



PPCES 2021 – MPI Course

HPC.NRW Competence Network



THE COMPETENCE NETWORK FOR HIGH PERFORMANCE COMPUTING IN NRW.

Wrap-up and Outlook

HPC.NRW Competence Network

MPI in Small Bites

- 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 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

– 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 3.1 Topics

- Persistent communication (a specialization of nonblocking communication)
- Nonblocking collective communication
- Virtual process topologies (Cartesian and general Graph topologies)
- Dynamic process management & the MPMD model
- One-sided communication
- File I/O

– MPI 4.0 Topics

- The Session Model vs. the World Model
- Partitioned Point-to-point communication