

- 1. Last time
- 2. Context switches (WeensyOS)
- 3. Synchronous vs. asynchronous I/O
- 4. User-level threading, intro.
- 5. Context switches (user-level threading)

switch()
yield()
I/O


sleep() implies that there was a context switches between the running thread yields the processor.

- 6. Cooperative multithreading
- 7. Preemptive user-level multithreading

2. (Context switches) in WeensyOS

changing %cr3
changing registers





kernel itself NEVER blocks when doing I/O

3.
blocking ==
synchronous

```
read (fd, buf, sz);
write (fd, buf, sz);
```

user process
view

non-blocking ==
asynchronous

```
read ( , , );
write ( , , );
```


"would
block"

fd repr a file

blocking

4. User-level threading

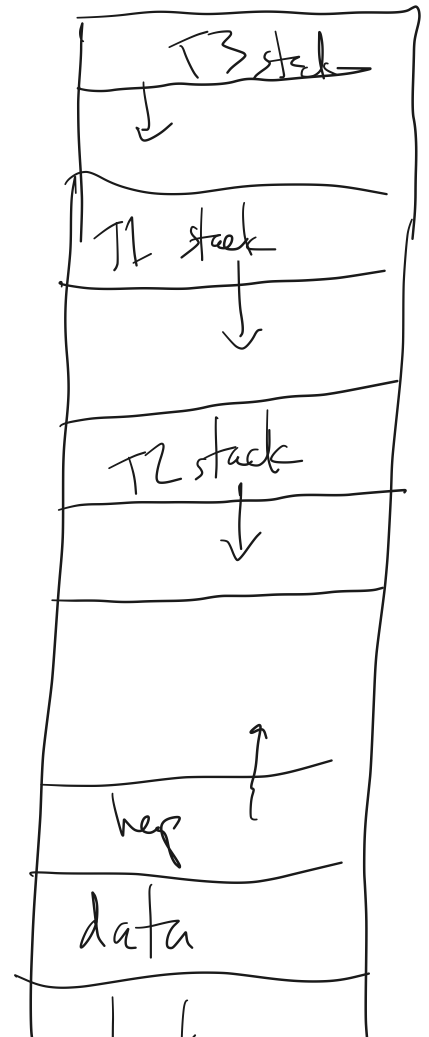
TCPs



5. Context switches (user space)

only change registers

switch ()



text

1 CS 202, Spring 2021
2 Handout 11 (Class 17)

1. User-level threads and switch()

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

Per-thread state in thread control block:

```

typedef struct tcb {
    unsigned long saved_rsp; /* Stack pointer of thread */
    char *t_stack;          /* Bottom of thread's stack */
    /* ... */
};

```

Machine-dependent thread initialization function:

```

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

```

Machine-dependent thread-switch function:

```

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

```

Implementation of switch(current, next):

```

# gcc x86-64 calling convention:
# on entering switch():
# register %rdi holds first argument to the function ("current")
# register %rsi holds second argument to the function ("next")



# Save call-preserved (aka "callee-saved") regs of 'current'
pushq %rbp
pushq %rbx
pushq %r12
pushq %r13
pushq %r14
pushq %r15

# store old stack pointer, for when we switch() back to "current" later
movq %rsp, (%rdi) /* %rdi->saved_rsp = %rsp */
movq (%rsi), %rsp /* %rsp = %rsi->saved_rsp */

# Restore call-preserved (aka "callee-saved") regs of 'next'
popq %r15
popq %r14
popq %r13
popq %r12
popq %rbx
popq %rbp

# Resume execution, from where "next" was when it last entered switch()
ret

```



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

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

```

void yield() {
    tcb* next = pick_next_thread(); /* get a runnable thread */
    tcb* current = get_current_thread();
    switch(current, next);
    /* when 'current' is later rescheduled, it starts from here */
}

```

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

This assumes a user-level threading package.

The thread calls something like "fake_blocking_read()". This looks to the _thread_ as though the call blocks, but in reality, the call is not blocking:

```

int fake_blocking_read(int fd, char* buf, int num) {
    int nread = -1;
    while (nread == -1) {
        /* this is a non-blocking read() syscall */
        nread = read(fd, buf, num);
        if (nread == -1 && errno == EAGAIN) {
            /*
             * read would block. so let another thread run
             * and try again later (next time through the
             * loop).
             */
            yield();
        }
    }
    return nread;
}

```

