

Locking in xv6

Mythili Vutukuru
CSE, IIT Bombay

Locking in xv6

- No threads in xv6, no two user programs can access same memory image
 - No need for userspace locks like pthreads mutex
- However, scope for concurrency in xv6 kernel
 - Two processes in kernel mode in different CPUs can access same kernel data structures like ptable
 - Even in single core, when a process is running in kernel mode, another trap occurs, 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)

Spinlocks in xv6

- Acquiring lock: uses xchg x86 atomic instruction (test and set)
 - Atomically set lock variable to new value 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

```
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
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 for kernel spinlocks (1)

- When acquiring kernel spinlock, **disables interrupts on CPU core**: why?
 - What if interrupt and handler requests same lock: **deadlock**
 - Interrupts disabled only on local core, OK to spin for lock on another core
 - Why disable interrupts before even acquiring lock? (otherwise, vulnerable window after lock acquired and before interrupts disabled)
- Disabling interrupts not needed for userspace locks like pthread mutex
 - Kernel interrupt handlers will not deadlock for userspace locks

Process in kernel mode

Kernel spinlock L acquired
Interrupt, switch to trap handler

Interrupt handler

Spin to acquire L
DEADLOCK

Process in kernel mode

Kernel spinlock L acquired

CRITICAL SECTION

Spinlock released

On another core

Spin to acquire L
Spin
Spin
Spinlock L acquired

Disabling interrupts for kernel spinlocks (2)

- Function `pushcli`: disables interrupts on CPU core before spinning for lock
 - 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, renables interrupts only when all locks released

```
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 }
```

Recap: Context switching in xv6 (1)

- Every CPU has a scheduler thread (special process that runs scheduler code)
- Scheduler goes over list of processes and switches to one of the runnable ones
- The special function “swtch” performs the actual context switch
 - Save context on kernel stack of old process
 - Restore context from kernel stack of new process

```
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 }
```

Recap: Context switching in xv6 (2)

- After running for some time, the process switches back to the scheduler thread, when:
 - Process has terminated (exit system call)
 - Process needs to sleep (e.g., blocking read system call)
 - Process yields after running for long (timer interrupt)
- Process calls “sched” which calls “swtch” to switch to scheduler thread again
- Scheduler thread runs its loop and picks next process to run, and the story repeats

```
2662 // Jump into the scheduler, never to return.
2663 curproc->state = ZOMBIE;
2664 sched();
2665 panic("zombie exit");
2666 }
```

```
2894 // Go to sleep.
2895 p->chan = chan;
2896 p->state = SLEEPING;
2897
2898 sched();
2899
```

```
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 }
```

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 in some way, 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



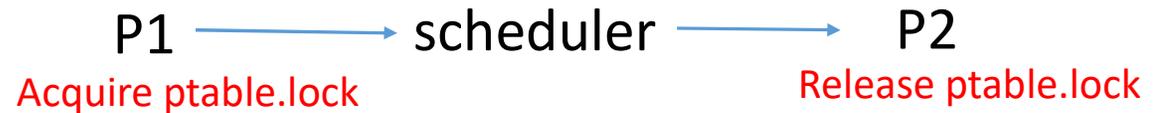
ptable.lock (2)

- Every function that calls sched() to give up CPU will do so with ptable.lock held
- Which functions invoke sched() to give up CPU?
 - Yield: process gives up CPU due to timer interrupt
 - Sleep: when process wishes to block
 - Exit: when process terminates
- Every function where a process resumes after being scheduled release ptable.lock
- What functions does a process resume after swtch?
 - Yield: resuming process after yield is done
 - Sleep: resuming process that is waking up after sleep
 - Forkret: for newly created processes
- Purpose of forkret: to release ptable.lock
 - New process then returns from trap like its parent

```
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 i
2861         // of a regular process (e.g., they call sleep
2862         // be run from main().
2863         first = 0;
2864         iinit(ROOTDEV);
2865         initlog(ROOTDEV);
2866     }
```

ptable.lock (3)



- Scheduler goes into loop with lock held
- Acquire ptable.lock in P1 \rightarrow scheduler picks P2 \rightarrow release in P2
- Later, acquire ptable.lock in P2 \rightarrow scheduler picks P3 \rightarrow release in P3
- 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 }
```