synch import *
 1
 2
 3
    def RollerCoaster(nseats): result = {
 4
         .mutex: Lock(), .nseats: nseats, .entered: 0, .left: nseats,
 5
         .empty: Condition(), .full: Condition()
 6
 7
    def enter(b):
 8
9
        acquire(?b->mutex)
        while b->entered == b->nseats: # wait for car to empty out
10
            wait(?b->empty, ?b->mutex)
11
12
        b->entered += 1
        if b->entered != b->nseats: # wait for car to fill up
13
14
            while b->entered < b->nseats:
15
                wait(?b->full, ?b->mutex)
16
        else:
                                         # car is ready to go
17
            b \rightarrow left = 0
            notify_all(?b->full)  # wake up others waiting in car
18
        release(?b->mutex)
19
20
    def exit(b):
21
22
        acquire(?b->mutex)
23
        b \rightarrow left += 1
24
        if b->left == b->nseats: # car is empty
25
             b->entered = 0
            notify_all(?b->empty)
                                     # wake up riders wanting to go
26
        release(?b->mutex)
27
```



JOE MCBRIDE / GETTY IMAGES



# Interrupt Safety





# Interrupt handling

- When executing in user space, a device interrupt is invisible to the user process
  - State of user process is unaffected by the device interrupt and its subsequent handling
  - This is because contexts are switched back and forth
  - So, the user space context is *exactly restored* to the state it was in before the interrupt

# Interrupt handling

- However, there are also "in-context" interrupts:
  - o kernel code can be interrupted
  - o user code can handle "signals"
- → Potential for race conditions

# "Traps" in Harmony



#### But what now?

```
count = 0
1
    done = False
2
 3
    finally count == 2
4
 5
    def handler():
 6
 7
         count += 1
        done = True
8
9
    def main():
10
        trap handler()
11
12
         count += 1
        await done
13
14
    spawn main()
15
```

#### But what now?

```
count = 0
 1
    done = False
2
 3
    finally count == 2
4
 5
    def handler():
 6
 7
         count += 1
         done = True
 8
 9
10
    def main():
         trap handler()
11
12
         count += 1
13
         await done
14
    spawn main()
15
```

#### Summary: something went wrong in an execution

- Schedule thread To: init()
  - Line 1: Initialize count to 0
  - Line 2: Initialize done to False
  - Thread terminated
- Schedule thread T1: main()
  - Line 12: Interrupted: jump to interrupt handler first
  - Line 12: Interrupts disabled
  - Line 7: Set count to 1 (was 0)
  - Line 8: Set done to True (was False)
  - Line 6: Interrupts enabled
  - Line 12: Set count to 1 (unchanged)
  - Thread terminated
- Schedule thread T2: finally()
  - Line 4: Harmony assertion failed

#### Locks to the rescue?

```
from synch import Lock, acquire, release
1
2
 3
    countlock = Lock()
    count = 0
4
    done = False
 5
 6
 7
    finally count == 2
 8
9
    def handler():
        acquire(?countlock)
10
        count += 1
11
12
        release(?countlock)
        done = True
13
14
15
    def main():
16
        trap handler()
17
        acquire(?countlock)
18
        count += 1
19
        release(?countlock)
20
        await done
21
22
    spawn main()
```

## Locks to the rescue?

```
Summary: some execution cannot terminate
     from synch import Lock, acq
 1

    Schedule thread T0: init()

 2
 3
     countlock = Lock()

    Line 3: Initialize countlock to False

     count = 0
 4

    Line 4: Initialize count to 0

     done = False
 5

    Line 5: Initialize done to False

 6
 7
     finally count == 2

    Schedule thread T1: main()

 8

    Line synch/36: Set countlock to True (was False)

 9
     def handler():

    Line 18: Set count to 1 (was 0)

10
           acquire(?countlock)
                                               • Line synch/39: Interrupted: jump to interrupt handler first
11
           count += 1
12
           release(?countlock)

    Line synch/39: Interrupts disabled

           done = True
13

    Preempted in main() --> release(?countlock) --> handler() --> acquire(?

14
                                                  countlock) about to execute atomic section in line synch/35
15
     def main():
16
           trap handler()
17
           acquire(?countlock)
                                           Final state (all threads have terminated or are blocked):
           count += 1
18
                                           • Threads:
19
           release(?countlock)

    T1: (blocked interrupts-disabled) main() --> release(?countlock) -->

20
           await done
                                                  handler() --> acquire(?countlock)
21

    about to execute atomic section in line synch/35

22
     spawn main()
```

# Enabling/disabling interrupts

```
count = 0
1
    done = False
2
 3
    finally count == 2
4
5
    def handler():
6
7
        count += 1
        done = True
8
9
    def main():
10
11
        trap handler()
                                disable interrupts
12
        setintlevel(True)
13
        count += 1
                                enable interrupts
        setintlevel(False)
14
15
        await done
16
    spawn main()
17
```

#### Interrupt-Safe Methods

```
count = 0
1
    done = False
2
3
    finally count == 2
4
5
    def increment():
6
        let prior = setintlevel(True):
 7
             count += 1
8
             setintlevel(prior)
9
10
11
    def handler():
12
        increment()
        done = True
13
14
    def main():
15
16
        trap handler()
        increment()
17
        await done
18
19
    spawn main()
20
```

disable interrupts restore old interrupt level
```
from synch import Lock, acquire, release
1
 2
 3
    count = 0
    countlock = Lock()
4
 5
    done = [ False, False ]
 6
 7
    finally count == 4
8
9
    def increment():
        let prior = setintlevel(True):
10
            acquire(?countlock)
11
12
            count += 1
13
            release(?countlock)
            setintlevel(prior)
14
15
    def handler(self):
16
17
        increment()
        done self = True
18
19
20
    def thread(self):
21
        trap handler(self)
22
        increment()
        await done[self]
23
24
25
    spawn thread(0)
    spawn thread(1)
26
```

```
from synch import Lock, acquire, release
1
 2
 3
    count = 0
4
    countlock = Lock()
 5
    done = [ False, False ]
 6
 7
    finally count == 4
 8
9
    def increment():
        let prior = setintlevel(True):
10
            acquire(?countlock)
11
12
            count += 1
13
            release(?countlock)
            setintlevel(prior)
14
15
    def handler(self):
16
17
        increment()
        done self = True
18
19
20
    def thread(self):
21
        trap handler(self)
        increment()
22
                                 wait for own interrupt
        await done[self]
23
24
25
    spawn thread(0)
26
    spawn thread(1)
```

```
from synch import Lock, acquire, release
1
 2
 3
    count = 0
4
    countlock = Lock()
 5
    done = [ False, False ]
 6
 7
    finally count == 4
                                            first disable interrupts
 8
9
    def increment():
        let prior = setintlevel(True):
10
            acquire(?countlock)
11
12
            count += 1
13
            release(?countlock)
            setintlevel(prior)
14
15
    def handler(self):
16
17
        increment()
        done self = True
18
19
20
    def thread(self):
21
        trap handler(self)
        increment()
22
                                wait for own interrupt
        await done[self]
23
24
25
    spawn thread(0)
    spawn thread(1)
26
```

```
from synch import Lock, acquire, release
1
 2
 3
    count = 0
4
    countlock = Lock()
 5
    done = [ False, False ]
 6
 7
    finally count == 4
                                            first disable interrupts
 8
    def increment():
9
        let prior = setintlevel(True):
10
            acquire(?countlock)
11
                                      then acquire a lock
12
            count += 1
13
            release(?countlock)
            setintlevel(prior)
14
15
    def handler(self):
16
17
        increment()
        done[self] = True
18
19
    def thread(self):
20
21
        trap handler(self)
        increment()
22
                                wait for own interrupt
        await done[self]
23
24
25
    spawn thread(0)
    spawn thread(1)
26
```



# Warning: very few C functions are interrupt-safe

- pure system calls are interrupt-safe
   e.g. read(), write(), etc.
- functions that do not use global data are interrupt-safe

o e.g. strlen(), strcpy(), etc.

- malloc() and free() are *not* interrupt-safe
- printf() is not interrupt-safe
- However, all these functions are thread-safe