

# Concurrency Bugs

Soham Chakraborty

16.02.2022

Order violation

Atomicity violation

Deadlock

Data race


# Order Violation

Cause: Programmer assumes certain ordering of events

# Order Violation

Cause: Programmer assumes certain ordering of events

Example:

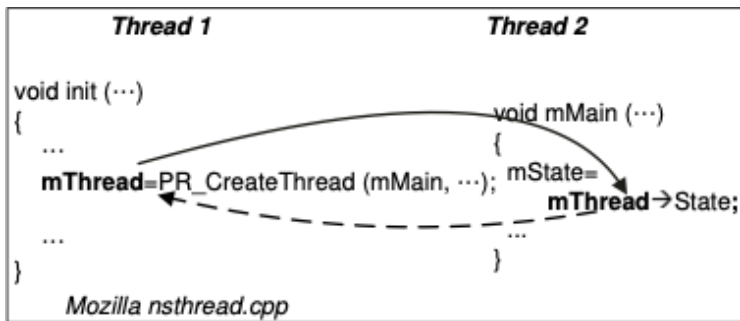


Thread 2 should not deref. **mThread** before Thread 1 initializes it

# Order Violation

Cause: Programmer assumes certain ordering of events

Example:



Thread 2 should not deref. **mThread** before Thread 1 initializes it  
Pattern:

```
X = 0;  
X = 1; || t = X; // 1
```

# Order Violation

Cause: Programmer assumes certain ordering of (W,R) events

Example:

<i>Thread 1</i>	<i>Thread 2</i>
<pre>void js_DestroyContext (...) {   /* last one entering this function */    <b>js_UnpinPinnedAtom(&amp;atoms);</b> }</pre>	<pre>void js_DestroyContext (...) {   /* non-last one entering this   function */    <b>js_MarkAtom(&amp;atoms, ...);</b> }</pre>

*Mozilla jsctx.c, jsgc.c*

`js_UnpinPinnedAtom` should happen after `js_MarkAtom`.

# Order Violation

Cause: Programmer assumes certain ordering of (W,R) events

Example:

<i>Thread 1</i>	<i>Thread 2</i>
<pre>void js_DestroyContext (...) {   /* last one entering this function */    <b>js_UnpinPinnedAtom(&amp;atoms);</b> }</pre>	<pre>void js_DestroyContext (...) {   /* non-last one entering this   function */    <b>js_MarkAtom(&amp;atoms, ...);</b> }</pre>

*Mozilla jsctx.c, jsgc.c*

`js_UnpinPinnedAtom` should happen after `js_MarkAtom`.


Pattern:

```
X = 0;  
X = 1; || t = X; // 0
```

# Order Violation

Cause: Programmer assumes certain ordering of (W,W) events

Example:



Assumption: S1 and S2 execute atomically


Unsafe ordering blocks thread 1



# Order Violation

Cause: Programmer assumes certain ordering of (W,W) events

Example:



Assumption: S1 and S2 execute atomically

Unsafe ordering blocks thread 1


Pattern:

```
X = 1;
while(X == 1) ; // 0 || X = 0;
```

# Atomicity Violation

Cause: Programmer assumes atomicity of certain code regions

Example:




Assumption: S1;S2 are executed atomically

S2 access NULL value

# Atomicity Violation

Cause: Programmer assumes atomicity of certain code regions

Example:



Assumption: S1;S2 are executed atomically

S2 access NULL value

Pattern:


$$\begin{array}{l} X = 0; \\ a = X; \\ b = X; \end{array} \parallel \begin{array}{l} X = 1; \end{array}$$

Desired:  $a = b = 0$  or  $a = b = 1$

# Multi-Variable Atomicity Bugs

Cause: variables are semantically connected which is violated

Example:




Assumption: `mOffset` and `mLength` are updated atomically wrt thread 1

Lack of synchronization  $\Rightarrow$  thread 1 read inconsistent value

# Multi-Variable Atomicity Bugs

Cause: variables are semantically connected which is violated

Example:



Assumption: *mOffset* and *mLength* are updated atomically wrt thread 1

Lack of synchronization  $\Rightarrow$  thread 1 read inconsistent value

$$Y = Z = 0;$$
$$t = X[Y + Z]; \parallel Y = 1; Z = 1;$$

Desired: access  $X[0]$  or  $X[2]$

# Timing Bugs

Cause: Programmer assumes the tasks would complete within certain time period

Example:

<i>Thread 1</i>	<i>Thread 2 ... Thread n</i>	<i>Monitor thread</i>
<pre>void buf_flush_try_page() {     ...     rw_lock(&amp;lock); }</pre>	<pre>rw_lock(&amp;lock);</pre>	<pre>void error_monitor_thread() {     if(lock_wait_time[i] &gt;         fatal_timeout)         assert(0, "We crash the server;         It seems to be hung."); }</pre>
<i>MySQL buf0flu.c</i>		<i>MySQL srv0srv.c</i>

Assumption:  $n$  tasks would complete before *fatal\_timeout*

Crash the server

Understand the semantics

Add/modify locks

Add/modify synchronizations

Revisit the examples

# Deadlock

A thread holds a lock and wait for another lock held by another thread and vice versa

<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>



## Deadlock: Another Scenario

Another challenge: encapsulation

```
Vector v1, v2;  
v1.AddAll(v2); || v2.AddAll(v1);
```

# Conditions for Deadlock

All conditions must hold:

- **Mutual exclusion:** Threads claim exclusive control of resources (e.g. lock) that they require.
- **Hold-and-wait:** Threads hold allocated resources while waiting for additional resources
- **No preemption:** Held resources cannot be forcibly removed from threads
- **Circular wait:** There exists a circular chain of threads where each thread holds a resource that are being requested by the next thread in the chain.

# Deadlock Prevention

**Prevent circular wait** Programming convention: total ordering on acquiring lock

- Prone to mistakes

**Prevent hold-and-wait** Acquire all locks at once

- Decreases concurrency significantly

## Prevent no-preemption

Hold locks only when all the locks are available

Challenge: encapsulation prevents the 'top' loop implementation

```
top :  
lock(L1);  
if(trylock(L2) == -1) {  
    unlock(L1);  
    goto top;  
}
```

# Deadlock Prevention

## Prevent no-preemption

Hold locks only when all the locks are available

Challenge: encapsulation prevents the 'top' loop implementation

```
top :
lock(L1);
if(trylock(L2) == -1) {
    unlock(L1);
    goto top;
}

|||

top :
lock(L2);
if(trylock(L1) == -1) {
    unlock(L2);
    goto top;
}
```

# Deadlock Prevention

## Prevent no-preemption

Hold locks only when all the locks are available

Challenge: encapsulation prevents the 'top' loop implementation

<pre>top : lock(L1); if(trylock(L2) == -1) {     unlock(L1);     goto top; }</pre>		<pre>top : lock(L2); if(trylock(L1) == -1) {     unlock(L2);     goto top; }</pre>
--	--	--

Problem: Livelock

**Prevent circular wait** Total ordering on acquiring lock

- Prone to mistakes

**Prevent hold-and-wait** Acquire all locks at once

- Decreases concurrency significantly

**Prevent no-preemption**

- Problem: Livelock
- Challenge: encapsulation prevents the 'top' loop implementation


**No mutual-exclusion**

- Lock free programming

# Deadlock Avoidance

Schedule threads that access same resources

	T1	T2	T3	T4
L1	yes	yes	no	no
L2	yes	yes	yes	no





## Deadlock Recovery


Deadlock detector automatically detect deadlock

If deadlock is detected; restart system

Event  $a$  and  $b$  is in data race if:

- $a$  and  $b$  are concurrent/in conflict
- $a$  and  $b$  access same location
- At least one of  $a$  and  $b$  is a write

# Concurrent Accesses



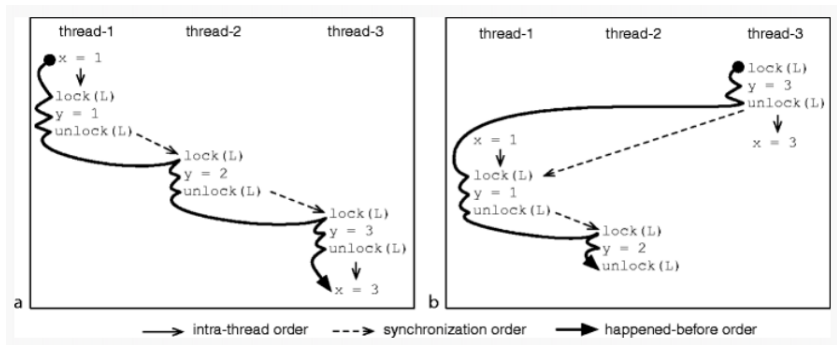
Concurrent:  $(e_1, e_2)$ ,  $(e_2, e_3)$

$e_3$  happens-before  $e_1$

- $end(e_3) \rightarrow start(e_1)$

# Happens-Before

concurrent/conflict  $\Rightarrow$  Not in happens-before (HB) order



Execution 1: No data race

Execution 2: data race on  $x$

## Lockset algorithm

Let  $locks\_held(t)$  be the set of locks held by thread  $t$ .

For each  $v$ , initialize  $C(v)$  to the set of all locks.

On each access to  $v$  by thread  $t$ ,

    set  $C(v) := C(v) \cap locks\_held(t)$ ;

    if  $C(v) = \{ \}$ , then issue a warning.

## Lockset algorithm

Let  $locks\_held(t)$  be the set of locks held by thread  $t$ .

For each  $v$ , initialize  $C(v)$  to the set of all locks.

On each access to  $v$  by thread  $t$ ,  
set  $C(v) := C(v) \cap locks\_held(t)$ ;  
if  $C(v) = \{ \}$ , then issue a warning.

## Example:

<i>Program</i>	<i>locks_held</i>	<i>C(v)</i>
	{}	{mu1, mu2}
lock(mu1);	{mu1}	
v := v+1;		{mu1}
unlock(mu1);	{}	
lock(mu2);	{mu2}	
v := v+1;		{}
unlock(mu2);	{}	

Learning from Mistakes – A Comprehensive Study on Real World Concurrency Bug Characteristics.

Shan Lu, Soyeon Park, Eunsoo Seo and Yuanyuan Zhou  
ASPLOS 2008.

Common Concurrency Problems (chapter 32)

Operating Systems: Three Easy Pieces

Remzi H. Arpaci-Dusseau and Andrea C. Arpaci-Dusseau

<https://pages.cs.wisc.edu/remzi/OSTEP/threads-bugs.pdf>

Race Detection Techniques

Christoph von Praun

[https://doi.org/10.1007/978-0-387-09766-4\\_38](https://doi.org/10.1007/978-0-387-09766-4_38)

Eraser: A Dynamic Data Race Detector for Multithreaded Programs

Stefan Savage, Michael Burrows, Greg Nelson, Patrick Sobalvarro,  
Thomas Anderson. ACM TOCS 1997.