

# CS 167 Final Exam

2pm May 14, 2016

Closed Book

Please Use a Pen, Not a Pencil!

**Do all questions.**

1. The implementation given of *thread\_switch* in class is as follows:

```
void thread_switch() {
    thread_t NextThread, OldCurrent;

    NextThread = dequeue(RunQueue);
    OldCurrent = CurrentThread;
    CurrentThread = NextThread;
    swapcontext(&OldCurrent->context, &NextThread->context);

    // We're now in the new thread's context
}
```

- a. Suppose we don't have a *swapcontext* routine to call. Instead, we have to use *getcontext* and *setcontext*. Recall that *getcontext* saves all of the calling thread's registers (including the stack pointer, but not the instruction pointer) into its argument, a pointer to a data structure of type *ucontext\_t*, and returns nothing. The routine *setcontext* restores the registers from its first (and only) argument (of type *ucontext\_t*). Show how to implement *thread\_switch* using *getcontext* and *setcontext*. If changes to the thread control block (of type *thread\_t*) are required, indicate what they are.

- b. Suppose *thread\_switch* takes an argument and this argument is referred to after *setcontext* returns. Explain what, if anything, must be done to ensure that what is referred to is what was passed in the most recent call to *thread\_switch*.

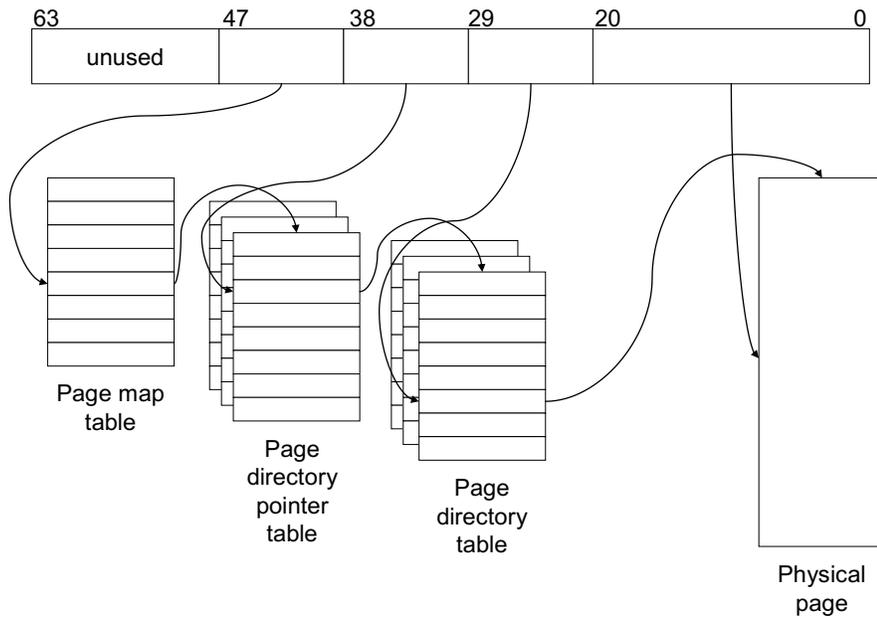
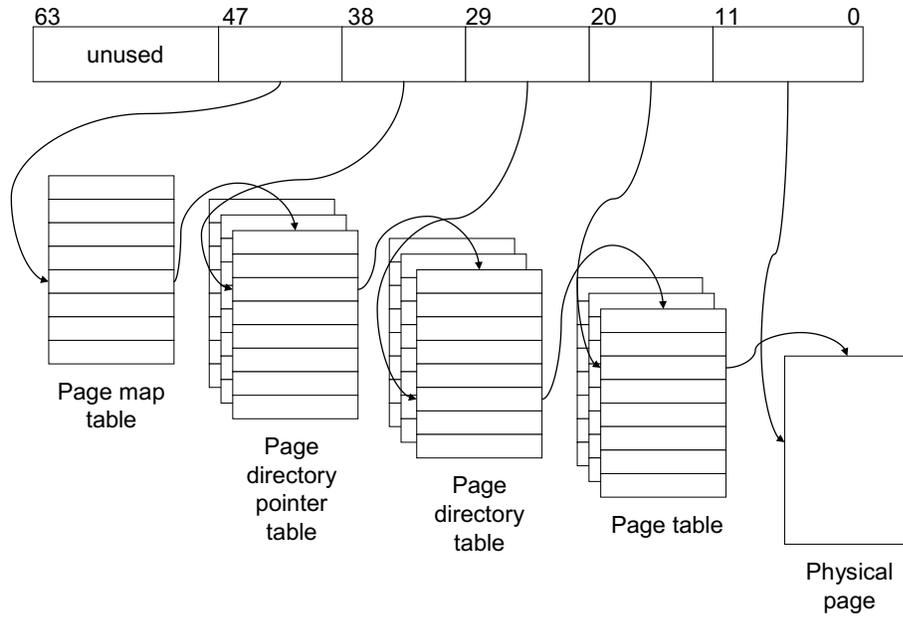
2. Recall that in Unix, when a file whose set-user-id bit is one is *exec*'d, the process's effective *user ID* and the saved *user ID* become the owner of the file, while the *real user ID* is not modified. A program may set its effective user ID to be anything if it's currently *superuser*, otherwise it may set its effective user ID only to its real, effective, or saved user IDs. Also, file access checks are done using a process's effective UID. A set-UID program uses these features to check if its caller has permission to access a given file, perhaps one passed as an argument. (It first sets its effective UID to that of the caller, opens the file, then sets its effective UID back to the saved UID.) A primary use of these features is to allow the owner of a file to keep others from directly accessing the file, but provide a program (whose set-UID bit is set) that performs operations on the file on the behalf of others, but under the control of the file's owner.
  - a. This approach works well if the set-UID program is owned by superuser — such programs are effectively extensions to the kernel. However, it's not that useful if the file containing the program is owned by a normal (untrusted) user. Explain why not.

- b. We'd like to fix this: we want to make it possible for some user, say Mary, to create a program that others may run as set-UID-to-Mary, with none of the problems mentioned in the answer to part a. Explain how this might be done. Hint: certain functionality may have to be turned off, and we may need to adopt a convention for passing files as parameters to Mary's program.





5. The x64 architecture supports both 4KB and 2MB pages as shown below.





- c. Consider two possible implementations of fork using copy on write. One, which we call shallow copy on write, involves sharing just the actual page frames, but both processes have their own private copies of the all page tables (including page directory tables, page directory pointer tables, etc.). The other, which we call deep copy on write, involves sharing both page frames and all page tables. Thus with deep copy on write, after fork, both parent and child share the page map table and nothing is immediately copied. When page frames are modified, all page tables (and page directory tables, etc.) that must be modified are also copied. But with shallow copy on write, when a process forks, all page tables, including the page map table, are copied at the time of the fork. It should be clear that, with deep copy on write, forks are cheaper than with shallow copy on write. Why might it be advantageous, nevertheless, to use shallow copy on write? While we're looking for a qualitative and not a quantitative explanation, you should mention any important data structures that might be required in one but not the other.
- d. How might one optimize shallow copy on write to make it much cheaper for the usual case of fork, in which it's followed quite soon by an exec? (Hint: this is yet another application of "lazy evaluation").



- c. In DFS, servers give clients tokens that allow clients to use cached data from files. These tokens have time limits associated with them, so that if the token times out, the client has to assume that any cached data it has from the server is invalid and must obtain a new token (and new data) from the server. Explain how tokens might be used to help with the situation of part b. (Hint: it won't necessarily solve the problem, but it might help client applications become aware of problems.)
- d. Explain why DFS does not support hard-mounted file systems.