







Parallel Programming Models

Shared Memory

 tasks share a common address space, which they read and write asynchronously.

- Threads
 - a single process can have multiple, concurrent execution paths. Example implementations: POSIX threads & OpenMP

Message Passing tasks exchange data through communications by sending and receiving messages. Example: MPI

- Data Parallel languages
 tasks perform the same operation on their partition of work. Example: Co-array Fortran (CAF), Unified Parallel C (UPC), Chapel
- Hybrid
 MPI + OpenMP

TUDelft













Domain Decomposition





Domain Decomposition Find the largest element of an array				
Core 0	Core 1	Core 2	Core 3	
5 13 1 9 26 13	6 49 34 50 22 12	68 12 98 16 78 31	83 51 94 27 74 64	
13	49	68	83	
TU Delft				



Domain Decomposition Find the largest element of an array Core 0 Core 1 Core 2 Core 3 5 13 1 9 6 49 34 50 22 12 66 12 98 16 78 18 51 94 27 74 64 Joint Core 2 Core 3 Joint Core 2 Joint Core 2 Joint Core 3 </

rúDelft







Domain Decomposition









Task/Functional Decomposition

- The problem is decomposed according to the work that must be done. Each task then performs a portion of the overall work.
- Divide computation based on natural set of independent tasks Assign data for each task as needed
- Example: pipeline seismic data pre-processing
 - static-correction
 - deconvolution
- nmo correction stacking



tuDelft







What Is OpenMP?

- Compiler directives for multithreaded programming
- Easy to create threaded Fortran and C/C++ codes
- Supports data parallelism model
- Portable and Standard
- Incremental parallelismCombines serial and parallel code in single source

TUDelft









Programming Model

• Explicit parallelism:

- > All processes starts at the same time at the same point in the code
- ▶ Full parallelism: there is no sequential part in the program





Data Model

- All variables and data structures are local to the process
- Processes can exchange data by sending and receiving messages

MPI advantages

- Mature and well understood
 - Backed by widely-supported formal standard (1992)
 - Porting is "easy"
- Efficiently matches the hardware
 - Vendor and public implementations available
- User interface:
 - Efficient and simple
 - Buffer handling
 - Allow high-level abstractions
- Performance

TUDelft

44

Work Distribution

• All processors run the same executable.

• Parallel work distribution must be explicitly programmed into the algorithm:

47

- domain decomposition
- master worker

Basic Concepts TUDelft 46

Data and Work Distribution • To communicate together mpi-processes need identifiers: rank = identifying number • all distribution decisions are based on the rank - i.e., which process works on which data myrank= (**size**-1) (myrank=2 myrank=0 (myrank=1 data data data data program program program <mark>program</mark> communication network **T**UDelft 48

51

TUDelft

Synchronous Sends in

53

- Check marks will appear next to each message you send. Here's what each one indicates:
- The message was successfully sent.
- *I* The message was successfully delivered to the recipient's phone or any of their linked devices.
- 🗸 The recipient has read your message.

tUDelft

Non-Blocking Operations

• Non-blocking operations return immediately and allow the subprogram to perform other work.

Collective Communications

- Collective communication routines are higher level routines.
- Several processes are involved at a time.
- May allow optimized internal implementations, e.g., tree based algorithms

rúDelft

56

Reduction Operations • Combine data from several processes to produce a single result. **Sum=? 15 10 30 30 5 T 5**

