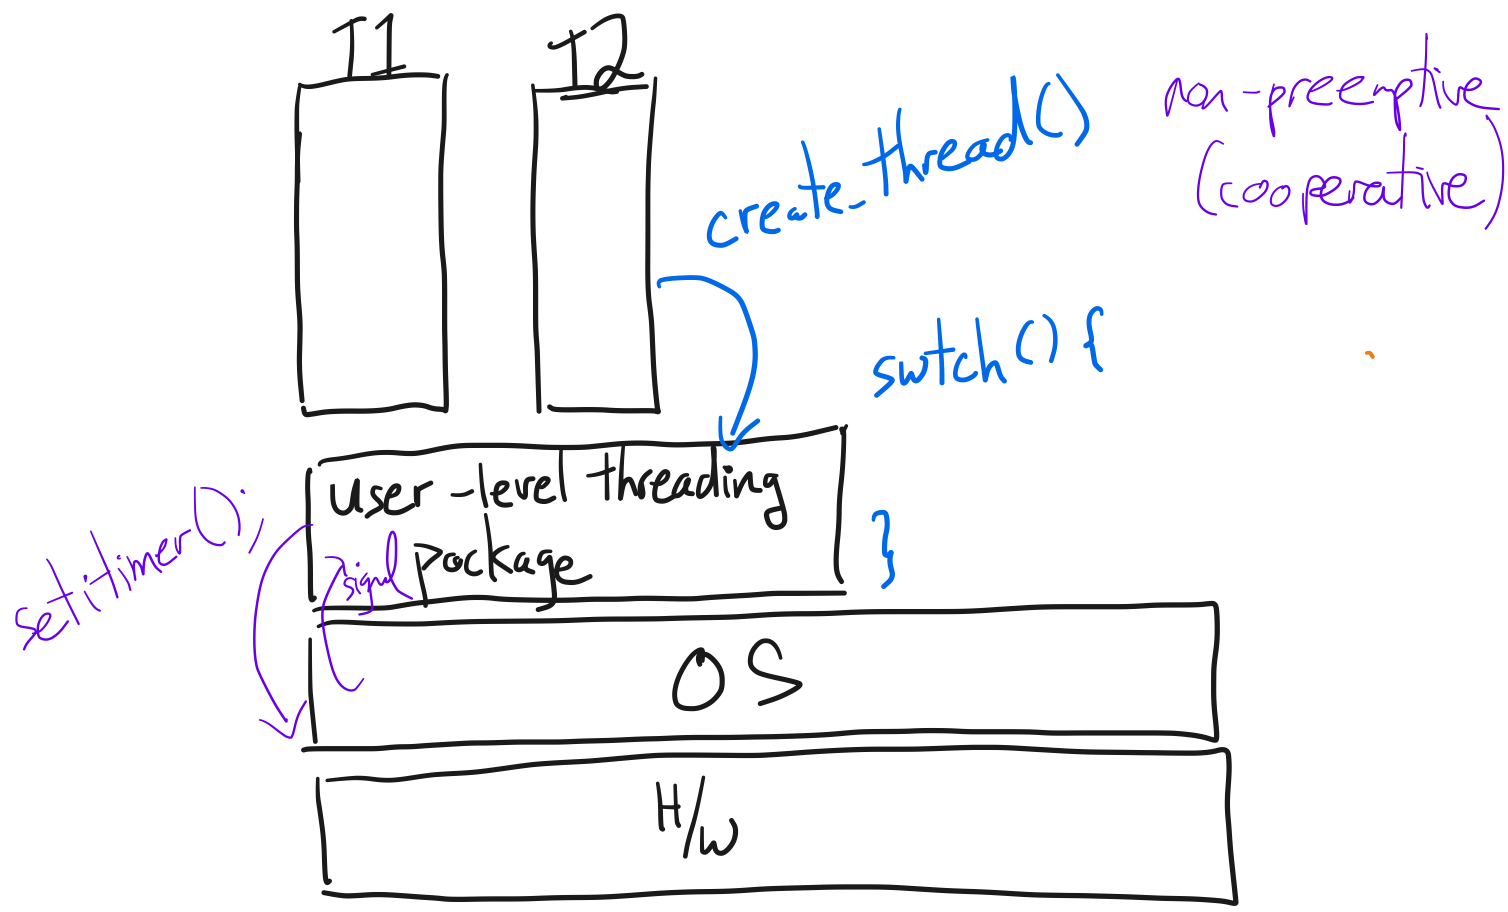


- 1. Last time
 - 2. Context switches (WeensyOS)
 - 3. User-level threading, intro
 - 4. Context switches (user-level threading)
 - switch()
 - yield()
 - I/O
 - 5. Cooperative multithreading
 - 6. Preemptive user-level multithreading
-

2. Context switches in WeensyOS


3. User-level threading

preemptive



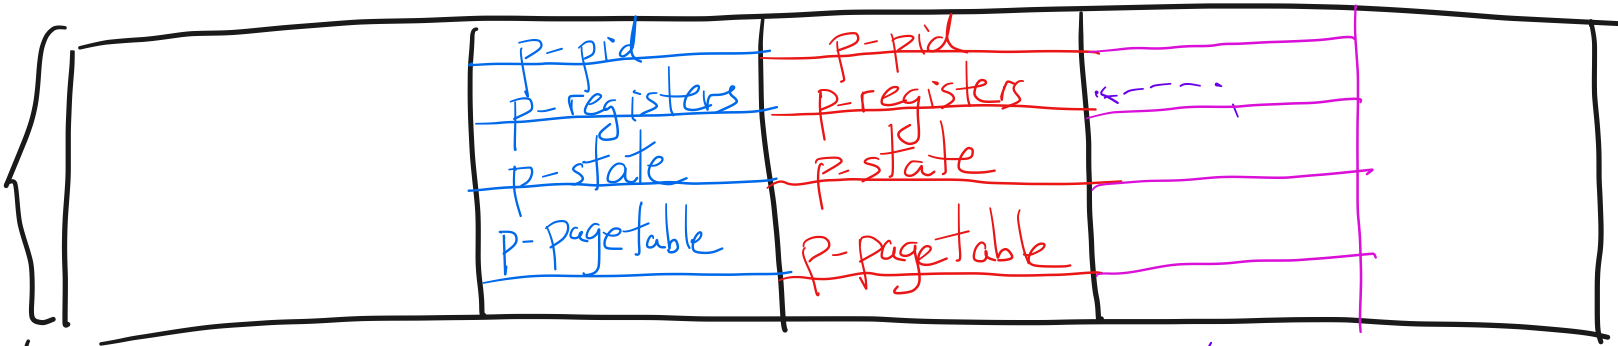
4. Context switches (user space)

- switch registers active
- switch page tables



over
text

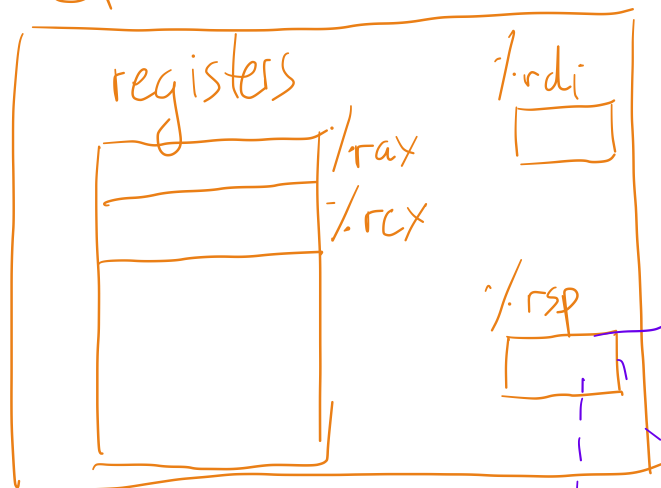
Context switches in Weensy OS



processes

Mem


CPU





Processes

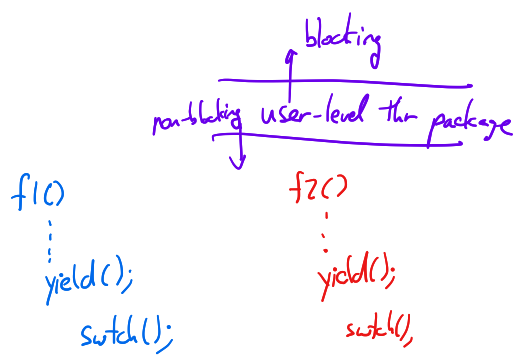
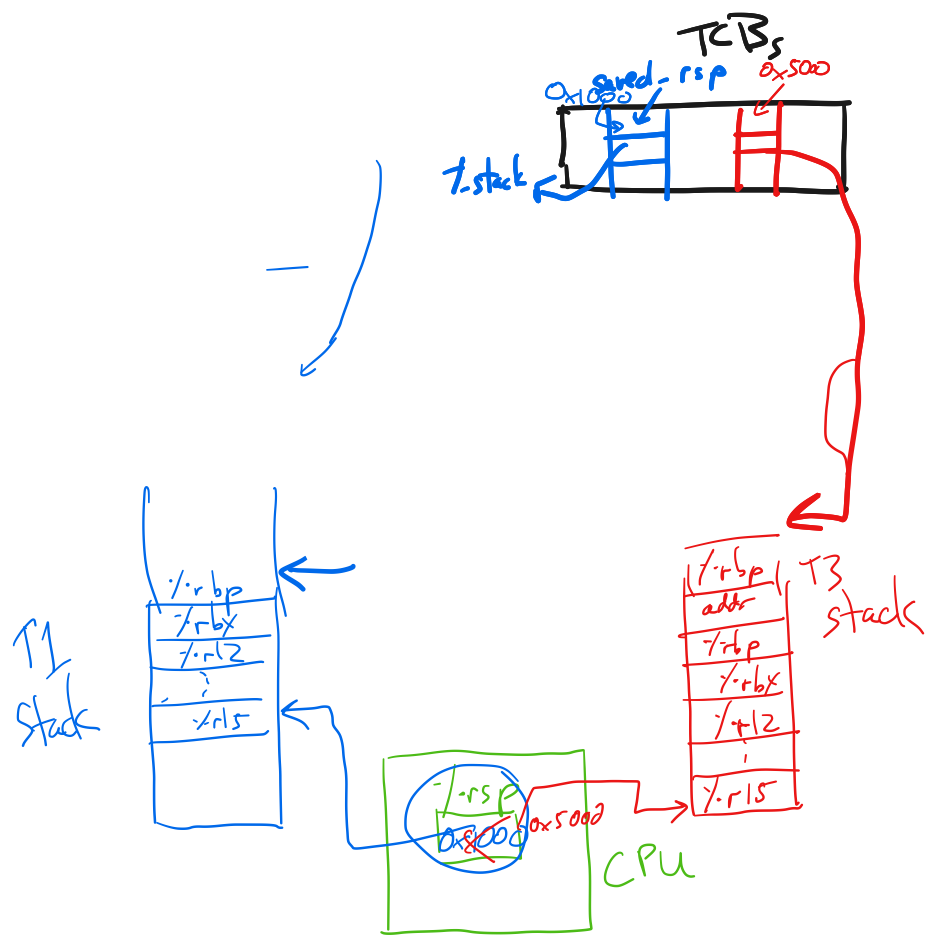
p-registers

save %rsp, %rdi



trapframe





1 CS 202, Spring 2022
2 Handout 10 (Class 17)

3
4 1. User-level threads and switch()

5
6 We'll study this in the context of user-level threads.

7 Per-thread state in thread control block:

```
8
9
10 typedef struct tcb {
11     unsigned long saved_rsp; /* Stack pointer of thread */
12     char *t_stack; /* Bottom of thread's stack */
13     /* ... */
14 };
```


15 Machine-dependent thread initialization function:

```
16 void thread_init(tcb **t, void (*fn) (void *), void *arg);
```

TCBs

17 Machine-dependent thread-switch function:



```
18 void switch(tcb *current, tcb *next);
```



19 Implementation of switch(current, next):

```
20
21 # gcc x86-64 calling convention:
22 # on entering switch():
23 # register %rdi holds first argument to the function ("current")
24 # register %rsi holds second argument to the function ("next")
25
26 # Save call-preserved (aka "callee-saved") regs of 'current'
27 pushq %rbp
28 pushq %rbx
29 pushq %r12
30 pushq %r13
31 pushq %r14
32 pushq %r15
33
34 # store old stack pointer, for when we switch() back to "current" later
35 movq %rsp, (%rdi) /* %rdi->saved_rsp = %rsp */
36 movq (%rsi), %rsp /* %rsp = %rsi->saved_rsp */
37
38 # Restore call-preserved (aka "callee-saved") regs of 'next'
39 popq %r15
40 popq %r14
41 popq %r13
42 popq %r12
43 popq %rbx
44 popq %rbp
45
46 # Resume execution, from where "next" was when it last entered switch()
47 ret
```

current -> saved_rsp = stack_ptr



0xf.f2000

0xf.f5000

55
56 2. Example use of switch(): the yield() call.

57
58 A thread is going about its business and decides that it's executed for
59 long enough. So it calls yield(). Conceptually, the overall system needs
60 to now choose another thread, and run it:


```
61 void yield() {
62
63     tcb* next = pick_next_thread(); /* get a runnable thread */
64     tcb* current = get_current_thread();
65
66     switch(current, next);
67
68     /* when 'current' is later rescheduled, it starts from here */
69 }
70
```

71 3. How do context switches interact with I/O calls?

72 This assumes a user-level threading package.

73
74 The thread calls something like "fake_blocking_read()". This looks
75 to the _thread_as though the call blocks, but in reality, the call
76 is not blocking:

```
77
78 int fake_blocking_read(int fd, char* buf, int num) {
79
80     int nread = -1;
81
82     while (nread == -1) {
83
84         /* this is a non-blocking read() syscall */
85         nread = read(fd, buf, num);
86
87         if (nread == -1 && errno == EAGAIN) {
88             /*
89              * read would block. so let another thread run
90              * and try again later (next time through the
91              * loop).
92              */
93             yield();
94         }
95     }
96
97     return nread;
98 }
99
```



f1() ... yield(); switch();

f2() ... yield(); switch();