

10/14 Threads Execution

Threads

→ TLB

→ Within the same process, heap, code, data is shared, stack is private

→ Kernel threads

→ can execute in user mode, the main thread in every single threaded process

→ each process has a kstack (interrupt stack)

↳ each thread has its own kstack

```
int main() {  
    test();  
    foo();  
    bar();  
}
```

Single threaded

```
int main() {  
    thread_create (foo);  
    thread_create (bar);  
    thread_join () * 2;  
}
```

Multi-threaded

Thread Abstraction

→ an abstraction for dedicated CPU

Why might a thread suspend?

- ① timer interrupts. (fair sharing).
- ② I/O.
- ③ thread error out (tenants)
- ④ thread yield (give up voluntarily)

Programmer's View	Possible Execution #1	Possible Execution #2	Possible Execution #3
·	·	·	·
·	·	·	·
·	·	·	·
↓ x = x + 1; ·	→ x = x + 1;	→ x = x + 1;	x = x + 1;
↓ y = y + x; ·	→ y = y + x;	y = y + x;
z = x + 5y; ·	→ z = x + 5y;	Thread is suspended.
·	·	Other thread(s) run.	Thread is suspended.
·	·	Thread is resumed.	Other thread(s) run.
·	·	Thread is resumed.
·	·	→ y = y + x;
·	·	z = x + 5y;	z = x + 5y;

★ need to support pause & resume.

Context
Switch

```
209 // Give up the CPU for one scheduling round.
210 void yield(void) {
211     acquire(&table.lock); // DOC: yieldlock
212     myproc()->state = RUNNABLE; = READY.
213     sched();
214     release(&table.lock);
215 }
```

```
187 void sched(void) {
188     int intena;
189
190     if (!holding(&table.lock))
191         panic("sched ptable.lock");
192     if (mycpu()->ncli != 1) {
193         printf("pid : %d\n", myproc()->pid);
194         printf("ncli : %d\n", mycpu()->ncli);
195         printf("intena : %d\n", mycpu()->intena);
196
197         panic("sched locks");
198     }
199     if (myproc()->state == RUNNING)
200         panic("sched running");
201     if (readeflags() & FLAGS_IF)
202         panic("sched interruptible");
203
204     intena = mycpu()->intena;
205     swtch(&myproc()->context, mycpu()->scheduler);
206     mycpu()->intena = intena;
207 }
```

```
148 // - eventually that process transfers control
149 // via swtch back to the scheduler.
150 void scheduler(void) {
151     struct proc *p;
152
153     for (;;) {
154         // Enable interrupts on this processor.
155         sti();
156
157         // Loop over process table looking for process to run.
158         acquire(&ptable.lock);
159         for (p = ptable.proc; p < &ptable.proc[NPROC]; p++) {
160             if (p->state != RUNNABLE)
161                 continue;
162
163             // Switch to chosen process. It is the process's job
164             // to release ptable.lock and then reacquire it
165             // before jumping back to us.
166             mycpu()->proc = p;
167             vspaceinstall(p);
168             p->state = RUNNING;
169             swtch(&mycpu()->scheduler, p->context);
170             vspaceinstallkern();
171
172             // Process is done running for now.
173             // It should have changed its p->state before coming back.
174             mycpu()->proc = 0;
175         }
176         release(&ptable.lock);
177     }
178 }
```

process table

→ installs VAS.

→ running

→ switches to the next proc.

```
1 # Context switch code. We only need to save callee save registers
2 # as all other registers are saved before we reach here.
3 # arg0 -- context (in struct process) for the old process
4 # arg1 -- context (in struct process) for the new process
5 #
6 # note that xk only switches to/from the scheduler process
7 # the scheduler process picks the next thread to run
8 #
9 .globl swtch
10 swtch:
11     push %rbp
12     push %rbx
13     push %r11
14     push %r12
15     push %r13
16     push %r14
17     push %r15
18
19     mov %rsp, (%rdi)
20     mov %rsi, %rsp
21
22     pop %r15
23     pop %r14
24     pop %r13
25     pop %r12
26     pop %r11
27     pop %rbx
28     pop %rbp
29
30     ret
31
```

saving old process registers

old process state

sets the new process state

pushing new process's registers

```

1  #include <stdio.h>
2  #include <pthread.h>
3
4  int global_x = 0;
5
6  void* increment() {
7  → global_x += 1;
8     return NULL;
9  }
10
11 main
12 int main(int argc, char** argv) {
13     pthread_t tid1, tid2;
14     pthread_create(&tid1, NULL, increment, NULL);
15     pthread_create(&tid2, NULL, increment, NULL);
16
17     pthread_join(tid1, NULL);
18     pthread_join(tid2, NULL);
19
20     printf("%d\n", global_x); // minimum? maximum?
21
22     return 0;
23 }

```

tid1
tid2

(gdb) disas /m increment
Dump of assembler code for function increment:

```

6  void* increment() {
   0x0000000000401146 <+0>:  push  %rbp
   0x0000000000401147 <+1>:  mov   %rsp,%rbp
7  global_x += 1;
   0x000000000040114a <+4>:  mov   0x2ee8(%rip),%eax    # 0x404038 <global_x>
   0x0000000000401150 <+10>: add  $0x1,%eax
   0x0000000000401153 <+13>: mov  %eax,0x2edf(%rip)    # 0x404038 <global_x>
8  return NULL;
   0x0000000000401159 <+19>:  mov  $0x0,%eax
9  }
   0x000000000040115e <+24>:  pop  %rbp
   0x000000000040115f <+25>:  ret

```

Actually 3 instructions.

Possible outcomes

1) - 2 = t1 t2

⌋

execute one after the other.

2) - 1

t1
read global_x (0)

t2

x=1

increment
write global_x (1)

writes the register value

increment the register value

(into a register)

```
1 #include <stdio.h>
2 #include <pthread.h>
3
4 int global_x = 0;
5 pthread_t tids[100];
6
7 void* increment() {
8     global_x += 1;
9     return NULL;
10 }
11
12 int main(int argc, char** argv) {
13     for (int i=0; i<100; i++) {
14         pthread_create(&tids[i], NULL, increment, NULL);
15     }
16     printf("%d\n", global_x); // minimum? maximum?
17
18     for (int i=0; i<100; i++) {
19         pthread_join(tids[i], NULL);
20     }
21     printf("%d\n", global_x); // minimum? maximum?
22
23     return 0;
24 }
```

read x to reg
add 1 to reg
write reg to x

0 = why?

100: all execute in turn

t_1
read $x=0$



$t_2 - t_{100}$
all execute in order

add 1 to reg
write reg to x
(1)

100

```
1  #include <stdio.h>
2  #include <pthread.h>
3
4  int global_x = 0;
5  pthread_t tids[2];
6
7  void* increment100() {
8      for (int i=0; i<100; i++) {
9          global_x += 1;
10     }
11     return NULL;
12 }
13
14 int main(int argc, char** argv) {
15     for (int i=0; i<2; i++) {
16         pthread_create(&tids[i], NULL, increment100, NULL);
17     }
18
19     for (int i=0; i<2; i++) {
20         pthread_join(tids[i], NULL);
21     }
22
23     printf("%d\n", global_x); // minimum? maximum?
24
25     return 0;
26 }
```

read x
add 1
write x



Smaller than 100, so what is it?

200

Problem: Unpredictable Output Based on Different Executions

- multiple threads access a shared variable w/out protections (data race)
- explicit synchronization primitives = locks