

CS 167 Midterm Exam Solutions

Spring 2016

Do all of questions 1 through 4.

1. *In the mthreads assignment you implemented user threads on top of POSIX threads (referred to as LWPs — lightweight processes). This is known as a two-level threads implementation. One detail is the number of LWPs that should be created. POSIX threads allows for two-level implementations (though the Linux implementation is not done this way). It provides the routine `pthread_setconcurrency` to provide “advice” to the pthreads library on how many LWPs should be created. How this advice is used is not specified, but it has been used as a specification of a lower bound on the number of LWPs that should be created. (Note that “`pthread_setconcurrency`” is mentioned only as background information. You don’t need to refer to it in your answers.)*

a. *Why is setting such a lower bound important on multicore systems?*

Since the OS schedules LWPs onto cores, there should be at least as many LWPs as cores so that the program can take advantage of the parallelism provided by the multiple cores.

b. *Could there be problems, even on single-core systems, if there are not enough LWPs? Either give a specific example of such a problem or provide a careful argument as to why there wouldn’t be problems. What we are looking for is a major problem that would go away with the addition of an LWP. Assume the system has at least one LWP.*

Yes, there could be a problem. Consider a situation in which there is one LWP, but two user threads. One thread, currently assigned to the only LWP, is blocked, attempting to read from a pipe. The other thread, when it next gets to run, will write to the pipe, thus unblocking the first thread. However, since there’s only one LWP, and it is blocked, the second thread will never get a chance to run and thus will never unblock the first thread. However, if one additional LWP were available, the second thread would get to run.

2. *You are to write a device driver for a data-collection device. Each time a new piece of data is gathered, the device interrupts the processor. The device driver must respond to the interrupt, and, in the interrupt handler, fetch the data from a device register and put the data into a queue from which the data will be consumed by user threads (that have entered the kernel to perform read system calls). Each user thread consumes one data item at a time. There may be multiple data-collection devices attached to the system; their interrupt handlers put their data on the same shared queue. All interrupt the processor at the same priority level (thus one interrupt is processed at a time). Assume that there are no special instructions for putting items into a queue atomically or for removing items from a queue atomically. Keep in mind that when an interrupt at a particular priority takes place, the processor automatically masks interrupts at that priority and lower.*

a. *Assume we have a uniprocessor system and the operating-system kernel is non-preemptible. Describe how you would synchronize access to the queue.*

Threads should mask interrupts while accessing the queue.

b. *Suppose now the kernel is preemptible. How would you change your answer to part a?*

There is now the additional concern that multiple threads might access the queue concurrently. To prevent this, a mutex should be associated with the queue and threads must lock the mutex before accessing the queue and unlock it afterwards.

c. *Suppose further that we have a multiprocessor system and that each device interrupt is delivered to a processor chosen at random (thus two devices can interrupt different processors simultaneously). How would you change your answer to part b? Note that a thread can mask interrupts only on the processor on which it is running.*

It’s no longer the case that masking interrupts is sufficient to keep an interrupt handler from accessing the queue. Since interrupt handlers cannot sleep waiting for a mutex, we use a spin lock. Thus all entities, both threads and interrupt handlers, must lock the spin lock before accessing the queue and unlock it afterwards. It’s now a problem if a thread, holding the spin lock, loses its time slice. So, to prevent this

from happening, threads disable preemption (perhaps by masking clock interrupts) while accessing the queue.

- d. Finally, suppose that each piece of data requires significant processing that must be completed before the data is made available to user threads. What additional concerns do we have? How should they be coped with?

Our concern is that we don't want the processing of data to interfere with other important interrupt-handling duties of the processor. One way of avoiding the problem is to defer the processing by use of a deferred procedure call (DPC) so that the processing is done at the lowest interrupt priority level, but before the system returns to the context of the interrupted thread.

3. In systems using x86 architectures, devices are controlled via registers that are mapped into the address space. So, for example, a disk device might have a memory-address register into which the operating system places the address of a buffer, a device-address register into which the operating system places the disk address of an operation, and a control register into which the operating system places the device command to be performed. All three registers are accessed by loading from and storing to what appear to be memory locations.

How would a virtual machine monitor (VMM) virtualize device I/O in such architectures? Assume that pure virtualization (and not paravirtualization) is being used, and thus unmodified operating systems that normally run on bare hardware are to run in virtual machines. (Hint: the address space is divided into fixed-size pieces called pages, and each page can be separately protected.)

The pages containing the device registers should be set up so that all accesses by the operating system running on the virtual machine result in faults. This way the VMM can respond to the faults, mediating then performing the operations requested.

4. Assume we're using the Rhinopias disk drive.

- a. Suppose we are using S5FS and we've doubled the block size from 512 bytes to 1024 bytes. Is the maximum expected transfer rate roughly doubled? Explain.

Yes. Since the blocks of a file are randomly scattered, fetching each requires an average seek time plus an average rotational delay. The time for the actual transfer of a block is negligible in comparison. By doubling the block size, after each seek and rotational delay, twice as much data is transferred as before.

- b. We've done the same with FFS. Is the maximum expected transfer rate roughly doubled? Explain. Assume that we're using two-way block interleaving: when we allocate successive blocks on a track, one block is skipped. Thus, in the best case, every other block of a track is allocated to a file.

In the best case, in which the blocks of a file are all within a cylinder group and arranged on tracks as efficiently as is permitted, each rotation of the disk causes half a track's worth of data to be transferred. This is unaffected by the block size.

- c. Why should we care about the block size in FFS?

Consider the case in which FFS is being used on a busy server, with concurrent requests for blocks from many files. The situation could be as bad as things are with S5FS, in which successive requests are to random disk locations. Thus the analysis given in part a for the speedup in S5FS applies.

If you do all of the following correctly, you'll get an A regardless of how well you do on the first four problems. If you miss any of the following, your grade will be based solely on how well you do on the first four problems.

5. Roy Tomlinson, who is credited with the invention of email and, in so doing, with the use of the @ sign for email addresses, recently passed away. He was a contractor working on the development of a computer network. Which network was this?

The ARPAnet.

6. Early Unix systems, which were interconnected via telephone lines, used a different syntax to specify an email address. What was this syntax?

MachineName!user (e.g., brunix!twd). This implied one's computer had the phone number of the destination computer. If not, addresses could be composed. For example, if your computer had the phone number of decvax, which had the phone number of brunix, you could email twd as decvax!brunix!twd.

7. *Dynamic linking was introduced into both Unix and Windows in the early 1990s. On which operating system, and roughly when, was the concept first developed?*

The MULTICS system in the mid 1960s.

8. *Gene Amdahl, who passed away recently, was the founder of the Amdahl Corporation, a maker of mainframe computers that competed with IBM. He was prominent in the computer world before forming his company. What was he most famous for before then?*

He was chief architect of the IBM/360, whose successors were the primary competition of his company.

9. *Douglas Engelbart passed away in 2013. He was a computer visionary who is most famous for what invention?*

The computer mouse.

10. *What female computer scientist (still living) is credited with the development of the first operating system for a personal computer?*

Mary Allen Wilkes. She developed LAP (Linc Assembly Program) for the LINC computer while at Lincoln Laboratory in the early 1960s.