

# Concurrency Primitives

Soham Chakraborty

11.02.2022

## Recap

- Previous example

## Concurrency Primitives

- Synchronization
- C/C++ primitives
- Mutex & Lock

## Properties & Reasoning

## Recap: Previous Example

*Init*

<code>lock(m)</code>		<code>lock(m)</code>
<code>S<sub>1</sub>;</code>		<code>S<sub>2</sub>;</code>
<code>unlock(m)</code>		<code>unlock(m)</code>

*Init*

<code>lock(m)</code>		<code>lock(m)</code>
<code>S<sub>1</sub>;</code>		<code>S<sub>2</sub>;</code>
<code>unlock(m)</code>		<code>unlock(m)</code>

Lock/Unlock

- How is lock/unlock defined?
- What are the properties?

*Init*

lock( <i>m</i> )		lock( <i>m</i> )
critical section		critical section
unlock( <i>m</i> )		unlock( <i>m</i> )

# Mutual Exclusion

*Init*

lock( <i>m</i> )		lock( <i>m</i> )
critical section		critical section
unlock( <i>m</i> )		unlock( <i>m</i> )

*m*: lock object

critical section: code block between lock(*m*) and unlock(*m*)

Mutual execution property: critical sections can be executed by only one thread at a time

# Thread Communication & Synchronization

Threads communicate to communicate information

```
*X = NULL, flag = 0;

X = new Obj();
flag = 1;
    || while(flag != 1) // waiting....
    || ;
    || t = X; // must be non-NULL
```

What happens if we reorder the statements in the first thread?

# Thread Communication & Synchronization

$*X = NULL, flag = 0;$

$*X = NULL, Y = 0;$

$X = new\ Obj();$   
 $flag = 1;$   $\left\| \begin{array}{l} while(flag \neq 1) \\ ; \\ t = X; \end{array} \right.$

$\rightsquigarrow$

$flag = 1;$   
 $X = new\ Obj();$   $\left\| \begin{array}{l} while(flag \neq 1) \\ ; \\ t = X; \end{array} \right.$

$t = NULL$  is NOT possible

$t = NULL$  is possible



*Init*

<code>lock(m);</code>		<code>lock(m);</code>
<code>X = X + 1;</code>		<code>X = X + 1;</code>
<code>unlock(m);</code>		<code>unlock(m);</code>

Can we reorder instructions with lock/unlock?

- $X$  and  $m$  are independent variables

*Init*

<code>lock(m);</code>		<code>lock(m);</code>
<code>unlock(m);</code>		<code>X = X + 1;</code>
<code>X = X + 1;</code>		<code>unlock(m);</code>

Is it problematic?

*Init*


<code>lock(m);</code>		<code>lock(m);</code>
<code>unlock(m);</code>		<code>X = X + 1;</code>
<code>X = X + 1;</code>		<code>unlock(m);</code>

Is it problematic?


Safe reordering in sequential program  
may be unsafe for concurrency

# Roach Motel Reordering

Safe



Unsafe



How to restrict these reorderings?

- Concurrency primitives

# C/C++ Concurrency Primitives

Introduced in 2011 C/C++ standard.

Provides platform independent abstraction

Consistency rules.

# Shared Memory Accesses

Non-atomic accesses: Read (Ld), Write (St)

Atomic accesses = operation + memory order

Operations:

- Read (Ld)
- Write (St)
- Atomic update (U)
- Fence (F)

Memory orders:

- Relaxed (rlx)
- Release (rel)
- Acquire (acq)
- Acquire-Release (acq\_rel)
- Sequentially consistent (sc)

# Shared Memory Accesses

Non-atomic accesses: Read (Ld), Write (St)

Atomic accesses = operation + memory order

Operations:

- Read (Ld)
- Write (St)
- Atomic update (U)
- Fence (F)

Memory orders:

- Relaxed (rlx)
- Release (rel)
- Acquire (acq)
- Acquire-Release (acq\_rel)
- Sequentially consistent (sc)

Example:

- X.load(memory order)
- X.store(val, memory order)
- X.CAS(oldval, nwval, success mem order, failure mem order)
- atomic\_thread\_fence(memory order)

# Access Types

Read.  $t = X_o$   
where  $o \in \{\text{na}, \text{rlx}, \text{acq}, \text{sc}\}$

Write.  $X_o = v$   
where  $o \in \{\text{na}, \text{rlx}, \text{acq}, \text{sc}\}$

Update.  $\text{CAS}(X, v, v', o_s, o_f)$   
where  $o_s, o_f \in \{\text{rlx}, \text{rel}, \text{acq}, \text{acq\_rel}, \text{U}_{\text{sc}}\}$

Fence  $F_o$   
where  $o \in \{\text{rel}, \text{acq}, \text{acq\_rel}, \text{sc}\}$

For now we consider only sc accesses



# Reordering Rules

$a(\ell) \downarrow / b(\ell') \Rightarrow$	$Ld_{na}/St_{na}$	$St_{sc}$	$Ld_{sc}$
$Ld_{na}/St_{na}$	✓	✗	✓
$St_{sc}$	✓	✗	✗
$Ld_{sc}$	✗	✗	✗

$a; b \rightsquigarrow b; a$  where  $\ell \neq \ell'$  and are independent

## Reordering Rules

$a(\ell) \downarrow / b(\ell') \Rightarrow$	Ld <sub>na</sub> /St <sub>na</sub>	St <sub>sc</sub>	Ld <sub>sc</sub>
Ld <sub>na</sub> /St <sub>na</sub>	✓	✗	✓
St <sub>sc</sub>	✓	✗	✗
Ld <sub>sc</sub>	✗	✗	✗

$a; b \rightsquigarrow b; a$  where  $\ell \neq \ell'$  and are independent

Revisiting the (simplified) synchronization example:

```
int X = 0,  
atomic_int flag = 0; // all accesses are sc
```

```
X = 1;    || while(flag ≠ 1)  
flag = 1; ||     ;  
          ||     t = X;
```

## Reordering Rules

$a(\ell) \downarrow / b(\ell') \Rightarrow$	Ld <sub>na</sub> /St <sub>na</sub>	St <sub>sc</sub>	Ld <sub>sc</sub>
Ld <sub>na</sub> /St <sub>na</sub>	✓	✗	✓
St <sub>sc</sub>	✓	✗	✗
Ld <sub>sc</sub>	✗	✗	✗

$a; b \rightsquigarrow b; a$  where  $\ell \neq \ell'$  and are independent

Revisiting the (simplified) synchronization example:

```
int X = 0,  
atomic_int flag = 0; // all accesses are sc
```

```
X = 1;    || while(flag ≠ 1)  
flag = 1; || ;  
          || t = X;
```

Note: the atomic accesses can be used to implement lock & unlock

```
void lock(){  
    i = tid();  
    j = 1 - i;  
    flag[i] = true; // I'm interested  
    while(flag[j]) {} // wait loop  
}
```

```
unlock() {  
    i = tid();  
    flag[i] = false; // I'm not interested  
}
```

Ensures mutual exclusion (Critical sections do not overlap)

A thread is well-formed if:

- ① each critical section is associated with a unique lock object.
- ② the thread calls `lock` method for that object when it is trying to enter the critical section, and
- ③ the thread calls the `unlock` method for that object when it leaves the critical section.

Threads are state machines:  $a_0 \rightarrow a_1 \rightarrow \dots$

Thread transitions are events

$(a_0, a_1)$ : Interval between events  $a_0$  and  $a_1$

$I_A = (a_0, a_1)$ : interval between  $a_0$  and  $a_1$  in thread  $A$

$I_B = (b_0, b_1)$ : interval between  $a_0$  and  $a_1$  in thread  $B$

$I_A \rightarrow I_B$ : interval  $I_A$  precedes  $I_B$ ; when  $a_1 \rightarrow b_0$

$a_i^j$ :  $j^{\text{th}}$  occurrence of an event  $a_i$

$I_A^j$ :  $j^{\text{th}}$  occurrence of an interval  $I_A$

# Mutual Exclusion

Critical sections do not overlap.

Given thread  $A$  and  $B$

Given the intervals  $CS_A^i$  and  $CS_B^j$ :

either  $CS_A^i \rightarrow CS_B^j$  or  $CS_B^j \rightarrow CS_A^i$

## Violation of Mutual Exclusion

```
void lock(){  
    i = tid();  
    j = 1 - i;  
    flag[i] = true;  
    while(flag[j]) {}  
}
```

```
unlock() {  
    i = tid();  
    flag[i] = false;  
}
```



# Violation of Mutual Exclusion

```
void lock(){  
    i = tid();  
    j = 1 - i;  
    flag[i] = true;  
    while(flag[j]) {}  
}
```

$CS_0^j \not\rightarrow CS_1^k$  and  $CS_1^k \not\rightarrow CS_0^j$

```
unlock() {  
    i = tid();  
    flag[i] = false;  
}
```

# Violation of Mutual Exclusion

```
void lock(){  
    i = tid();  
    j = 1 - i;  
    flag[i] = true;  
    while(flag[j]) {}  
}
```

```
unlock() {  
    i = tid();  
    flag[i] = false;  
}
```

$CS_0^j \not\rightarrow CS_1^k$  and  $CS_1^k \not\rightarrow CS_0^j$

a:  $W_0(flag[0], true)$



b:  $R_0(flag[1], false)$



$CS_0$

# Violation of Mutual Exclusion

```
void lock(){  
    i = tid();  
    j = 1 - i;  
    flag[i] = true;  
    while(flag[j]) {}  
}
```

```
unlock() {  
    i = tid();  
    flag[i] = false;  
}
```

$CS_0^j \not\rightarrow CS_1^k$  and  $CS_1^k \not\rightarrow CS_0^j$

a:  $W_0(flag[0], true)$

c:  $W_1(flag[1], true)$

b:  $R_0(flag[1], false)$

d:  $R_1(flag[0], false)$

$CS_0$

$CS_1$

# Violation of Mutual Exclusion

```
void lock(){  
    i = tid();  
    j = 1 - i;  
    flag[i] = true;  
    while(flag[j]) {}  
}
```

```
unlock() {  
    i = tid();  
    flag[i] = false;  
}
```

$CS_0^j \not\rightarrow CS_1^k$  and  $CS_1^k \not\rightarrow CS_0^j$

a:  $W_0(flag[0], true)$

c:  $W_1(flag[1], true)$

b:  $R_0(flag[1], false)$

d:  $R_1(flag[0], false)$

$CS_0$

$CS_1$



# Violation of Mutual Exclusion

```
void lock(){  
    i = tid();  
    j = 1 - i;  
    flag[j] = true;  
    while(flag[j]) {}  
}
```

```
unlock() {  
    i = tid();  
    flag[i] = false;  
}
```

$CS_0^j \not\rightarrow CS_1^k$  and  $CS_1^k \not\rightarrow CS_0^j$

a:  $W_0(flag[0], true)$

c:  $W_1(flag[1], true)$

b:  $R_0(flag[1], false)$

d:  $R_1(flag[0], false)$

$CS_0$

$CS_1$

Contradiction:  $a \rightarrow d$  and no intermediate  $W(flag[0], false)$  between  $a$  and  $b$ .

## Alternative Lock Unlock

```
void lock(){  
    i = tid();  
    victim = i; // let the other go first  
    while(victim == i) {} // wait  
}
```

```
unlock() {}
```

Ensures mutual exclusion

# Mutual Exclusion

```
void lock(){  
    i = tid();  
    victim = i;  
    while(victim == i)  
        {}  
}
```

```
unlock() {}
```

# Mutual Exclusion

$CS_0^j \not\rightarrow CS_1^k$  and  $CS_1^k \not\rightarrow CS_0^j$

```
void lock(){  
    i = tid();  
    victim = i;  
    while(victim == i)  
        {}  
}
```

```
unlock() {}
```



# Mutual Exclusion

```
void lock(){  
    i = tid();  
    victim = i;  
    while(victim == i)  
        {}  
}
```

```
unlock() {}
```

$CS_0^j \not\rightarrow CS_1^k$  and  $CS_1^k \not\rightarrow CS_0^j$

$W_0(victim, 0)$



$R_0(victim, 1)$



$CS_0$

# Mutual Exclusion

```
void lock(){  
    i = tid();  
    victim = i;  
    while(victim == i)  
        {}  
}
```

```
unlock() {}
```

$CS_0^j \not\rightarrow CS_1^k$  and  $CS_1^k \not\rightarrow CS_0^j$

$W_0(victim, 0)$

$W_1(victim, 1)$

$R_0(victim, 1)$

$R_1(victim, 0)$

$CS_0$


$CS_1$

# Mutual Exclusion

```
void lock(){  
    i = tid();  
    victim = i;  
    while(victim == i)  
        {}  
}
```

```
unlock() {}
```

$CS_0^j \not\rightarrow CS_1^k$  and  $CS_1^k \not\rightarrow CS_0^j$




# Mutual Exclusion

```
void lock(){  
    i = tid();  
    victim = i;  
    while(victim == i)  
        {}  
}
```

```
unlock() {}
```

$CS_0^j \not\rightarrow CS_1^k$  and  $CS_1^k \not\rightarrow CS_0^j$



Problem: What if the threads are not running concurrently?

# Peterson's Lock Algorithm

```
void lock(){  
    i = tid();  
    j = 1 - i;  
    flag[i] = true; // I am interested  
    victim = i; // you go first  
    while(flag[j] && victim == i) {} // I am waiting  
}
```

```
unlock() {  
    i = tid();  
    flag[i] = false; // I am not interested  
}
```

# Peterson's Lock Algorithm

```
void lock(){  
    i = tid();  
    j = 1 - i;  
    flag[i] = true;  
    victim = i;  
    while(flag[j] &&  
          victim == i)  
        {}  
}
```

```
unlock() {  
    i = tid();  
    flag[i] = false;  
}
```

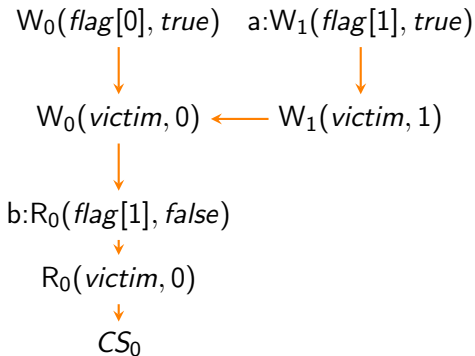
w.l.o.g assume thread 0  
writes to `victim`



# Peterson's Lock Algorithm

```
void lock(){
  i = tid();
  j = 1 - i;
  flag[i] = true;
  victim = i;
  while(flag[j] &&
        victim == i)
    {}
}
```

```
unlock() {
  i = tid();
  flag[i] = false;
}
```



w.l.o.g assume thread 0  
writes to **victim**





# Properties of Mutual Exclusion

(Essential) **Deadlock freedom.** If one or multiple process trying to enter critical section, then some process eventually will enter critical section

**Starvation freedom.** Any process/thread trying to enter critical section, will eventually enter critical section

- Requires fairness

# Properties of Mutual Exclusion

(Essential) **Deadlock freedom.** If one or multiple process trying to enter critical section, then some process eventually will enter critical section

**Starvation freedom.** Any process/thread trying to enter critical section, will eventually enter critical section

- Requires fairness

## Waiting.

If one process/thread delays in critical section then other processes/threads also get delayed

- What if a thread acquires a lock and crashes?
- Requires fault tolerance



## Alternative Lock Unlock

```
void lock(){  
    i = tid();  
    victim = i; // let the other go first  
    while(victim == i) {} // wait  
}
```

```
unlock() {}
```

Is it deadlock-free?

It deadlocks if one thread runs completely before the other

# Peterson's Lock Algorithm

```
void lock(){
    i = tid();
    j = 1 - i;
    flag[i] = true; // I am interested
    victim = i; // you go first
    while(flag[j] && victim == i) {} // I am waiting
}
```

```
unlock() {
    i = tid();
    flag[i] = false; // I am not interested
}
```

It is deadlock free and starvation free.

## Nested Locking

<i>lock(m<sub>1</sub>);</i>		<i>lock(m<sub>2</sub>);</i>
<i>lock(m<sub>2</sub>);</i>		<i>lock(m<sub>1</sub>);</i>
:		:
<i>unlock(m<sub>2</sub>);</i>		<i>unlock(m<sub>1</sub>);</i>
<i>unlock(m<sub>1</sub>);</i>		<i>unlock(m<sub>2</sub>);</i>

May lead to deadlock

## Nested Locking

<i>lock(m<sub>1</sub>);</i>		<i>lock(m<sub>2</sub>);</i>
<i>lock(m<sub>2</sub>);</i>		<i>lock(m<sub>1</sub>);</i>
:		:
<i>unlock(m<sub>2</sub>);</i>		<i>unlock(m<sub>1</sub>);</i>
<i>unlock(m<sub>1</sub>);</i>		<i>unlock(m<sub>2</sub>);</i>

May lead to deadlock

Solution: whenever threads lock multiple mutexes, they do so in the same order



The Art of Multiprocessor Programming (chapter 2)

2nd Edition - September 8, 2020

Authors: Maurice Herlihy, Nir Shavit, Victor Luchangco, Michael Spear