HPC.NRW

MPI in Small Bites

PPCES 2024

HPC.NRW Competence Network

Wrap-up and Outlook

HPC.NRW Competence Network

MPI in Small Bites

MPI Course Summary

HPC.NRW

- MPI Concepts
  - Blocking vs. Nonblocking procedures
  - Asynchronous vs. Synchronous operations
  - Communication paradigms: point-to-point, collective, one-sided
- Point-to-point communication
  - Data exchange between **two** processes
  - Use standard send (MPI_Send) by default and synchronous send (MPI_Ssend) where needed
    - Standard send provides the best of buffered vs. synchronous protocols
    - Correctness of your application should never rely on a specific behavior (buffered vs. synchronous)
  - Nonblocking communication …
    - … can help to avoid deadlocks in data exchanges (communication-communication overlap)
    - … may (or **may not**) progress communication on the side (communication-computation overlap)

MPI Course Summary

HPC.NRW

- MPI Datatypes
  - Describe memory access patterns (type map)
  - Are local to a process
  - Type signatures (sequence of elementary types) need to match for a communication

- Blocking Collective Communication
  - Implement common parallel data distribution patterns
  - All ranks in the communicator must participate
  - Standard behavior (across platforms/implementation)
    - However, **no** performance portability

MPI Course Summary

HPC.NRW

- Communicator Basics
  - Communicators provide isolation of communication contexts
  - Addressing based on ranks
    - Ranks of a process may differ in distinct communicators/groups
  - Ranks in communicator are a characteristics of the underlying process group
  - Communicators and groups can be derived from each other

- Hybrid Programming
  - MPI is largely thread-agnostic (provides selection of different thread levels)
  - High thread level may impact overall performance
  - Be aware of potential data races when combining MPI with threading

MPI Course Outlook (What we didn't cover)

HPC.NRW

- Advanced MPI Topics
  - Persistent communication (a specialization of nonblocking communication)
  - Nonblocking collective communication
  - Virtual process topologies (Cartesian and general Graph topologies)
  - Neighborhood Collectives (based on virtual process topologies)
  - Dynamic process management & the MPMD model
  - One-sided communication
  - File I/O
  - The Session Model vs. the World Model (introduced with MPI 4.0)
  - Partitioned Point-to-point communication (introduced with MPI 4.0)

Evaluation

Feedback Wanted!

PPCES Evaluation Form

(valid until March 29th, 2024)

https://www.soscisurvey.net/ppces2024