



Department of Electrical Engineering and Computer Science

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

6.S081/6.828 Fall 2019

Quiz II

All problems are open-ended questions. In order to receive credit you must answer the question as precisely as possible. You have 120 minutes to finish this quiz.

Write your name on this cover sheet AND at the bottom of each page of this booklet.

Some questions may be harder than others. Read them all through first and attack them in the order that allows you to make the most progress. If you find a question ambiguous, be sure to write down any assumptions you make. Be neat. If we can't understand your answer, we can't give you credit!

THIS EXAM IS OPEN BOOK AND OPEN LAPTOP, but CLOSED NETWORK.

Please do not write in the boxes below.

I (xx/15)	II (xx/15)	III (xx/5)	IV (xx/10)	V (xx/4)	VI (xx/15)	VII (xx/5)	VIII (xx/5)	IX (xx/2)	(xx/76)

Kerberos ID:

I Xv6 file system

Alyssa adds the statement:

```
printf("bwrite %d\n", b->blockno);
```

to xv6's `bwrite` in `bio.c`. She then makes a fresh `fs.img`, boots xv6, and runs the following command:

```
$ cat README > z
bwrite 3
bwrite 4
bwrite 5
bwrite 2
bwrite 33
bwrite 46
bwrite 32
bwrite 2
bwrite 3
bwrite 4
bwrite 5
bwrite 2
bwrite 45
bwrite 801
bwrite 33
bwrite 2
bwrite 3
bwrite 4
bwrite 2
bwrite 801
bwrite 33
bwrite 2
bwrite 3
bwrite 4
bwrite 5
bwrite 2
bwrite 45
bwrite 802
bwrite 33
bwrite 2
bwrite 3
bwrite 4
bwrite 2
bwrite 802
bwrite 33
bwrite 2 // XXX
$
```

Name:

II Lab Lock

Ben is working on parallelizing xv6's memory allocator for 6.828's lock lab. He modifies the kernel's page allocator to use per-CPU free lists, using the `cpuid()` function to determine which CPU the code is running on. Following the hint in the lab text, Ben sees that `cpuid()`'s documentation (in `proc.c`) says "Must be called with interrupts disabled, to prevent a race with the process being moved to a different CPU." He writes his `kalloc()` implementation as follows:

```
1 void *
2 kalloc(void)
3 {
4     struct run *r;
5
6     push_off();
7     int me = cpuid();
8     pop_off();
9     acquire(&kmem[me].lock);
10    r = kmem[me].freelist;
11    if(r)
12        kmem[me].freelist = r->next;
13    release(&kmem[me].lock);
14
15    // if r == NULL, search other CPUs' free lists
16    // ...
17
18    // clear page and return it
19    // ...
20 }
```

4. [5 points]: Ben expects line 10, `r = kmem[me].freelist;`, to access the free list of the CPU that the code is currently executing on. It turns out that this is **not** always the case. Describe a concrete sequence of events that violate Ben's expectation.

Name:

5. [5 points]: How could Ben change his code so that his expectation holds, i.e. `r = kmem[me].freelist;` is guaranteed to access the free list of the CPU that the code is executing on?

6. [5 points]: Alyssa points out that Ben can remove the calls to `push_off()` (line 6) and `pop_off()` (line 8), even though that violates the `cpuid()` function's specification. Ben modifies xv6 and it passes all the usertests, despite the missing `push_off()` and `pop_off()`. Explain why deleting these two lines doesn't break xv6.

III Lab Syscall

The two parts of the 6.828 user-level threads and alarm lab both involve saving and restoring contexts. For user-level threading, this happens in `uthread_switch`, and for the alarm system call, saving the context happens in `usertrap()` and restoring the context happens in `sys_sigreturn()`. While all registers are saved/restored for the alarm system call, this is not necessary for `uthread_switch`, which only needs to save `sp`, `s0` through `s11`, and `ra`.

7. [5 points]: Explain why `uthread_switch` can get away with not saving and restoring certain general-purpose registers.

IV xv6 i-node counts

Alyssa is implementing symbolic links in xv6, as part of the 6.828 file system lab. She observes that each in-memory inode (`struct inode` in `file.h`) contains two similar fields: `nlink` and `ref`.

8. [5 points]: Suppose a bug accidentally changed a file's `nlink` field from 2 to 1. What bad thing(s) would happen as a result?

9. [5 points]: Suppose a bug accidentally changed a file's `ref` field from 2 to 1. What bad thing(s) would happen as a result?

Name:

V EXT3

Recall the Linux EXT3 journaling file system from *Journaling the Linux ext2fs Filesystem* and Lecture 14. The paper's "ext2fs" is the same as EXT3.

Suppose you run the following program on Linux with EXT3:

```
int
main()
{
    int fd;

    fd = open("a", O_CREAT|O_WRONLY, 0666); // create a
    if(fd < 0) exit(1);
    close(fd);

    fd = open("b", O_CREAT|O_WRONLY, 0666); // create b
    if(fd < 0) exit(1);
    close(fd);

    printf("done\n");
}
```

The two `open()` system calls create files. Before you run the program, the two files did not exist.

The program prints `done`, so you know the file creations succeeded. Moments after the program finishes, there's a power failure and your computer (which has no battery) stops executing. After a while the power comes back on, and your computer reboots and runs EXT3's recovery code. You look for files `a` and `b`.

10. [4 points]: Which of the following situations could exist at this point?
(Circle True or False for each choice.)

- **True / False** Neither `a` nor `b` exists.
- **True / False** File `a` exists, but not `b`.
- **True / False** File `b` exists, but not `a`.
- **True / False** Both `a` and `b` exist.

Name:

VI RCU

Consider *RCU Usage In the Linux Kernel: One Decade Later*, by McKenney *et al.*

Suppose Figure 6's `setsockopt()` were modified to copy the new options into the socket structure, rather than changing a pointer, like this:

```
if (opt == IP_OPTIONS) {
    memmove(sock->opts, arg, ...the correct size...);
    return;
}
```

This modification would require that `sock->opts` be a buffer of the appropriate size. `udp_sendmsg()` remains the same.

- 11. [5 points]:** Explain why this modification would break RCU. What kinds of problems would it cause to `udp_sendmsg()`?

Ben Bitdiddle is thinking about adding RCU to his xv6 kernel, which he runs on real RISC-V hardware. His RISC-V hardware has a timer that ticks every 10 milliseconds, and Ben sees that xv6's `kerneltrap()` causes a context switch if a timer interrupt arrives while a kernel thread is running. He reasons that since the point of `synchronize_rcu()` is to wait for a context switch on each CPU, he could change `synchronize_rcu()` to simply wait 10 milliseconds rather than schedule itself on each CPU in turn (as in Figure 2).

12. [5 points]: Explain why Ben's idea of waiting 10 milliseconds would break RCU.

Alyssa is excited about Biscuit as described in *The benefits and costs of writing a POSIX kernel in a high-level language* by Cutler et al. She notices that Biscuit uses RCU for its directory cache. In studying the Biscuit code, she notices that there are no calls to `synchronize_rcu()` (or to any similar function).

13. [5 points]: Explain why Biscuit has no need for `synchronize_rcu()`.

VII Networking lecture/reading

Consider *Eliminating Receive Livelock in an Interrupt-driven Kernel*, by Mogul *et al*, as well as Lecture 17.

Ben notices that a networking stack is failing to make much progress sending packets when it is receiving packets at a high rate. To solve this problem, he configures the network interface card (NIC) to **not** generate receive interrupts, and instead polls the receive descriptor ring during each timer tick (every millisecond). More progress is now made in sending packets, but a new problem has appeared where some incoming packets are dropped even at load below saturation.

14. [5 points]: Explain why packets are dropped because of this change.

VIII Virtualization

Consider *A Comparison of Software and Hardware Techniques for x86 Virtualization*, by Adams *et al.*

The designers of a new RISC-V processor want to more easily debug the state of the running CPU, so they allow reads to the `sstatus` (supervisor status) register to be performed from user mode without trapping. Normally making `sstatus` readable doesn't create any problems, but suppose a trap-and-emulate hypervisor (in supervisor mode) is running `xv6` as a guest (in user mode). After running for a short time the guest `xv6` kernel panics and prints "kerneltrap: not from supervisor mode".

Note: `sret` does not modify `SSTATUS.SPP`.

- 15. [5 points]:** Explain why allowing `sstatus` to be read in user mode without trapping results in this specific panic in the `xv6` guest kernel.

IX 6.828

16. [1 points]: Please indicate which of the labs you found to be the most helpful, and which the least.

- Alarm / uthreads

- Locking

- File system

- mmap

- Networking (if you did this lab)

17. [1 points]: Which paper is the best candidate for deletion in future years?

End of Quiz II — Happy holidays!

Name: