


Peterson's lock: A derivation

- Combines ideas from Like-to and Selfless locks
- Two variables:
 - in_i : thread T_i is executing in CS, or trying to do so
 - $turn$: id of thread allowed to enter CS under contention ($turn = i \Leftrightarrow victim = 1 - i$)

28

Establishing mutual exclusion


- Claim:** If the following **invariant** holds when T_i is in CS, so does mutual exclusion



29

Establishing mutual exclusion

- Claim:** If the following **invariant** holds when T_i is in CS, so does mutual exclusion



Why?

- If we instantiate the invariant for both T_1 and T_2 ...
- ...and take the conjunction of the two invariants (i.e., assume both threads in CS)...
- ...it boils down to $(turn = 0 \wedge turn = 1) = false$

30

Towards a solution

- The problem then reduces to establishing the following:

$$in_i \wedge (\neg in_j \vee (in_j \wedge turn = i)) = in_i \wedge (\neg in_j \vee turn = i)$$

- How can we do that?

31

Towards a solution

- The problem then boils down to establishing the following:

$$in_i \wedge (\neg in_j \vee (in_j \wedge turn = i)) = in_i \wedge (\neg in_j \vee turn = i)$$

- How can we do that?

```
lock.acquire() : in_i := true
                while (in_j ∧ turn ≠ i);
```

if we can get past this loop, $(\neg in_j \vee turn \neq i)$ must hold!

32

Establishing the invariant

Thread T₀

```
while(!terminate) {
  in_0 := true
  {in_0}

  while (in_1 ∧ turn ≠ 0);
  {in_0 ∧ (¬in_1 ∨ turn = 0)}
  CS_0
  ...
  in_0 := false
  NCS_0
}
```

Thread T₁

```
while(!terminate) {
  in_1 := true
  {in_1}

  while (in_0 ∧ turn ≠ 1);
  {in_1 ∧ (¬in_0 ∨ turn = 1)}
  CS_1
  ...
  in_1 := false
  NCS_1
}
```

33

Establishing the invariant

- Add assignment to *turn* to establish second disjunct

Thread T₀

```
while(!terminate) {
  in_0 := true
  {in_0}

  while (in_1 ∧ turn ≠ 0);
  {in_0 ∧ (¬in_1 ∨ turn = 0)}
  CS_0
  ...
  in_0 := false
  NCS_0
}
```

Thread T₁

```
while(!terminate) {
  in_1 := true
  {in_1}

  turn := 0
  while (in_0 ∧ turn ≠ 1);
  {in_1 ∧ (¬in_0 ∨ turn = 1)}
  CS_1
  ...
  in_1 := false
  NCS_1
}
```

34

Establishing the invariant

- Add assignment to *turn* to establish second disjunct

Thread T₀

```
while(!terminate) {
  in_0 := true
  {in_0}

  turn := 1
  while (in_1 ∧ turn ≠ 0);
  {in_0 ∧ (¬in_1 ∨ turn = 0)}
  CS_0
  ...
  in_0 := false
  NCS_0
}
```

Thread T₁

```
while(!terminate) {
  in_1 := true
  {in_1}

  turn := 0
  while (in_0 ∧ turn ≠ 1);
  {in_1 ∧ (¬in_0 ∨ turn = 1)}
  CS_1
  ...
  in_1 := false
  NCS_1
}
```

35

Are we there yet?

Thread T₀

```

while(!terminate) {
  in0 := true
  {in0}
  turn := 1
  while (in1 ∧ turn ≠ 0);
  {in0 ∧ (¬in1 ∨ turn = 0) ∨ at(α1)}
  CS0
  ...
  in0 := false
  NCS0
}

```

Thread T₁

```

while(!terminate) {
  in1 := true
  {in1}
  turn := 0
  while (in0 ∧ turn ≠ 1);
  {in1 ∧ (¬in0 ∨ turn = 1)}
  CS1
  ...
  in1 := false
  NCS1
}

```

Diagram: T₁'s PC → auxiliary variable → α₁

36

...and we are gold!

Thread T₀

```

while(!terminate) {
  in0 := true
  {in0}
  α0 turn := 1
  while (in1 ∧ turn ≠ 0);
  {in0 ∧ (¬in1 ∨ turn = 0 ∨ at(α1))}
  CS0
  ...
  in0 := false
  NCS0
}

```

Thread T₁

```

while(!terminate) {
  in1 := true
  {in1}
  α1 turn := 0
  while (in0 ∧ turn ≠ 1);
  {in1 ∧ (¬in0 ∨ turn = 1 ∨ at(α0))}
  CS1
  ...
  in1 := false
  NCS1
}

```

37

Peterson: Safety

Thread T₀

```

while(!terminate) {
  in0 := true
  {in0}
  α0 turn = 1
  while (in1 ∧ turn ≠ 0);
  {in0 ∧ (¬in1 ∨ turn = 0 ∨ at(α1))}
  CS0
  ...
  in0 := false
  NCS0
}

```

Thread T₁

```

while(!terminate) {
  in1 := true
  {in1}
  α1 turn = 0
  while (in0 ∧ turn ≠ 1);
  {in1 ∧ (¬in0 ∨ turn = 1 ∨ at(α0))}
  CS1
  ...
  in1 := false
  NCS1
}

```

If both in the critical section, then:

$$in_0 \wedge (\neg in_1 \vee turn = 0 \vee at(\alpha_1)) \wedge in_1 \wedge (\neg in_0 \vee turn = 1 \vee at(\alpha_0)) \wedge \neg at(\alpha_0) \wedge \neg at(\alpha_1) = (turn = 0) \wedge (turn = 1) = false$$

38

Peterson: Non-blocking

Thread T₀

```

while(!terminate) {
  {R1 : ¬in0 ∧ (turn = 1 ∨ turn = 0)}
  in0 = true
  {R2 : in0 ∧ (turn = 1 ∨ turn = 0)}
  α0 turn = 1
  while (in1 ∧ turn ≠ 0);
  {R3 : in0 ∧ (¬in1 ∨ turn = 0 ∨ at(α1))}
  CS0
  {R3}
  in0 = false
  {R1}
  NCS0
}

```

Thread T₁

```

while(!terminate) {
  {S1 : ¬in1 ∧ (turn = 1 ∨ turn = 0)}
  in1 := true
  {S2 : in1 ∧ (turn = 1 ∨ turn = 0)}
  α1 turn := 0
  {S2}
  while (in0 ∧ turn ≠ 1);
  {S3 : in1 ∧ (¬in0 ∨ turn = 1 ∨ at(α0))}
  CS1
  {S3}
  in1 = false
  {S1}
  NCS1
}

```

Diagram: T₀'s PC → T₁'s PC

Blocking Scenario: T₀ before NCS₀, T₁ stuck at while loop

$R_1 \wedge S_2 \wedge in_0 \wedge (turn = 0) = \neg in_0 \wedge in_1 \wedge in_0 \wedge (turn = 0) = false$

39

Peterson: Deadlock-free

```

while(!terminate) {
  {R1 : ¬in0 ∧ (turn = 1 ∨ turn = 0)}
  in0 = true
  {R2 : in0 ∧ (turn = 1 ∨ turn = 0)}
  α0 turn = 1
  {R2}
  T0's PC → while(in1 ∧ turn ≠ 0);
  {R3 : in0 ∧ (¬in1 ∨ turn = 0 ∨ at(α1))}
  CS0
  {R3}
  in0 = false
  {R1}
  NCS0
}

while(!terminate) {
  {S1 : ¬in1 ∧ (turn = 1 ∨ turn = 0)}
  in1 := true
  {S2 : in1 ∧ (turn = 1 ∨ turn = 0)}
  α1 turn := 0
  {S2}
  T1's PC → while(in0 ∧ turn ≠ 1);
  {S3 : in1 ∧ (¬in0 ∨ turn = 1 ∨ at(α0))}
  CS1
  {S3}
  in1 = false
  {S1}
  NCS1
}

```

Blocking Scenario: T₀ and T₁ at the while loop, before entering critical section

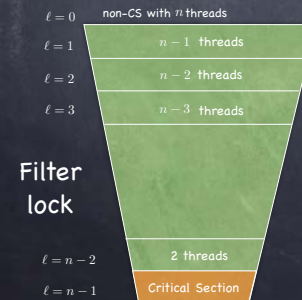
$R_2 \wedge S_2 \wedge in_1 \wedge (turn = 1) \wedge in_0 \wedge (turn = 0) \Rightarrow (turn = 0) \wedge (turn = 1) = false$

40

Taking Stock – I

- The critical section problem is very subtle...
- ...and our solution only covers two threads!

□ with multiple threads it gets even harder!




Bakery algorithm



41

Taking Stock – II

- Our current solutions implement spinlocks
 - threads busy-wait while waiting for lock to be released
 - wastes cycle
 - can lead to **priority inversion**



A Better Way

- How can we do better?
 - Use hardware to **support atomic operations beyond load and store**
 - **Define higher-level programming abstractions** that leverage hardware support

43

Only
on

Disabling Interrupts for Mutual Exclusion

uni-
processors

```
lock.acquire() { disable interrupts;
lock.release() { enable interrupts;
```

- Simple, but flawed
 - thread may never give up CPU!
 - even if it does, it could take too long to respond to an interrupt

44

Only
on


Disabling Interrupts: A Refinement

uni-
processors

- Use a variable to implement the lock; enforce mutual exclusion only on the operations that test and set that variable

```
class lock { int value := FREE }
```

```
lock.acquire() {
  Disable interrupts;
  while (value == BUSY) {
    Enable interrupts;
    Disable interrupts;
  }
  value := BUSY;
  Enable interrupts;
}
```



```
lock.release() {
  Disable interrupts;
  value := FREE;
  Enable interrupts;
}
```

45

Only
on

Disabling Interrupts: No Busy-Waiting

uni-
processors

- Keep a queue of threads waiting for the lock

```
class lock { int value := FREE }
```

```
lock.acquire() {
  Disable interrupts;
  if (value == BUSY) {
    Put on queue of threads
    waiting for lock;
    Go to sleep;
  } else {
    value := BUSY;
  }
  Enable interrupts;
}
```

```
lock.release() {
  Disable interrupts;
  if anyone on wait queue {
    Take a waiting thread off wait queue
    and put it at the front of ready queue;
  } else {
    value := FREE;
  }
  Enable interrupts;
}
```

46

Only
on

Disabling Interrupts: No Busy-Waiting

uni-
processors

- Keep a queue of threads waiting for the lock

```
class lock { int value := FREE }
```

```
lock.acquire() {
  Disable interrupts;
  if (value == BUSY) {
    Put on queue of threads
    waiting for lock;
  }
```

```
lock.release() {
  Disable interrupts;
  if anyone on wait queue {
    Take a waiting thread off wait queue
    and put it at the front of ready queue;
  } else {
    value := FREE;
  }
  Enable interrupts;
}
```

47

Only
on

Disabling Interrupts: No Busy-Waiting

uni-
processors

- Keep a queue of threads waiting for the lock

```
class lock { int value := FREE }
```

```
lock.acquire() {  
  Disable interrupts;  
  if (value == BUSY) {  
    Enable interrupts?  
    Put on queue of threads  
    waiting for lock;  
  }
```

```
lock.release() {  
  Disable interrupts;  
  if anyone on wait queue {  
    Take a waiting thread off wait queue  
    and put it at the front of ready queue;  
  } else {  
    value := FREE;  
  }  
  Enable interrupts;  
}
```

48

Only
on

Disabling Interrupts: No Busy-Waiting

uni-
processors

- Keep a queue of threads waiting for the lock

```
class lock { int value := FREE }
```

```
lock.acquire() {  
  Disable interrupts;  
  if (value == BUSY) {  
    Enable interrupts?  
    Put on queue of threads  
    waiting for lock;  
  }
```

NO!

Suppose T_2 concurrently
calls `lock.release`

- T_2 could check the queue before
 T_1 is placed on the queue
- T_1 would be waiting on the
queue even if lock is free

49

Only
on

Disabling Interrupts: No Busy-Waiting

uni-
processors

- Keep a queue of threads waiting for the lock

```
class lock { int value := FREE }
```

```
lock.acquire() {  
  Disable interrupts;  
  if (value == BUSY) {  
    Put on queue of threads  
    waiting for lock;  
    Enable interrupts?  
    Go to sleep;  
  } else {  
    value := BUSY;  
  }  
  Enable interrupts;  
}
```

```
lock.release() {  
  Disable interrupts;  
  if anyone on wait queue {  
    Take a waiting thread off wait queue  
    and put it at the front of ready queue;  
  } else {  
    value := FREE;  
  }  
  Enable interrupts;  
}
```

50

Only
on

Disabling Interrupts: No Busy-Waiting

uni-
processors

- Keep a queue of threads waiting for the lock

```
class lock { int value := FREE }
```

```
lock.acquire() {  
  Disable interrupts;  
  if (value == BUSY) {  
    Put on queue of threads  
    waiting for lock;  
    Enable interrupts?  
    Go to sleep;  
  }
```

NO!

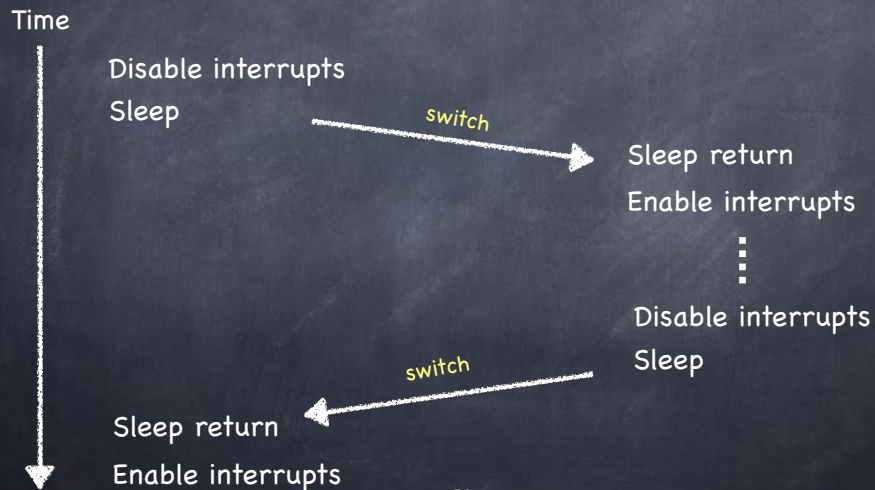
Suppose T_2 put T_1 on
the ready queue

- When T_1 runs it will
immediately go to sleep!!

When are interrupts re-enabled??

51

Passing the baton



52

Atomic Read/Modify/Write

- On a multiprocessor, disabling interrupts does not ensure atomicity
 - other CPUs could still enter the critical section
 - costly to disable interrupts on all CPUs
- Hardware provides special machine instructions
 - Test-and-Set (TAS)**
 - TAS reads value of a memory location, writes back 1 in its place
 - Compare-and-Swap (CAS)**
 - CAS compares contents of a memory location with a given value; if they are the same, set the memory location to a new given value
 - Load link/Store Conditional (LL/SC)**
 - LL reads value of a memory location; SC writes a new value to that location only if it has not been updated since LL

53

Hey! What's wrong with this?

```
lock value := 0
lock.acquire() {while (value = 1);}
lock.release() value := 0;
```


These new hardware-supported instructions can change atomically the variable that implements the lock

54

TAS Locks

- Simple, but flawed

```
lock value := 0
lock.acquire() { while (TAS(value) = 1); }
lock.release() { value := 0; }
```



well, certainly not for time needed to complete an arbitrarily long CS!

55

TAS Locks: Try 2

Going meta!

- Use TAS (on a second variable **guard**) to build a new "secondary CS" where the variable **value** that stores status of the lock associated with the original ("primary") CS can be atomically read and set.

Why bother?

- Instead on spinning on TAS for the unknown length of the primary CS, we are now spinning at worst for the known, short length of the secondary CS!

56

TAS Locks: Try 2

```
lock.acquire() {  
    while (TAS(guard)) ;  
    if (value == BUSY) {  
        Put on queue of threads  
        waiting for lock;  
        Set guard to 0;  
        Go to sleep;  
    } else {  
        value := BUSY;  
        guard := 0  
    }  
}
```

```
lock.release() {  
    while (TAS(guard)) ;  
    if anyone on wait queue {  
        Take a waiting thread off lock's  
        wait queue and put it at the  
        front of ready queue;  
    } else {  
        value := FREE;  
    }  
    guard := 0  
}
```

57