



Friedrich-Alexander-Universität Erlangen-Nürnberg

# Introduction to Parallel Programming with MPI

#### Dr. Alireza Ghasemi, Dr. Georg Hager

Erlangen National High Performance Computing Center

**Nonblocking Collectives** 



## Nonblocking Collectives in MPI

Similar to blocking collectives: nonblocking collective calls including all ranks of a communicator All ranks must call the function!

- Nonblocking variants (since MPI 3.0): buffer can be used after completion (MPI\_Wait\*/MPI\_Test\*)
- Local: not synchronization
- Multiple outstanding collective communications on same communicator
- Cannot interfere with point-to-point communication
  - Completely separate modes of operation!
- Cannot interfere with blocking collective communication
  - Such interference was allowed in point-to-point communication

## **Collectives in MPI**

- Rules for all collectives
  - Data type matching
  - No tags
  - Count must be exact, i.e., there is only one message length, buffer must be large enough

## **Collectives in MPI**

- Rules for all collectives
  - Data type matching
  - No tags
  - Count must be exact, i.e., there is only one message length, buffer must be large enough
- Types:
  - Synchronization (barrier)
  - Data movement (broadcast, scatter, gather, all to all)
  - Collective computation (reduction, scan)
  - Combinations of data movement and computation (reduction + broadcast)

## **Collectives in MPI**

- Rules for all collectives
  - Data type matching
  - No tags
  - Count must be exact, i.e., there is only one message length, buffer must be large enough
- Types:
  - Synchronization (barrier)
  - Data movement (broadcast, scatter, gather, all to all)
  - Collective computation (reduction, scan)
  - Combinations of data movement and computation (reduction + broadcast)
- General assumption: MPI does a better job at collectives than you trying to emulate them with point-to-point calls

Nonblocking synchronization

MPI\_Ibarrier(MPI\_Comm comm, MPI\_Request \*request)

- Must be followed by an MPI\_Wait
- Calling process enters the barrier, no synchronization happens
- Synchronization may happen asynchronously
- Overlapping synchronization with work: reducing idle time
- Comparison:
  - 1) **MPI\_Ibarrier** Work MPI\_Wait

2) Work — MPI\_Barrier idle times before and after Work differ on each process and their sum!

## Collectives: Blocking vs. Nonblocking

- Broadcast:
  - MPI\_Bcast(buf,count,datatype,root,comm);
  - MPI\_Ibcast(buf,count,datatype,root,comm,request);
- Scatter:
  - MPI\_Scatter(sendbuf,sendcount,sendtype,recvbuf,recvcount, recvtype,root,comm);
  - MPI\_Iscatter(sendbuf,sendcount,sendtype,recvbuf,recvcount, recvtype,root,comm,request);
- Gather:
  - MPI\_Gather(sendbuf,sendcount,sendtype,recvbuf,recvcount, recvtype,root,comm);
  - MPI\_Igather(sendbuf,sendcount,sendtype,recvbuf,recvcount, recvtype,root,comm,request);

## Collectives: Blocking vs. Nonblocking

- Similarly many blocking collective calls have nonblocking analogues:
  - MPI\_Iallgather, MPI\_Ialltoall, MPI\_Ireduce, ...

Remarks for MPI\_Ireduce:

- Both send and receive buffers can be used only after completion!
- Similar to nonblocking point-to-point calls, MPI\_Wait\* or MPI\_Test\* must be used to halt or examine for the completion of a request, respectively
- root (if available) and comm must be the same on all processes
- Type signature of send and receive variables must match

## Nonblocking reduction on all ranks

- No root
- sendbuf, recvbuf can be reused after completion: requires MPI\_Wait or MPI\_Test
- Recvbuf is significant on all processes

- 1. Nonblocking collectives are useful when there exist multiple collective calls on the same communicator so overlapping different collective calls becomes possible.
  - a. Correct b. Incorrect

- 1. Nonblocking collectives are useful when there exist multiple collective calls on the same communicator so overlapping different collective calls becomes possible.
  - a. Correct b. Incorrect

Answer: a.

- 1. Nonblocking collectives are useful when there exist multiple collective calls on the same communicator so overlapping different collective calls becomes possible.
  - a. Correct b. Incorrect

Answer: a.

- 2. Nonblocking collectives can interfere with blocking collectives, that is, some processes of the communicator call the nonblocking binding and the rest blocking counterpart.
  - a. Correct b. Incorrect

- 1. Nonblocking collectives are useful when there exist multiple collective calls on the same communicator so overlapping different collective calls becomes possible.
  - a. Correct b. Incorrect

Answer: a.

- 2. Nonblocking collectives can interfere with blocking collectives, that is, some processes of the communicator call the nonblocking binding and the rest blocking counterpart.
  - a. Correct b. Incorrect

Answer: b.