

CMPU-224 Lab12 Quiz

Spring 2026

Name: _____

This is a closed book, closed notes quiz. No electronic devices are allowed. You have until 3:30pm to complete the quiz. There are a total of 5 questions and 10 points.

There should be enough space on the quiz for your answers. If you need more space to work out a problem, blank paper will be available, just ask.

If you finish with time remaining, raise your hand and I will come and collect your quiz. You may then work on the lab assignment.

Good Luck!

- (2 points) Which statement best describes what `pthread_join` does?
 - Creates a new thread that runs the given start function.
 - Blocks the calling thread until the target thread terminates, optionally collecting its return value.
 - Acquires a mutex so the calling thread can safely enter a critical section.
 - Detaches a thread so its resources are reclaimed automatically when it exits.
- (2 points) Two threads each execute `g_counter++` in a loop, where `g_counter` is an unprotected global `int`. After both threads finish, the final value is sometimes lower than expected. Which statement best explains why?
 - The compiler emits incorrect code for the `++` operator on shared variables.
 - The variable overflows because `int` is too small to hold the result.
 - `g_counter++` compiles to a load / modify / store sequence; if two threads interleave those steps, one thread's update can overwrite the other's and an increment is silently lost.
 - `pthread_create` occasionally fails to launch one of the threads, so its iterations never run.
- (2 points) Consider two implementations a parallel computation. In implementation A, every thread acquires a shared mutex *inside* its inner loop and updates a shared variable on each iteration. In implementation B, every thread accumulates into a local variable inside its inner loop and acquires the mutex *once at the end* to merge its private result into the shared variable. Why is B typically much faster than A?
 - B uses fewer threads than A.
 - B is actually incorrect; B only appears faster because it skips work.
 - B shrinks the critical section and removes nearly all lock contention — threads no longer fight for the mutex on every iteration.
 - B avoids using `malloc`, while A must allocate memory on every iteration.

4. (2 points) Consider the following thread function:

```
typedef struct { int x; int y; } pair_t;

void *worker(void *arg) {
    pair_t result;          /* declared here */
    result.x = 1;
    result.y = 2;
    return (void *) &result;
}
```

The main thread joins this worker and dereferences the returned pointer. What is wrong with this code?

- A. Nothing — the runtime preserves a thread's stack until every other thread has read from it.
 - B. The cast `(void *)&result` is illegal in C.
 - C. `result` lives on the worker's stack, which is reclaimed when `worker` returns. The pointer the parent receives is dangling, so dereferencing it is undefined behavior.
 - D. `pthread_join` cannot pass a pointer back to the parent thread; it can only return an integer status.
5. (2 points) A program calls `pthread_create` several times to launch worker threads, but *never* calls `pthread_join` on any of them. `main` then prints "done" and returns. Which statement is most accurate?
- A. `pthread_join` only matters when you need a worker's return value; if you don't, omitting it is harmless and any `printf` side effects are still guaranteed to appear.
 - B. The parent has no way to wait for the workers, so the process may exit before some of them finish; their output can be missing, truncated, or interleaved unpredictably.
 - C. The program fails to compile because every `pthread_create` requires a matching `pthread_join`.
 - D. The output is identical to a version that does call `pthread_join`, just slightly faster.