

Slide 1

Administrivia

- (None?)

Slide 2

Minute Essay From Last Lecture

- About Homework 3, it sounds like most people didn't have too much trouble with the sequential and OpenMP versions, but the MPI version was another story. No surprise! but I hope once you get good working code you will understand MPI much better.

Slide 3

Example Application: Matrix Multiplication

- Basic problem is easily turned into code.
- A variant of the obvious algorithm is to take a block-based approach. Can be significantly faster, probably because it makes better of caches.
- (Look at code for both versions. Notice that it represents matrices as 1D arrays. At least in part this is because it was originally written by a very-old-school C programmer.)

Slide 4

Parallelization — *Finding Concurrency*

- Obvious decomposition for simple approach is task-based, with one task per point. Tasks are completely independent.
- For block-based approach, may make more sense to think in terms of decomposing data into blocks; then tasks correspond to computing blocks of C . Again, though, they're independent.

Slide 5

Parallelization — *Algorithm Structure* (Shared Memory)

- For the simple approach, we have many mostly-independent tasks, forming a flat set rather than a hierarchy, so *Task Parallelism* seems like a good choice. Block-based program is similar.
- Key design decision is how to assign tasks to UEs.
- Probably makes sense to group tasks by rows rather than individual points and to use a simple static assignment of tasks to UEs, and group tasks by — what? for simple approach, two obvious choices; for block-based approach, more. Could try several and see which seems to work best.

Slide 6

Parallelization — *Supporting Structures and Code* (Shared Memory)

- For program structure, *Loop Parallelism* makes sense.
- Code in OpenMP is straightforward.
- (Look at code.)

Slide 7

Parallelization — *Algorithm Structure (Distributed Memory)*

- For distributed memory we have to think about how to distribute C and how to duplicate/distribute A and B . Might work better to think in terms of block-based approach and data decomposition — so *Geometric Decomposition* might be a better fit.
- Key design decisions here are how to decompose data and assign chunks to UEs, and then how to manage synchronization/communication for update operation.
- Probably makes sense to decompose data so we can assign one block of C to each UE — amount of work per block is pretty much constant.

Slide 8

Parallelization — *Algorithm Structure (Distributed Memory), Continued*

- For each block of C , computation can be thought of a sequence of update operations, each involving a different combination of blocks of A and B . (Compare how this fits overall idea of *Geometric Decomposition* with how heat-diffusion example fits.)
- This tells us what kind of communication we need. Simple approach is to broadcast two blocks at each step, one for “row” and one for “column”. More complex, but more efficient, version involves rotating blocks among processes.

Parallelization — *Supporting Structures* (Distributed Memory)

- For program structure, we probably want *SPMD* (especially if using MPI or similar programming environment).
- *Distributed Array* is relevant, especially for parts of sample/test program that initialize and print array (since they use each array element's global indices).

Slide 9

Parallelization — Code (Distributed Memory)

- If we distribute all three arrays (which seems like a good idea), we have to make changes in code to initialize and print as well as matrix-multiplication. As is often the case with programs using *Distributed Array*, the ideas are simple but the code inclined to be messy.
- For actual multiplication, each process will update one "chunk", doing the same computation done in the block-based sequential program, but with communication operations to broadcast two blocks per step.
- (Look at code.)

Slide 10

Minute Essay

- None really — sign in?

Slide 11

Minute Essay Answer

- FIXME

Slide 12