

Administrivia

- (Reminders about what's due when.)
- Sample solutions for all regular homeworks to be on Web.

Slide 1

More Administrivia

- "What about our grades?" You will get information by e-mail as soon as I have it.
- Recall(?) weights from syllabus.

Slide 2

Programming Environments, Revisited

Slide 3

- Choice of environments for book was based on how things were when it was written — MPI fairly dominant for distributed memory and OpenMP for shared memory, with Java not so widely used for parallel programming but more familiar/available.
- All three include more than we had time to cover in class, and have continued to evolve, and then there's a whole new hardware platform (GPUs) . . .
- (I will leave the rest of the slides in, but we didn't go over them in class.)

OpenMP Revisited

Slide 4

- OpenMP worksharing constructs define "implicit tasks" (one per thread). We looked only at parallel loops, but there are also "parallel sections", which allow for nesting/recursion.
- OpenMP 3.0 adds support for explicit tasks, which may help with some kinds of problems (irregular and recursive).

MPI Revisited

- Even MPI 1.0 includes far more than we could cover in class — many collective communication operations, communicators, process topologies, and support for user-defined data types in messages.
- MPI 2.0 and later versions add more — e.g., process spawning and one-sided communication.

Slide 5

Java Revisited

- Java 1.5 brought into the standard library a lot of classes previously available as third-party additions — thread pools, locks, various shared-data classes, etc.
- Java memory model also cleaned up a bit.
- (Curiously enough, though, the need for explicit multithreading in GUIs seems to have declined, with the notion of the EDT and new classes such as `SwingWorker` and timers.)

Slide 6

OpenCL Revisited

Slide 7

- Graphics processors emerging as a new platform for parallel computing — hardware is becoming sophisticated enough to support computation beyond the cards' original purpose, so why not put it to use?
- No consensus yet about programming environments, but OpenCL might emerge as a semi-standard, as MPI and OpenMP did.
- We barely scratched the surface of this environment but perhaps did enough to get past the initial intimidation factor. A brief recap . . .

A Little About GPU Hardware

Slide 8

- Processing hardware typically includes many processors working more or less in lockstep, each able to do pipelined/vector operations — i.e., SIMD, making a comeback!
- Typical hardware also includes a possibly-complex memory hierarchy separate from the memory hierarchy of the "host computer".

A Little About Programming for GPU Hardware

Slide 9

- SIMD hardware makes a data-parallel style of programming a good fit. Not something we really address in our pattern language (yet!), but conceptually similar to *Geometric Decomposition* but more closely synchronized. A.k.a. “stream processing”?
- So, you might express computations as a sequence of whole-array operations, or in terms of applying a “computational kernel” in parallel to many data elements. Whole-array operations included in some programming environments (e.g., Fortran). Current programming environments for GPUs (NVIDIA’s CUDA, e.g., and OpenCL) use the computational-kernel idea.
- Currently moving data back and forth between host’s memory and GPU’s memory must be done explicitly. Actually maybe not a bad idea given that it does take time?

Review of Course

Slide 10

- “PAD I for parallel programming”? We covered:
 - Three languages/libraries — OpenMP, MPI, Java.
 - How to find and exploit concurrency in programs.
- We also did several running examples and some homeworks . . .

Review of Homeworks

Slide 11

- Homeworks 1 and 2 — estimating π with Monte Carlo methods. Basic structure is *Task Parallelism*. Complication is that you need a thread-safe RNG.
- Homework 3 — Conway's game of life. Basic structure is *Geometric Decomposition*. Basic idea easy, details a bit messy (particularly for MPI).
- Homework 4 — quicksort. Basic structure is *Divide and Conquer*.
- For all programs, probably need large problem sizes to get any benefit from multiple UEs. Even then performance may not be amazingly good, but the primary goal is pedagogical rather than practical.

Minute Essay

Slide 12

- None — sign in.