Locking

Youjip Won



Contents

- Outline
- Race condition
- Spinlock
- Using lock
- Deadlock
- Interrupt handler
- sync_synchronize
- Sleep Lock
- Limitations of locks

Race condition

Situation in which a memory location is accessed concurrently, and at least one access is a write



```
struct list {
    int data;
    struct list *next;
};
struct list *list = 0;
void
insert(int data)
{
    struct list *1;
    l = malloc(sizeof *1);
    l->data = data;
    l->next = list;
    list = l;
}
```

 Situation in which a memory location is accessed concurrently, and at least one access is a write



```
struct list {
    int data;
    struct list *next;
};
struct list *list = 0;
void
insert(int data)
{
    struct list *l;
    l = malloc(sizeof *l);
    l->data = data;
    l->next = list;
    list = l;
}
```

 Situation in which a memory location is accessed concurrently, and at least one access is a write



```
struct list {
    int data;
    struct list *next;
};
struct list *list = 0;
void
insert(int data)
{
    struct list *l;
    l = malloc(sizeof *l);
    l->data = data;
    l->next = list;
    list = l;
}
```

 Situation in which a memory location is accessed concurrently, and at least one access is a write



```
struct list {
    int data;
    struct list *next;
};
struct list *list = 0;
void
insert(int data)
{
    struct list *l;
    l = malloc(sizeof *l);
    l->data = data;
    l->next = list;
    list = l;
}
```

 Situation in which a memory location is accessed concurrently, and at least one access is a write



```
struct list {
    int data;
    struct list *next;
};
struct list *list = 0;
void
insert(int data)
{
    struct list *l;
    l = malloc(sizeof *l);
    l->data = data;
    l->next = list;
    list = l;
}
```

Race conditions can be solved by locking

```
struct list *list = 0;
struct lock listlock;
                                                                 Lock L',
void
insert(int data)
                                                                 insert(){
{
  struct list *l;
                                                                   lock (1);
1/code for insert.
  acquire(&listlock);
  l = malloc(sizeof *1);
                                                                    unlode (L)',
  1 - > data = data;
  l->next = list;
                                                                  3
  list = 1;
  release(&listlock);
}
```









xchg(mem,	val);
-----------	-------

- 1. Exchange
- 2. Return the value that was in the memory



Spinlock

- The xv6 has 2 types of locks: spinlock and sleep-lock
- Spinlock structure
- primitives
 - acquire(struct spinlock *lk)
 - release(struct spinlock *lk)

acquire in spinlock

```
void
acquire(struct spinlock *lk)
  pushcli(); // disable interrupts to avoid deadlock.
  if(holding(lk))
   panic("acquire");
 // The xchg is atomic.
 while(xchg(&lk->locked, 1) != 0)
    ;
  // Tell the C compiler and the processor not to move loads or stores
  // past this point, to ensure that the critical section's memory
  // references happen after the lock is acquired.
   sync synchronize();
 // Record info about lock acquisition for debugging.
 lk->cpu = mycpu();
  getcallerpcs(&lk, lk->pcs);
```

Spinlock (Cont.)

```
9 getcallerpcs(void *v, uint pcs[])
```

Spinlock (Cont.)

- getcallerpcs() follows the latest ten stack frames, and records the caller address (return address) to pcs.
- XV6 uses getcallerpcs() for debugging purpose.
- Unrecorded elements of pcs array are set to 0.

Spinlock (Cont.)

getcallerpcs (void *v, uint pcs[])



```
void
release(struct spinlock *lk)
  if(!holding(lk))
    panic("release");
 lk - pcs[0] = 0;
  lk \rightarrow cpu = 0;
  // Tell the C compiler and the processor to not move loads or stores
  // past this point, to ensure that all the stores in the critical
  // section are visible to other cores before the lock is released.
  // Both the C compiler and the hardware may re-order loads and
  // stores; sync synchronize() tells them both not to.
  sync synchronize();
  // Release the lock, equivalent to lk->locked = 0.
  // This code can't use a C assignment, since it might
  // not be atomic. A real OS would use C atomics here.
  asm volatile("movl $0, %0" : "+m" (lk->locked) : );
 popcli();
```

Caution

- Locking should be used properly, not too much.
- Too many locks reduce parallelism.
- xv6 uses a few data-structure specific locks.

Lock	Description
bcache.lock	Protects allocation of block buffer cache entries
cons.lock	Serializes access to console hardware, avoids intermixed output
ftable.lock	Serializes allocation of a struct file in file table
icache.lock	Protects allocation of inode cache entries
idelock	Serializes access to disk hardware and disk queue
kmem.lock	Serializes allocation of memory
log.lock	Serializes operations on the transaction log
pipe's p->lock	Serializes operations on each pipe
ptable.lock	Serializes context switching, and operations on proc->state and proctable
tickslock	Serializes operations on the ticks counter
inode's ip->lock	Serializes operations on each inode and its content
buf's b->lock	Serializes operations on each block buffer

Deadlock

Deadlock: The situation that two threads require each other's lock and wait indefinitely.



To avoid deadlock, the callers must invoke functions in a way that causes locks to be acquired in the agreed-on order.



Interrupt and lock

Spinlocks are used by the interrupt handler as well as the user thread.



• Important!!!

- You have to disable the interrupt when you hold the spinlock that is used by the interrupt handler.
- It is difficult to determine whether the given lock is used by the interrupt handler or not.
- XV6 blindly disables the interrupt when you hold the spinlock.
- Turn off interrupt before you acquire spinlock.





```
void
acquire(struct spinlock *lk)
{
    pushcli(); // disable interrupts to avoid deadlock. Disable interrupt.
    if(nolding(lk))
    panic("acquire");
    // The xchg is atomic.
    while(xchg(&lk->locked, 1) != 0)
    ;
    // Release the lock, equivalent to lk->locked = 0.
    // This code can't use a C assignment, since it might
    // not be atomic. A real OS would use C atomics here.
    asm volatile("movl $0, %0" : "+m" (lk->locked) : );
    popcli();
    Enable interrupt.
```

Disable and enable interrupt

	• pushcli/popcli are lik	<mark>e</mark> cli/sti excep	ot that they are ma	atched: it
	takes two popcli to undo	two pushcli.		
void pushcli(void)	• If interrupts are off, then pa	ishcli, popcli lea	ves them off.	
{ int eflags;				
<pre>eflags = readeflags cli();</pre>	();			
<pre>if(mycpu()->ncli == 0)</pre>				
<pre>mycpu()->intena = mycpu()->ncli += 1; }</pre>	erlags & FL_IF;			
void				
popcli(void)				
{ if(readeflags()&FL_3	IF)			
panic("popcli - in if(mycpu()->ncli panic("popcli");	nterruptible"); < 0)			
if(mycpu()->ncli ==	0 && mycpu()->intena)			
sti(); }	#define FL_IF	0x00000200	// Interrupt	Enable

Instructions and Memory Ordering



- Compilers may reorder instructions.
- Many complier and processor execute code out of order to achieve higher performance.

Instructions and Memory Ordering (Cont.)

• There is bad example. Line 6 can runs before Line 3, 4, 5 by other cores.

```
acquire(&listlock);
l = malloc(sizeof *1);
l->data = data;
l->next = list;
list = l;
release(&listlock);
```

- sync_synchronize tell the hardware and compiler not to perform re-orderings.
- sync_synchronize can guarantees execution order in critical section by being called in acquire and release.

```
acquire(struct spinlock *lk)
{
    pushcli(); // disable interrupts to avoid deadlock.
    if(holding(lk))
    panic("acquire");
    // The xchg is atomic.
    while(xchg(&lk->locked, 1) != 0)
    ;
    // Tell the C compiler and the processor to not move loads or stores
    // past this point, to ensure that the critical section's memory
    // references happen after the lock is acquired.
    ___sync_synchronize();
```



Sleep lock

- Sometimes a process needs to hold a lock for a long time.
- When a thread that wants to hold lock is looping, another thread that is holding lock cannot release lock.
- Sleep lock: yield CPU to others while waiting for the lock!



<pre>// Long-term locks for processes struct sleeplock { wint locked:</pre>				
struct spinlock lk: // spinlock protocting this sloop lock				
Struct Spiniock ik, // Spiniock protecting this steep lock				
// For debugging:				
char *name;	// Name of lock.			
int pid;	// Process holding lock			
};	For debuggin			
	i or debuggin			

- lock: put itself into sleep.
- ⁹ unlock: wake up one of the processes that have been waiting for the lock.

process structure for the sleep lock

specify the lock which it is waiting for. : called sleep channel.

// Per-process state	
<pre>struct proc {</pre>	
uint sz;	<pre>// Size of process memory (bytes)</pre>
<pre>pde_t* pgdir;</pre>	// Page table
char *kstack;	// Bottom of kernel stack for this process
enum procstate state;	// Process state
int pid;	// Process ID
<pre>struct proc *parent;</pre>	// Parent process
<pre>struct trapframe *tf;</pre>	// Trap frame for current syscall
<pre>struct context *context:</pre>	// swtch() here to run process
void *chan;	<pre>// If non-zero, sleeping on chan</pre>
int killed;	// If non-zero, have been killed
<pre>struct file *ofile[NOFILE];</pre>	// Open files
<pre>struct inode *cwd;</pre>	// Current directory
char name[16];	// Process name (debugging)
};	

sleep lock

o acquiresleep

going to sleep with holding a lock?

```
void
acquiresleep(struct sleeplock *lk)
{
    acquire(&lk->lk);
    while (lk->locked) {
        sleep(lk, &lk->lk);
      }
      lk->locked = 1;
      lk->pid = myproc()->pid;
      release(&lk->lk);
    }
}
```



void sleep(void *chan, struct spinlock *lk)

Place the caller to chan, release 1k and put the state of the caller process to sleep.



peek on sleep()



sleep unlock

- releasesleep
- Clears the sleeplock lk.
- Wakes up all sleeping processes by calling wakeup().

```
void
releasesleep(struct sleeplock *lk)
{
    acquire(&lk->lk);
    lk->locked = 0;
    lk->pid = 0;
    wakeup(lk);
    release(&lk->lk);
}
```

```
void
wakeup(void *chan)
{
    acquire(&ptable.lock);
    wakeup1(chan);
    release(&ptable.lock);
}
```

Lock the process table and call wakeup1().

After wakeup1 (), unlock the process table.



Better Sleep lock design



lock: If lock is not available, insert itself to the linked list and 'sleep'.
Mock: scan the linked list and wake one up.

lock ordering (mm/filemap.c)

```
Lock ordering:
  ->i mmap rwsem
                     (truncate pagecache)
    ->private lock
                         ( free pte-> set page dirty buffers)
      ->swap lock
                         (exclusive swap page, others)
        ->i pages lock
  ->i mutex
    ->i mmap rwsem
                         (truncate->unmap mapping range)
  ->mmap sem
    ->i mmap rwsem
      ->page table lock or pte lock (various, mainly in memory.c)
        ->i pages lock
                        (arch-dependent flush dcache mmap lock)
  ->mmap sem
    ->lock page
                     (access process vm)
*
  ->page table lock or pte lock
    ->swap lock
                     (try to unmap_one)
    ->private lock
                         (try to unmap one)
    ->i pages lock
                        (try to unmap one)
    ->zone lru lock(zone)
                             (follow page->mark page accessed)
*
    ->zone lru lock(zone)
                             (check pte range->isolate lru page)
                         (page remove rmap->set_page_dirty)
*
    ->private lock
    ->i pages lock
*
                         (page remove rmap->set page dirty)
*
    bdi.wb->list lock
                             (page remove rmap->set page dirty)
*
    ->inode->i lock
                         (page remove rmap->set page dirty)
                         (page remove rmap->lock_page_memcg)
    ->memcg->move lock
    bdi.wb->list lock
                             (zap_pte_range->set_page_dirty)
                         (zap pte range->set_page_dirty)
    ->inode->i lock
    ->private lock
                         (zap pte range-> set page dirty buffers)
```

Limitations of locks

- Function uses a data that needs to be protected by a lock.
 - The caller may hold a lock.
 - The caller may not hold a lock.
- Solution
 - Function with a lock and function without a lock: wakeup1 and wakeup
 - Ask caller to hold lock always.: sched
 - Recursive lock: the caller hold the lock and callee re-acquire the lock.

A process calls sched() to yield CPU.

Hand CPU over to other processes.

Put itself to the ready queue.

Cases

- A process uses up its time quantum.
- A process goes to sleep.
- A process exits.

peek on sched()

```
void
sched(void)
  int intena;
  struct proc *p = myproc();
 if(!holding(&ptable.lock))
    panic("sched ptable.lock");
  if(mycpu()->ncli != 1) {
    cprintf("panic! ncli: %d\n", mycpu()->ncli);
    panic("sched locks");
  if (p->state == RUNNING)
    panic("sched running");
  if(readeflags()&FL IF)
    panic("sched interruptible");
  intena = mycpu()->intena;
  swtch(&p->context, mycpu()->scheduler);
  mycpu() ->intena = intena;
```

Importance of system software and os

Toyota's killer firmware: Bad design and its consequences

Michael Dunn -October 28, 2013

129 Comments

- Toyota's electronic throttle control system (ETCS) source code is of unreasonable quality.
- Toyota's source code is defective and contains bugs, including bugs that can cause unintended acceleration (UA).
- Code-quality metrics predict presence of additional bugs.
- Toyota's fail safes are defective and inadequate (referring to them as a "house of cards" safety architecture).
- Misbehaviors of Toyota's ETCS are a cause of UA.

FUEL CELL CARS

Toyota calls back all the Mirais for software bug

It only affects about 2,800 vehicles around the world.

BY ANDREW KROK | FEBRUARY 16, 2017 11:19 AM PST

f y r 🗇 🎵





What is the programming language u jets?

Ad by JetBrains

Level up your code with IntelliJ IDEA.

An IDE built for professional development. Make developing enjoy

Free trial at jetbrains.com

6 Answers



Ajay Pandey, Equal Opportunity Skull Buster Answered Jun 23 2018 · Author has 87 answers and 304.8k answe

I am saddened reading this answer from a narrative builder what : experts in Quora. I can't resist to write a counter instead of an ans

Programming language doesn't really matter. Giving emphasis to : language is just a non-expert attempt to add credibility. I am not saying C++ is not powerful. I myself is a C++ programmer for 24 years and saw the evolution from C to C++98 up to all modern standards of C++11/14/17.

Domain driven and specially designed propriety languages starts taking over any complex domain. Specially when you can meet the program-ability requirement of hardware by porting/cross-compilation or automatic generation of code. I myself programmed in Matlab and designed systems in Simulink and auto generated C++ code to feed into Hardware; once simulated system succeeds.

The series of false claims is revealing that the poster appears don't have any experience in basic Embedded systems. Expecting aviation and defense specific expertise is way too much.

Here is small list from Ahmed Siddique's answer:

- 1. Most of the fighters jets use Ada. Not true, initial libraries were developed in Ada hence leading defense vendors carried it as far as possible. Most Nonwestern systems are not Ada unless they purchased Ada component libraries under ToT agreement.
- 2. Pakistan air force made a unique decision. Instead of using ADA programming language they used C++. - Not a unique decision, C++ is widely used language for interaction with Hardware even in defense applications.
- 3. Because there were literally tens of thousands of young highly



Using Linux in Air Traffic Control

Hardware and Operating System Platforms

Gerolf Ziegenhain

Janson 2017-02-05



London City Airport and NATS to introduce the UK's first digital air traffic control tower