

Slide 1

### Administrivia

- (You all got my e-mail of last week, about appendices, right?)

Slide 2

### Numerical Integration Problem in OpenMP, Revisited

- Anyone have interesting results to report?

## MPI — the Message Passing Interface

Slide 3

- Idea was to come up with a single standard (concepts and library) for message-passing programs, then allow many implementations. Similar to language standards (C, C++, etc.). Good for portability.
- MPI Forum — international consortium — began work in 1992. MPI 1.1 and MPI 2.0 standards defined. Huge! 1.1 specification is 500+ pages, and 2.0 standard even bigger.
- Original reference implementation — MPICH (Argonne National Lab). LAM/MPI (Local Area Multicomputer) is another free implementation. Latest / most popular may be OpenMPI (installed here).

## What's an MPI Program Like?

Slide 4

- “SPMD” (Single Program, Multiple Data) model — many processes, all running the same source code, but each with its own memory space and each with a different ID. Could take different paths through the code depending on ID.
- Source code in C/C++/Fortran, with calls to MPI library functions.
- How programs get started isn't specified by the (first) standard! (for historical/political reasons — some early target platforms were very restrictive, would not have supported what academic-CS types wanted).
- (Compare and contrast all of the above with OpenMP.)

Slide 5

### What's in the MPI Library?

- Setup and bookkeeping — initialization, cleanup, environment query, etc.
- Data management — pack/unpack, derived data types.
- Point-to-point communication — several varieties, differing mostly in how much synchronization.
- Collective operations — e.g., broadcast.

Slide 6

### MPI “Communicators”

- (One more thing to define before we can write simple code.)
- MPI allows grouping processes; group plus associated context called a “communicator”. Makes it easier to write “safe” parallel libraries.
- Predefined communicator `MPI_COMM_WORLD` includes all processes. Programmers can create additional ones.

### Simple Examples / Compiling and Executing

Slide 7

- Look at sample program `hello.c`. (All sample programs from class should be on the Web, linked from course “sample programs” page, with short instructions on how to use MPI. You will need to do some setup before MPI programs will run.)
- We'll use OpenMPI as installed on the lab machines. There should be man pages for all commands and functions.
- Compile with `mpicc`.
- Execute with `mpirun`.

### Simple (Blocking) Point-to-Point Communication in MPI

Slide 8

- Send with `MPI_Send` — returns as soon as data has been copied to system buffer, buffer in program can be reused.
- Receive with `MPI_Recv` — waits until message has been received.
- Can use “tags” to distinguish between kinds of messages. Can receive selectively or not (`MPI_ANY_TAG`). Received tag is in returned `MPI_Status` variable (e.g., `status.MPI_TAG`).
- Can receive from specific sender or from any sender. (`MPI_ANY_SOURCE`). Sender is in returned `MPI_Status` variable (e.g., `status.MPI_SOURCE`).
- For `MPI_Recv`, “length” parameter specifies buffer length. Use `MPI_Get_count` to get actual count.
- Look at sample program `send-recv.c`.

Slide 9

### Not-So-Simple Point-to-Point Communication in MPI

- For not-too-long messages and when readability is more important than performance, `MPI_Send` and `MPI_Recv` are probably fine.
- If messages are long, however, buffering can be a problem, and can even lead to deadlock. Also, sometimes it's nice to be able to overlap computation and communication.
- Therefore, MPI offers several other kinds of send/receive functions — “synchronous” (blocks both sender and receiver until communication can take place), “non-blocking” (doesn't block at all, program must later test/wait for communication to take place).  
(More about these later.)

Slide 10

### Collective Communication in MPI

- “Collective communication” operation — one that involves many processes (typically all, or all in MPI “communicator”).
- Could implement using point-to-point message passing, but some operations are common enough to be library functions — broadcast (`MPI_Bcast`), “reduction” (`MPI_Reduce`), etc.

## Minute Essay

- None — sign in.

Slide 11