

Message Passing Interface

Lecture 2

Manfred Liebmann
Technische Universität München
Chair of Optimal Control
Center for Mathematical Sciences, M17
`manfred.liebmann@tum.de`



Technische Universität München



Fakultät für Mathematik

November 23, 2015

Communication Models of MPI

- Two-sided communication
 - Blocking and non-blocking point-to-point communication
 - Collective communication routines
- One-sided communication
 - Only one process is involved actively in the communication
 - Put and get operations to and from remote memory

Blocking and Non-Blocking Point-to-Point Communication

- General initialization
 - MPI_Init, MPI_Finalize (initialization)
 - MPI_Comm_rank, MPI_Comm_size (process information)
 - MPI_COMM_WORLD (communicator)
- Blocking point-to-point communication
 - MPI_Send, MPI_Recv (blocking send and receive)
 - MPI_Status (struct)
 - MPI_Get_count
- Non-blocking point-to-point communication
 - MPI_Isend, MPI_Irecv (blocking send and receive)
 - MPI_Status, MPI_Request (structs)
 - MPI_Waitall (synchronization)

Collective Operations: Synchronization




Figure 1: MPI_Barrier: Barrier synchronization

Collective Operations: Broadcast




Figure 2: MPI_Bcast: Broadcast operation

Collective Operations: Scatter




Figure 3: MPI_Scatter: Scatter operation

Collective Operations: Gather and Allgather




Figure 4: MPI_Gather: Gather operation




Figure 5: MPI_Allgather: All gather operation

Collective Operations: Reduce and Allreduce

MPI_Reduce




Figure 6: MPI_Reduce: Reduce operation

MPI_Allreduce




Figure 7: MPI_Allreduce: All reduce operation

Collective Operations: Reduce Operation Types

- Reduce operation types
 - MPI_LAND - logical AND
 - MPI_LOR - logical OR
 - MPI_LXOR - logical XOR
 - MPI_BAND - binary AND
 - MPI_BOR - binary OR
 - MPI_BXOR - binary XOR
 - MPI_MIN - minimum
 - MPI_MAX - maximum
 - MPI_SUM - sum
 - MPI_PROD - product
 - MPI_MINLOC - minimum and rank
 - MPI_MAXLOC - maximum and rank

Collective Operations: Alltoall and Alltoallv

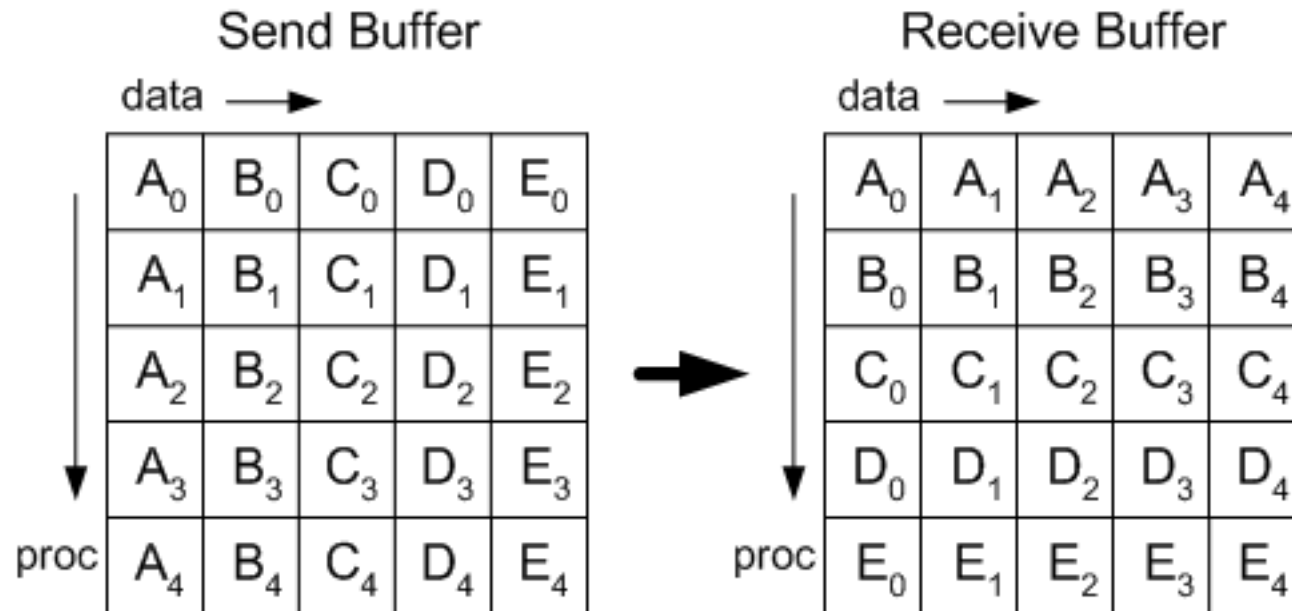


Figure 8: MPI_Alltoall: Alltoall communication pattern

The MPI_Alltoallv variant supports different lengths of the vectors sent to the other processes encoded in additional count and displacement vectors. Other variants are MPI_Scatterv, MPI_Gatherv, and MPI_Allgatherv.