

Slide 1

Administrivia

- Open lab hours announced via e-mail. Updated office hours posted on my Web page, etc.
- Homework 1 coming soon (by Wednesday).
- A comment about the 8/28 minute essay (about CSCI 2321): No, it's not a prerequisite. Most people said "no".

Slide 2

Basics of Multithreaded (Shared-Memory Parallel) Programming

- Multithreaded programs (in contrast to message-passing programs, as we'll see) typically start out with one "thread" and then create others as they execute.
- All threads have access to the same address space, so variables can be shared. Avoids need to communicate data explicitly, but means sometimes we need to "synchronize" among threads, e.g., to make sure only thread at a time can access a particular variable.

OpenMP

Slide 3

- Early work on message-passing programming resulted in many competing programming environments — but eventually, MPI emerged as a standard.
- Similarly, many different programming environments for shared-memory programming, but OpenMP may be emerging as a standard.
- In both cases, idea was to come up with a single standard, then allow many implementations. For MPI, standard defines concepts and library. For OpenMP, standard defines concepts, library, *and compiler directives*.
- First release 1997 (for Fortran, followed in 1998 by version for C/C++).
- Several production-quality commercial compilers available. Up until very recently, free compilers were, um, “research software” or in work. Latest versions of GNU compilers, though, offer support. !!

What’s an OpenMP Program Like?

Slide 4

- Fork/join model — “master thread” spawns a “team of threads”, which execute in parallel until done, then rejoin main thread. Can do this once in program, or multiple times.
- Source code in C/C++/Fortran, with OpenMP compiler directives (`#pragma` — ignored if compiling with a compiler that doesn’t support OpenMP) and (possibly) calls to OpenMP functions.
Compiler must translate compiler directives into calls to appropriate functions (to start threads, wait for them to finish, etc.)
- A plus — can start with sequential program, add parallelism incrementally — usually by finding most time-consuming loops and splitting them among threads.
- Number of threads controlled by environment variable (roughly analogous to “number of processes” parameter for `mpirun`), or from within program.

Simple Example / Compiling and Executing

- Look at simple program — `hello.c` on sample programs page.
- Compile with compiler supporting OpenMP.
- Execute like regular program. Can set environment variable `OMP_NUM_THREADS` to specify number of threads. Default value seems to be one thread per processor.

Slide 5

How Do Threads Interact?

- With MPI, processes don't share an address space, so to communicate they must use messages. With OpenMP, threads do share an address space, so they communicate by sharing variables.
- Sharing variables is more convenient, may seem more natural.
- However, "race conditions" are possible — program's outcome depends on scheduling of threads, often giving wrong results.
What to do? use synchronization to control access to shared variables.
Works, but takes (execution) time, so good performance depends on using it wisely.

Slide 6

Slide 7

OpenMP Constructs — Basic Categories

- Parallel regions (“replicate the following in all threads”).
- Worksharing (“divide the following among threads”).
- Data environment (shared variables versus per-thread variables).
- Synchronization.
- Runtime functions / environment variables.

Slide 8

Parallel Regions in OpenMP

- `#pragma omp parallel` tells compiler to do following block in all threads (starting team of threads if necessary). Execution doesn't proceed in main thread until all are done. Example — “hello world” shown earlier.
- Block must be a “structured block” — block with one point of entry (at top) and one point of exit (at bottom). In C/C++, this is a statement or statements enclosed in brackets (with no `gotos` into / out of block).

Worksharing Constructs in OpenMP

Slide 9

- `#pragma omp parallel for` tells compiler to split iterations of following `for` loop among threads. By default, main thread doesn't continue until all are done, but can override that (might be useful if you have two consecutive such loops).
- How loop iterations are mapped onto threads — controlled by `schedule` clause. More about this later.
- To make different threads do different things — `#pragma parallel sections`, etc. (More in standard.)

A Little About Variables in OpenMP

Slide 10

- Most variables are shared by default, including any global variables.
- Some things, though, aren't — variables within a statement block, stack (local) variables in subprograms called from parallel region.
- Can specify that each thread gets its own copy with `private` clause.
- Can specify that each thread gets its own copy, and copies are combined at the end, with `reduction` clause.

Example — Numerical Integration

- Compute π by integrating $\int_0^1 \frac{4}{1+x^2} dx$.
- Do this numerically by approximating area under curve by many small rectangles, computing their area, adding results.
- Sequential program fairly straightforward. (`num-int-seq.c` on “sample programs” page).
- “Parallelize” how? (`num-int-par.c` on “sample programs” page).

Slide 11

Minute Essay

- None — sign in.

Slide 12