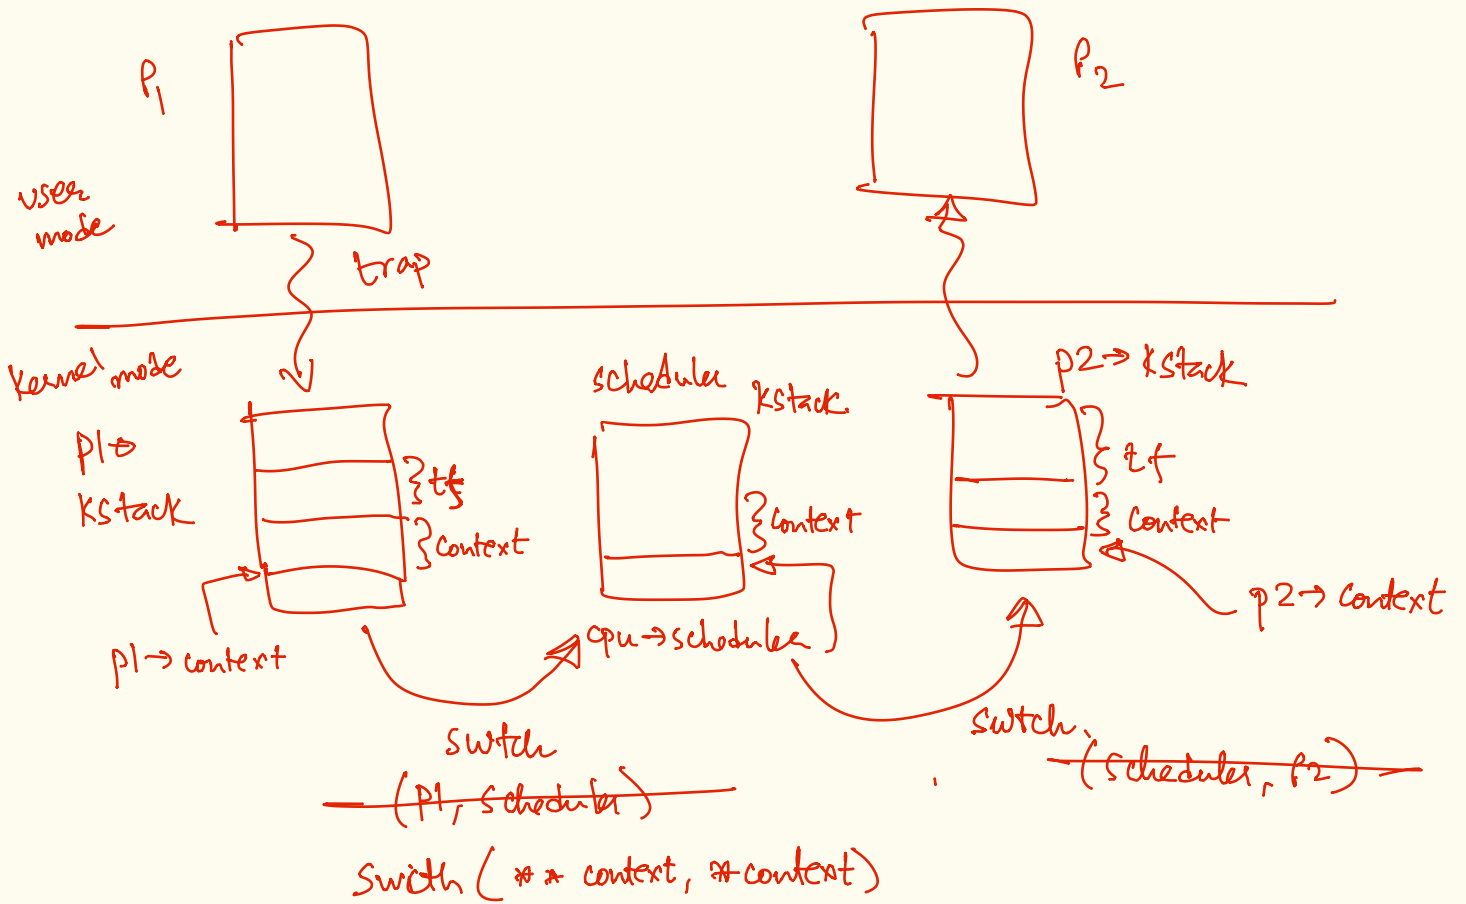


# # Lecture 16.



```

320 // - eventually that process transfers control
321 // via switch back to the scheduler.
322 void
323 scheduler(void)
324 {
325     struct proc *p;
326     struct cpu *c = mycpu();
327     c->proc = 0;
328
329     for(;;){
330         // Enable interrupts on this processor.
331         sti();
332
333         // Loop over process table looking for process to run.
334         acquire(&ptable.lock);
335         for(p = ptable.proc; p < &ptable.proc[NPROC]; p++){
336             if(p->state != RUNNABLE)
337                 continue;
338
339             // Switch to chosen process. It is the process's job
340             // to release ptable.lock and then reacquire it
341             // before jumping back to us.
342             c->proc = p;
343             switchvm(p);
344             p->state = RUNNING;
345
346             swtch(&(c->scheduler), p->context);
347             switchkvm();
348
349             // Process is done running for now.
350             // It should have changed its p->state before coming back.
351             c->proc = 0;
352         }
353         release(&ptable.lock);
354
355     }
356 }
357

```

switch (p, scheduler)

```

365 void
366 sched(void)
367 {
368     int intena;
369     struct proc *p = myproc();
370
371     if(!holding(&ptable.lock))
372         panic("sched ptable.lock");
373     if(mycpu()->ncli != 1)
374         panic("sched locks");
375     if(p->state == RUNNING)
376         panic("sched running");
377     if(readeflags() & FL_IF)
378         panic("sched interruptible");
379     intena = mycpu()->intena;
380     swtch(&p->context, mycpu()->scheduler);
381     mycpu()->intena = intena;
382 }
383
384 // Give up the CPU for one scheduling round.
385 void
386 yield(void)
387 {
388     acquire(&ptable.lock); //DOC: yieldlock
389     myproc()->state = RUNNABLE;
390     sched();
391     release(&ptable.lock);
392 }
393
394 // A fork child's very first scheduling by scheduler()
395 // will switch here. "Return" to user space.
396 void
397 forkret(void)
398 {
399     static int first = 1;
400     // Still holding ptable.lock from scheduler.
401     release(&ptable.lock);
402
403     if (first) {
404         // Some initialization functions must be run in the context
405         // of a regular process (e.g., they call sleep), and thus cannot
406         // be run from main().
407         first = 0;
408         iinit(ROOTDEV);
409         initlog(ROOTDEV);
410     }

```

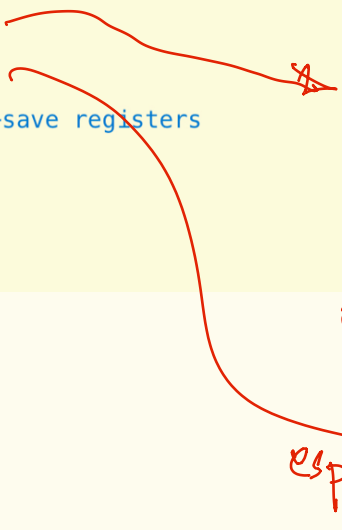

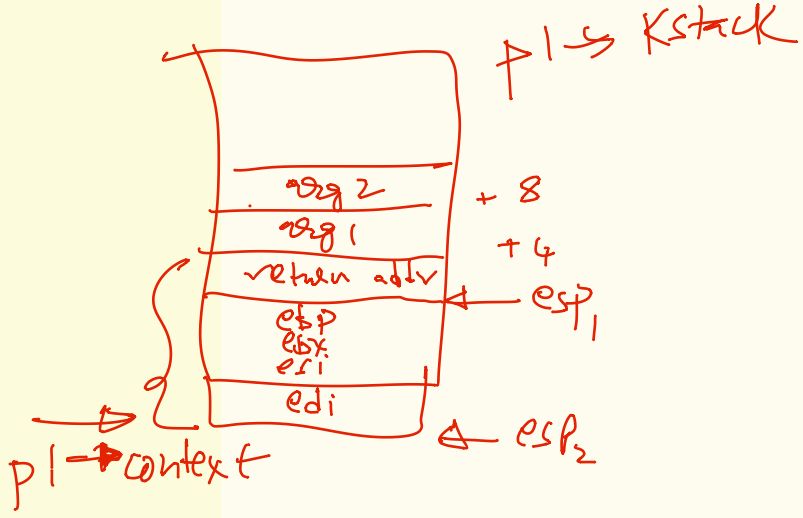

switch from process  
in kernel mode  
to  
the scheduler.

(any cpu) cpu on  
which process  
was executing  
in kernel  
mode.


pl → scheduler  
 switch (&pl → context,  
 cpu → scheduler)

```

1 # Context switch
2 #
3 # void swtch(struct context **old, struct context *new);
4 #
5 # Save the current registers on the stack, creating
6 # a struct context, and save its address in *old.
7 # Switch stacks to new and pop previously-saved registers.
8
9 .globl swtch
10 swtch:
11     movl 4(%esp), %eax
12     movl 8(%esp), %edx
13
14     # Save old callee-save registers
15     pushl %ebp
16     pushl %ebx
17     pushl %esi
18     pushl %edi
19
20     # Switch stacks
21     movl %esp, (%eax)
22     movl %edx, %esp
23
24     # Load new callee-save registers
25     popl %edi
26     popl %esi
27     popl %ebx
28     popl %ebp
29     ret
  
```



# # Concurrency / Synchronization.



```

C-statement
counter = counter + 1;

mov ebx, (%mem)
add ebx, 1
mov (%mem), ebx

read-update-write
↑      ↑
interrupt interrupt
    
```

Annotations: 'mov content from memory addr to ebx' points to the first instruction. 'non-deterministic sequence' is written next to the C-statement. 'race condition!' is written below the assembly code.

Threads of execution racing to ~~perform~~ read-update-write operations.

- outcome depends on when access is obtained.

~~non-~~ atomicity = fails!

all or nothing property!

↳ on what?

↳ critical section.