


Lecture 29: Locking in xv6

Mythili Vutukuru

IIT Bombay

<https://www.cse.iitb.ac.in/~mythili/os/>

Why locking in xv6?



- No threads in xv6, so no two user programs can access same userspace memory image
 - No need for userspace locks like pthreads mutex
- However, scope for concurrency in xv6 kernel
 - Two processes in kernel mode in different CPU cores can access same kernel data structures
 - When a process is running in kernel mode, another trap occurs, and the trap handler can access data that was being accessed by previous kernel code
- Solution: spinlocks used to protect critical sections
 - Limit concurrent access to kernel data structures that can result in race conditions
- xv6 also has a sleeping lock (built on spinlock, not discussed)


P1 → kernel → intr

Spinlocks in xv6

- Acquiring lock: uses xchg x86 atomic instruction (test and set)
 - Atomically set lock variable to 1 and returns previous value
 - If previous value is 0, it means free lock has been acquired, success!
 - If previous value is 1, it means lock is held by someone, continue to spin in a busy while loop till success
- Releasing lock: set lock variable to 0
- Must disable interrupts on CPU core before spinning for lock
 - Interrupts disabled only on this CPU core to prevent another trap handler running and requesting same lock, leading to deadlock
 - OK for process on another core to spin for same lock, as the process on this core will release it
 - Disable interrupts before starting to spin (otherwise, vulnerable window after lock acquired and before interrupts disabled)

```
1500 // Mutual exclusion lock.
1501 struct spinlock {
1502     uint locked; // Is the lock held?
1503
1504     // For debugging:
1505     char *name; // Name of lock.
1506     struct cpu *cpu; // The cpu holding the lock.
1507     uint pcs[10]; // The call stack (an array of program counters;
1508                  // that locked the lock.
1509 };
```

```
1573 void
1574 acquire(struct spinlock *lk)
1575 {
1576     pushcli(); // disable interrupts to avoid deadlock.
1577     if(holding(lk))
1578         panic("acquire");
1579
1580     // The xchg is atomic.
1581     while(xchg(&lk->locked, 1) != 0)
1582         ;
1583
1584     // Tell the C compiler and the processor to not move loads or stores
1585     // past this point, to ensure that the critical section's memory
1586     // references happen after the lock is acquired.
1587     __sync_synchronize();
1588
1589     // Record info about lock acquisition for debugging.
1590     lk->cpu = mycpu();
1591     getcallerpcs(&lk, lk->pcs);
1592 }
```



Disabling interrupts

L1 L2

- Must disable interrupts on CPU core before beginning to spin for spinlock
- Interrupts stay disabled until lock is released
- What if multiple spinlocks are acquired?
 - Interrupts must stay disabled until all locks are released
- Disabling/enabling interrupts:
 - pushcli disables interrupts on first lock acquire, increments count for future locks
 - popcli decrements count, reenables interrupts only when all locks released and count is zero

```
1662 // Pushcli/popcli are like cli/sti except that they are matched:
1663 // it takes two popcli to undo two pushcli. Also, if interrupts
1664 // are off, then pushcli, popcli leaves them off.
1665
1666 void
1667 pushcli(void)
1668 {
1669     int eflags;
1670
1671     eflags = readeflags();
1672     cli();
1673     if(mycpu()->ncli == 0)
1674         mycpu()->intena = eflags & FL_IF;
1675     mycpu()->ncli += 1;
1676 }
1677
1678 void
1679 popcli(void)
1680 {
1681     if(readeflags() & FL_IF)
1682         panic("popcli - interruptible");
1683     if(--mycpu()->ncli < 0)
1684         panic("popcli");
1685     if(mycpu()->ncli == 0 && mycpu()->intena)
1686         sti();
1687 }
```

ptable.lock (1)

```
2409 struct {  
2410     struct spinlock lock;  
2411     struct proc proc[NPROC];  
2412 } ptable;
```

- The process table protected by a lock, any access to ptable must be done with ptable.lock held
- Normally, a process in kernel mode acquires ptable.lock, changes ptable, releases lock
 - Example: when allocproc allocates new struct proc
- But during context switch from process P1 to P2, ptable structure is being changed all through context switch, so when to release lock?
 - P1 acquires lock, switches to scheduler, switches to P2, P2 releases lock

ptable.lock (2)

- Every function that calls sched() to give up CPU will do so with ptable.lock held. Which functions invoke sched()?
 - Yield, when a process gives up CPU due to timer interrupt
 - Sleep, when process wishes to block
 - Exit, when process terminates
- Every function that swtch switches to will release ptable.lock. What functions does swtch return to?
 - Yield, when switching in a process that is resuming after yielding is done
 - Sleep, when switching in a process that is waking up after sleep
 - Forkret for newly created processes
- Purpose of forkret: to release ptable.lock after context switch, before returning from trap to userspace

```
2826 // Give up the CPU for one scheduling round.
2827 void
2828 yield(void)
2829 {
2830     acquire(&ptable.lock);
2831     myproc()->state = RUNNABLE;
2832     sched();
2833     release(&ptable.lock);
2834 }
```



```
2852 void
2853 forkret(void)
2854 {
2855     static int first = 1;
2856     // Still holding ptable.lock from scheduler.
2857     release(&ptable.lock);
2858
2859     if (first) {
2860         // Some initialization functions must be run in the context
2861         // of a regular process (e.g., they call sleep), and thus cannot
2862         // be run from main().
2863         first = 0;
2864         iinit(ROOTDEV);
2865         initlog(ROOTDEV);
2866     }
```

trapret

ptable.lock (3)

- Scheduler goes into loop with lock held
- Switch to P1, P1 switches back to scheduler with lock held, scheduler switches to P2, P2 releases lock
- Periodically, end of looping over all processes, releases lock temporarily
 - What if no runnable process found due to interrupts being disabled? Release lock, enable interrupts, allow processes to become runnable.

```
2757 void
2758 scheduler(void)
2759 {
2760     struct proc *p;
2761     struct cpu *c = mycpu();
2762     c->proc = 0;
2763
2764     for(;;){
2765         // Enable interrupts on this processor.
2766         sti();
2767
2768         // Loop over process table looking for process to run.
2769         acquire(&ptable.lock);
2770         for(p = ptable.proc; p < &ptable.proc[NPROC]; p++){
2771             if(p->state != RUNNABLE)
2772                 continue;
2773
2774             // Switch to chosen process. It is the process's job
2775             // to release ptable.lock and then reacquire it
2776             // before jumping back to us.
2777             c->proc = p;
2778             switchvm(p);
2779             p->state = RUNNING;
2780
2781             swtch(&(c->scheduler), p->context);
2782             switchkvm();
2783
2784             // Process is done running for now.
2785             // It should have changed its p->state before coming back.
2786             c->proc = 0;
2787         }
2788         release(&ptable.lock);
2789     }
2790 }
2791 }
```

```
graph TD
    S((S)) -- acquire --> P1((P1))
    P1 --> S
    S --> P2((P2))
    P2 --> S
    S -- release --> End(( ))
```

Summary

- Spinlocks in xv6 based on xchg atomic instruction
- Processes in kernel mode hold spinlock when accessing shared data structures, disabling interrupts on that core while lock is held
- Special ptable.lock held across context switch