Module 3
Distributed-memory parallel programming with MPI

Slides are based on Chapters 9 & 10 of the following book:

Hager & Wellein, *Introduction to High Performance Computing for Scientists and Engineers*

Overview of Module 3

- The concept of explicit "message passing"
- Introduction to **MPI** (the Message Passing Interface)
- Guidelines for efficient MPI programming

Message passing

- Shared-memory programming (such as OpenMP) does not work for distributed-memory parallel computers
  - There is no way for one process to directly access the address space of another process that resides on a different node
- Explicit "**message passing**" is required on distributed-memory systems
  - This is a programming model that will also work for shared-memory or hybrid systems

Basic features of message-passing programming

- The same program runs on all processes (Single Program Multiple Data, or **SPMD**)—no difference from OpenMP programming in this regard
- The work of each process is implementation in a sequential language (such as C)
  - Data exchange (sending and receiving messages) is done via calls to an appropriate library
- All variables in a process are **local** to this process (nothing is shared)

"Messages" carry data to be exchanged between processes

Information needed about a message:

- Which process is the sender of the message?
- Where is the data located on the sending process?
- What kind of data is to be sent?
- How much data constitutes the message?
- Which process is the receiver of the message?
- Where should the data be placed on the receiving process?
- What amount of data is the receiving process prepared to accept?

MPI (Message Passing Interface)

MPI is a *library standard* for programming distributed memory

- MPI implementation(s) available on almost every major parallel platform (also on shared-memory machines)
- Portability, good performance & functionality
- Collaborative computing by a group of individual processes
- Each process has its own local memory
- Explicit message passing enables information exchange and collaboration between processes

MPI C language binding

```
#include <mpi.h>

rv = MPI_Xxxxx(parameter, ... )
```

Beware of the **case-sensitive** naming pattern.

The return value (rv) transports information about the success of the MPI operation. (**MPI_SUCCESS** is returned if the MPI routine completed successfully.)

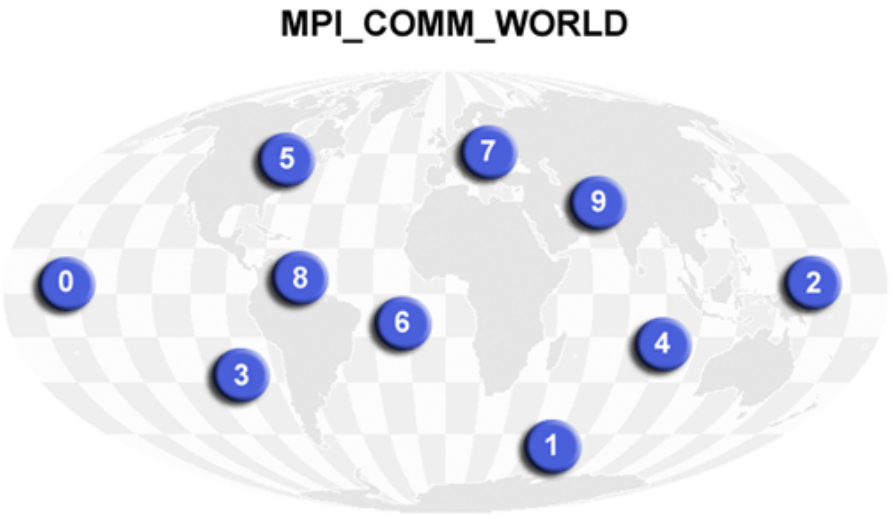
MPI_Init

The `MPI_Init` function initializes the parallel environment.

```
int MPI_Init( int *argc, char ***argv )
```

The `MPI_Init` function takes pointers to the `main()` function's input arguments so that the library can evaluate and remove any additional command line arguments that may have been added by the MPI startup process.

MPI communicator



https://computing.llnl.gov/tutorials/mpi/

- An MPI communicator works as a "communication universe" for a group of processes
- `MPI_COMM_WORLD` – name of the default MPI communicator, i.e., the collection of all the processes started by an execution
- Each process in a communicator is identified by its unique rank
- Almost every MPI command needs to provide a communicator as input argument

MPI process rank

- The unique rank for each process is an integer identifier within a communicator
- The rank value is between `0` and `#procs-1`
- The rank value is used to distinguish one process from another
- Commands `MPI_Comm_size` & `MPI_Comm_rank` are very useful

```
int size, my_rank;
MPI_Comm_size (MPI_COMM_WORLD, &size);
MPI_Comm_rank (MPI_COMM_WORLD, &my_rank);
```

MPI_Finalize

An MPI parallel program is shut down by a call to `MPI_Finalize()`.

Note that no MPI process except rank 0 is guaranteed to execute any code beyond `MPI_Finalize()`.

“Hello world” in MPI

```
#include <stdio.h>
#include <mpi.h>

int main (int nargs, char** args)
{
  int size, my_rank;
  MPI_Init (&nargs, &args);
  MPI_Comm_size (MPI_COMM_WORLD, &size);
  MPI_Comm_rank (MPI_COMM_WORLD, &my_rank);
  printf("Hello world, I've rank %d out of %d procs.\n",
         my_rank,size);
  MPI_Finalize ();
  return 0;
}
```

Example running result of "Hello world"

- Example of compilation: `mpicc hello_world_mpi.c`
- Example of parallel execution: `mpirun -np 4 ./a.out`
- Example running result (note: no deterministic order of the output):

```
Hello world, I've rank 2 out of 4 procs.
Hello world, I've rank 1 out of 4 procs.
Hello world, I've rank 3 out of 4 procs.
Hello world, I've rank 0 out of 4 procs.
```

Note: Compilation and execution can vary from system to system.

An abstract "picture" of parallel execution

Process 0 | Process 1 | ··· | Process P-1

```
#include <stdio.h>
#include <mpi.h>

int main (int nargs, char** args)
{
  int size, my_rank;
  MPI_Init (&nargs, &args);
  MPI_Comm_size (MPI_COMM_WORLD, &size);
  MPI_Comm_rank (MPI_COMM_WORLD, &my_rank);
  printf("Hello world, I've rank %d out of %d procs.\n",
         my_rank,size);
  MPI_Finalize ();
  return 0;
}
```

```
#include <stdio.h>
#include <mpi.h>

int main (int nargs, char** args)
{
  int size, my_rank;
  MPI_Init (&nargs, &args);
  MPI_Comm_size (MPI_COMM_WORLD, &size);
  MPI_Comm_rank (MPI_COMM_WORLD, &my_rank);
  printf("Hello world, I've rank %d out of %d procs.\n",
         my_rank,size);
  MPI_Finalize ();
  return 0;
}
```

```
#include <stdio.h>
#include <mpi.h>

int main (int nargs, char** args)
{
  int size, my_rank;
  MPI_Init (&nargs, &args);
  MPI_Comm_size (MPI_COMM_WORLD, &size);
  MPI_Comm_rank (MPI_COMM_WORLD, &my_rank);
  printf("Hello world, I've rank %d out of %d procs.\n",
         my_rank,size);
  MPI_Finalize ();
  return 0;
}
```

MPI message

An MPI message is simply an array of elements of a particular MPI data type.

Data types can either be standard types (pre-defined) or *derived types*, which must be defined beforehand by appropriate MPI calls.

MPI type	C type
MPI_CHAR	signed char
MPI_INT	signed int
MPI_LONG	signed long
MPI_FLOAT	float
MPI_DOUBLE	double
MPI_BYTE	N/A

Table 9.2: A selection of the standard MPI data types for C. Unsigned variants exist where applicable.

Point-to-point communication

- From one "sender" to one "receiver": point-to-point communication
- Both the sender and receiver are identified by their ranks within an MPI communicator
- Each point-to-point message can carry an additional integer label, the so-called **tag**

The simplest MPI send command

```
int MPI_Send(void *buffer, int count,
             MPI_Datatype datatype,
             int dest, int tag,
             MPI_Comm comm);
```

This *blocking* send function returns when the data has been delivered to the system and **the message buffer containing the outgoing message can be safely modified**. The message may not have been received by the destination process.

The simplest MPI receive command

```
int MPI_Recv(void *buffer, int count
             MPI_Datatype datatype,
             int source, int tag,
             MPI_Comm comm,
             MPI_Status *status);
```

- This *blocking* receive function waits until a matching message is received from the system so that the buffer contains the incoming message.
- Match of data type, source process (or `MPI_ANY_SOURCE`), message tag (or `MPI_ANY_TAG`).
- Receiving fewer `datatype` elements than `count` is ok, but receiving more elements is an error.

MPI_Status

The `MPI_Recv` function has an additional output argument: the `status` object, which can be used to determine "unknown" parameters (if any).

For example, the source or tag of a received message may not be known if `MPI_ANY_SOURCE` or `MPI_ANY_TAG` was used in the receive function.

To query the information stored in a `status` object (after the receive function is completed):

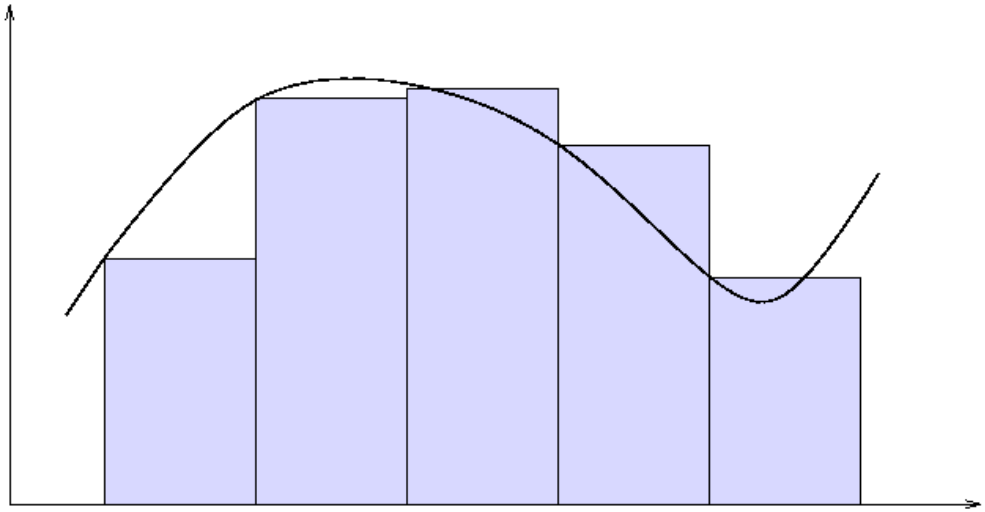
```
status.MPI_SOURCE

status.MPI_TAG

MPI_Get_count (MPI_Status *status,
               MPI_Datatype datatype,
               int *count);
```

Example of MPI parallelization; numerical integration

A quick recap of numerical integration



http://spiff.rit.edu/classes/phys317/lectures/num_integ2/num_integ2.html

The basic idea:

- Divide the integral domain into many small intervals
- Evaluate the function inside each small interval
- Sum up all the evaluated results (and then multiply with the interval width)

MPI parallelization; overall idea

- Each MPI process is assigned with a "subdomain' to work on
- Each MPI process computes the numerical integration (a partial result) over its subdomain
- Thereafter, all the subdomain results need to be summed up
- One possible strategy is to let each process (except rank 0) send its partial result to process rank 0, which sums up all the subdomain results

MPI implementation

```
int rank, size;
double a=0.0, b=1.0, mya, myb, psum;
MPI_Status status;

MPI_Init (&nargs, &args);
MPI_Comm_size (MPI_COMM_WORLD, &size);
MPI_Comm_rank (MPI_COMM_WORLD, &rank);

mya = a + rank*(b-a)/size;
myb = mya + (b-a)/size;
psum = integrate(mya,myb); // Each integrates over its "subdomain"

if (rank==0) {
  double res = psum;
  for (i=1; i<size; i++) {
    MPI_Recv(&psum,1,MPI_DOUBLE,i,0,MPI_COMM_WORLD,&status);
    res += psum;
  }
  printf("Result: %g\n", res);
}
else {
  MPI_Send(&psum,1,MPI_DOUBLE,0,0,MPI_COMM_WORLD);
}
MPI_Finalize();
```

Additional remarks

The preceding MPI example can be improved in several ways:

- Rank 0 receives in total `size-1` messages, one from each of the other processes. The order of receiving the messages is fixed, probably `not` the same as the incoming order of the messages. Using `MPI_ANY_SOURCE` (wildcard) instead of a prescribed sender rank is more appropriate.
- Rank 0 calls `MPI_Recv` after its own calculation (`integrate(mya,myb)`). If some of the other processes finish their computation earlier, communication cannot proceed, and it cannot be overlapped with computation on rank 0. (Non-blocking point-to-point communication will be more appropriate.)
- Rank 0 can easily become a "bottleneck", because it is responsible for receiving all the partial results. (Use of collective communication, such as that specifically designed for "reduction", will be more appropriate.)

“Uncertainties/risks” with MPI_Send

While MPI_Send is easy to use, one should be aware that the MPI standard allows for a considerable amount of freedom in its actual implementation.

- Internally it may work completely **synchronously**, meaning that the call can not return until a message transfer has at least started after a “handshake” with the receiver.
- However, it may also copy the message to an intermediate system buffer and return right away, leaving the “handshake” and data transmission to another mechanism, like a background thread.
- It may even change its behavior depending on any explicit or hidden parameters.

One scenario for deadlock

Deadlocks may occur if the possible synchronousness of MPI_Send is not taken into account. A typical communication pattern where this may become crucial is a "ring shift":

```
int rank, size, left, right,in_buf[N], out_buf[N];
MPI_Status status;

// .......

MPI_Comm_size (MPI_COMM_WORLD, &size);
MPI_Comm_rank (MPI_COMM_WORLD, &rank);

right = rank==size-1 ? 0 : rank+1;
left = rank==0 ? size-1 : rank-1;

MPI_Send(out_buf,N,MPI_INT,right,0,MPI_COMM_WORLD);
MPI_Recv(in_buf,N,MPI_INT,left,0,MPI_COMM_WORLD,&status);
```

Depiction of the "ring shift" pattern

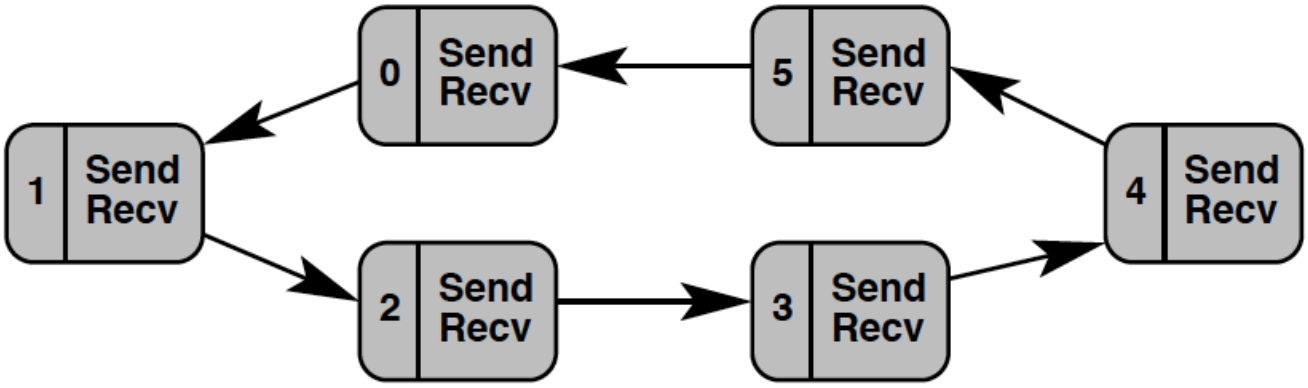


Figure 9.1: A ring shift communication pattern. If sends and receives are performed in the order shown, a deadlock can occur because `MPI_Send()` may be synchronous.

- If MPI_Send is **synchronous** (and there is no system buffering), all processes call it first and then wait forever for a matching receive to be posted—deadlock.
- However, the ring shift **may** run without problems if the messages are sufficiently short. In fact, most MPI implementations provide a (small) internal buffer for short messages and switch to synchronous mode when the buffer is full or too small.

One simple solution

```
int rank, size, left, right,in_buf[N], out_buf[N];
MPI_Status status;

// .......

MPI_Comm_size (MPI_COMM_WORLD, &size);
MPI_Comm_rank (MPI_COMM_WORLD, &rank);

right = rank==size-1 ? 0 : rank+1;
left = rank==0 ? size-1 : rank-1;

if (rank%2) {
  MPI_Recv(in_buf,N,MPI_INT,left,0,MPI_COMM_WORLD,&status);
  MPI_Send(out_buf,N,MPI_INT,right,0,MPI_COMM_WORLD);
}
else {
  MPI_Send(out_buf,N,MPI_INT,right,0,MPI_COMM_WORLD);
  MPI_Recv(in_buf,N,MPI_INT,left,0,MPI_COMM_WORLD,&status);
}
```

Depiction of the simple solution

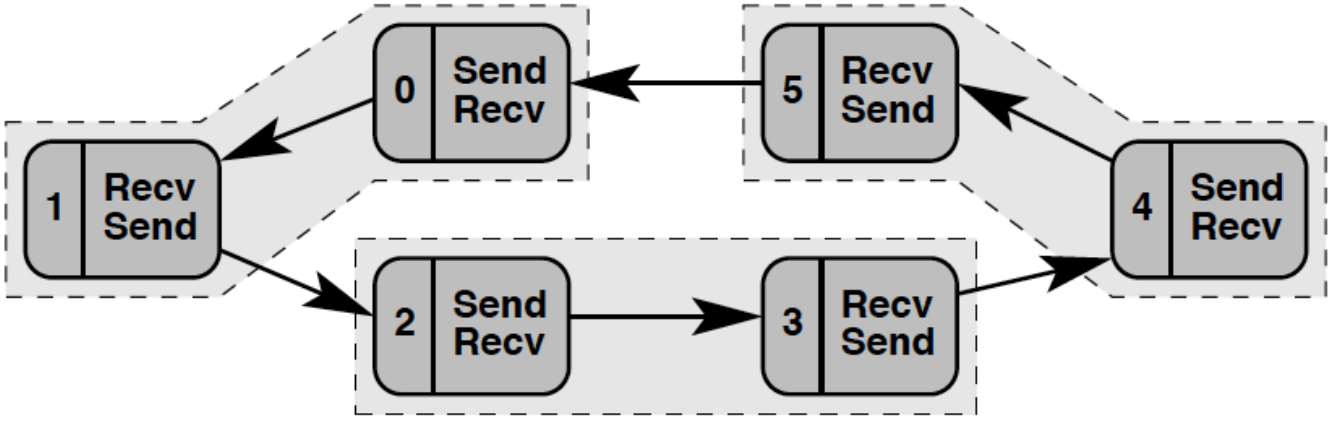


Figure 9.2: A possible solution for the deadlock problem with the ring shift: By changing the order of `MPI_Send()` and `MPI_Recv()` on all odd-numbered ranks, pairs of processes can communicate without deadlocks because there is now a matching receive for every send operation (dashed boxes).

Special MPI functions for "ring shifts"

```
int MPI_Sendrecv(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
                 int dest, int sendtag,
                 void *recvbuf, int recvcount, MPI_Datatype recvtype,
                 int source, int recvtag,
                 MPI_Comm comm, MPI_Status *status)
```

```
int MPI_Sendrecv_replace(void *buf, int count, MPI_Datatype datatype,
                         int dest, int sendtag, int source, int recvtag,
                         MPI_Comm comm, MPI_Status *status)
```

Both routines use blocking communication, are guaranteed to not be subject to the deadlock effects that may occur with separate send and receive.

Collective communication in MPI

Recall the numerical integration example: summing up the partial sums is a *reduction* operation.

MPI has mechanisms that make reductions much simpler and in most cases more efficient than looping over all ranks and collecting results.

Since a reduction is a procedure involves all ranks in a communicator, it belongs to the so-called **collective** communication operations in MPI.

A collective MPI routine must be called by all ranks in a communicator.

MPI_Barrier

MPI_Barrier is the simplest collective in MPI, it does not actually perform any real data transfer.

```
int MPI_Barrier(MPI_Comm comm)
```

The barrier *synchronizes* the members of the communicator, i.e., all processes must call it before they are allowed to return to the user code.

Don't over-use MPI_Barrier! There are other MPI routines that allow for implicit or explicit synchronization with finer control.

MPI_Bcast

```
int MPI_Bcast( void *buffer, int count, MPI_Datatype datatype, int root,
               MPI_Comm comm )
```

MPI_Bcast sends a message from one process (the "root") to all others in the communicator.

The **buffer** argument to MPI_Bcast() is a send buffer on the root and a receive buffer on any other process.

Depiction of MPI_Bcast

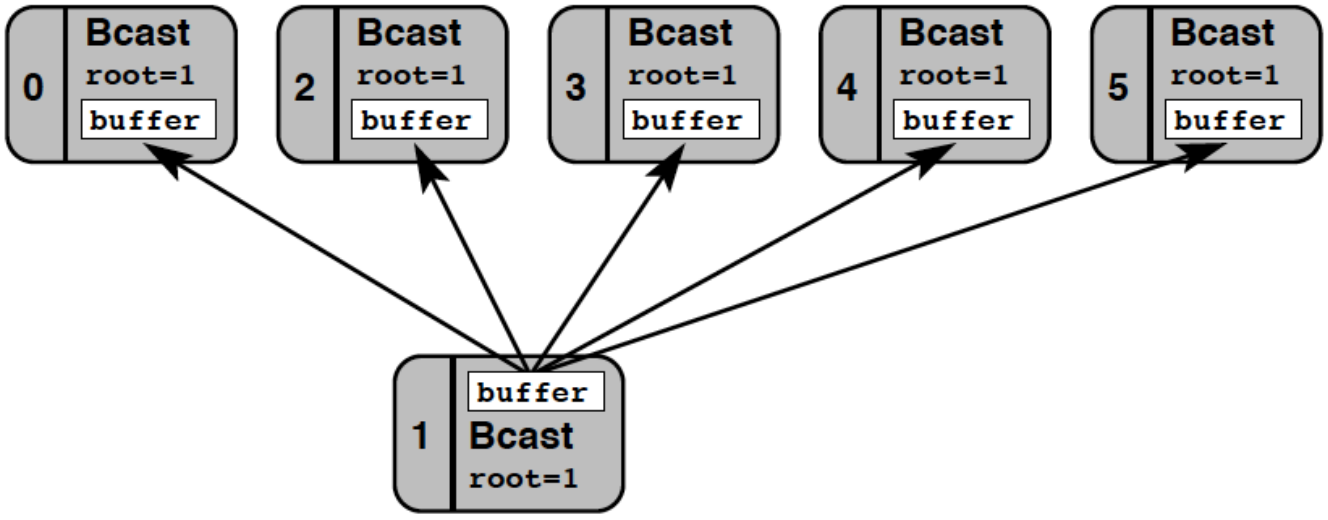


Figure 9.3: An MPI broadcast: The "root" process (rank 1 in this example) sends the same message to all others. Every rank in the communicator must call `MPI_Bcast()` with the same `root` argument.

MPI_Gather & MPI_Scatter

Examples of more advanced MPI collective calls that are concerned with global data distribution:

- MPI_Gather() collects the send buffer contents of all processes and concatenates them in rank order into the receive buffer of the root process
- MPI_Scatter() does the reverse, distributing equal-sized chunks of the root's send buffer
- MPI_Gatherv() and MPI_Scatterv() support arbitrary per-rank chunk sizes

MPI_Gather & MPI_Scatter syntax

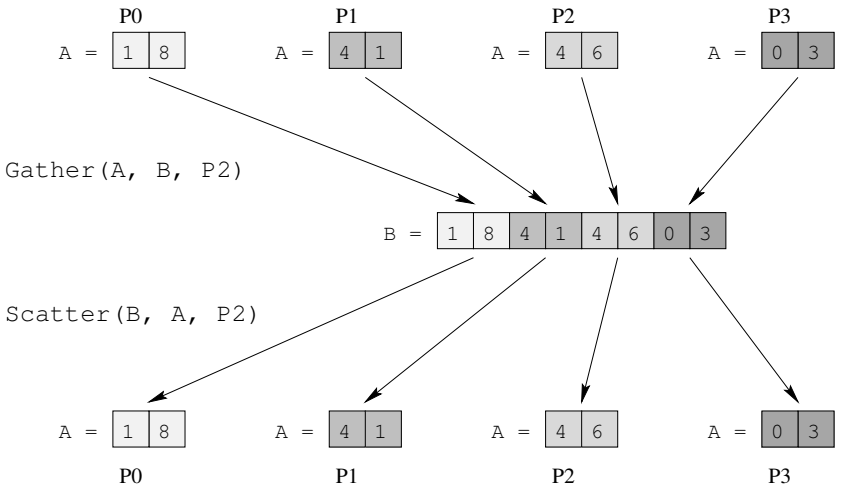
```
int MPI_Gather(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
               void *recvbuf, int recvcount, MPI_Datatype recvtype,
               int root, MPI_Comm comm)
```

```
int MPI_Scatter(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
               void *recvbuf, int recvcount, MPI_Datatype recvtype, int root,
               MPI_Comm comm)
```

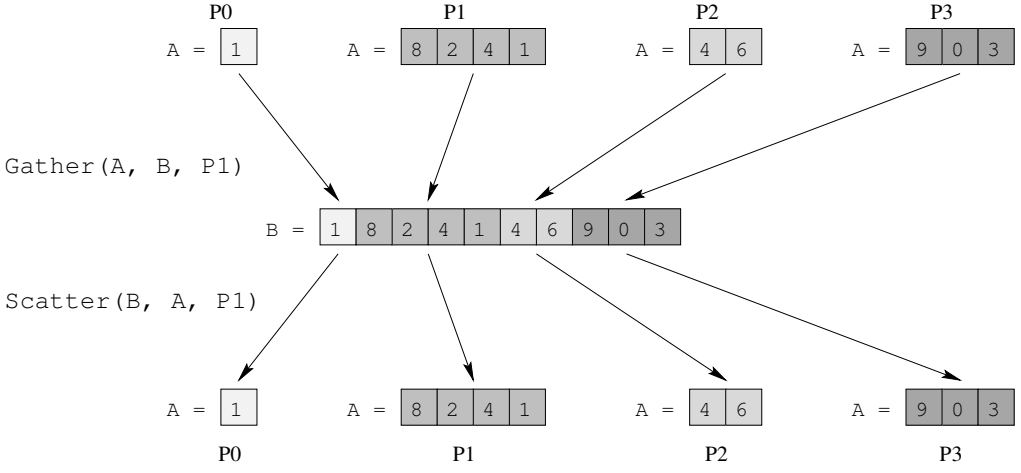
The syntax of MPI_Gatherv() and MPI_Scatterv() is more complex, not listed here.

Depiction of MPI_Gather & MPI_Scatter

(a) Equal-Size Gather and Scatter Operations



(b) Unequal-Size Gather and Scatter Operations



MPI_Reduce

```
int MPI_Reduce(const void *sendbuf, void *recvbuf,
               int count, MPI_Datatype datatype,
               MPI_Op op, int root, MPI_Comm comm)
```

MPI_Reduce() combines the contents of the sendbuf array on all processes, element-wise, using an operator encoded by the op argument, and stores the result in recvbuf on root

Examples of predefined operators: MPI_MAX, MPI_MIN, MPI_SUM and MPI_PROD

Depiction of MPI_Reduce

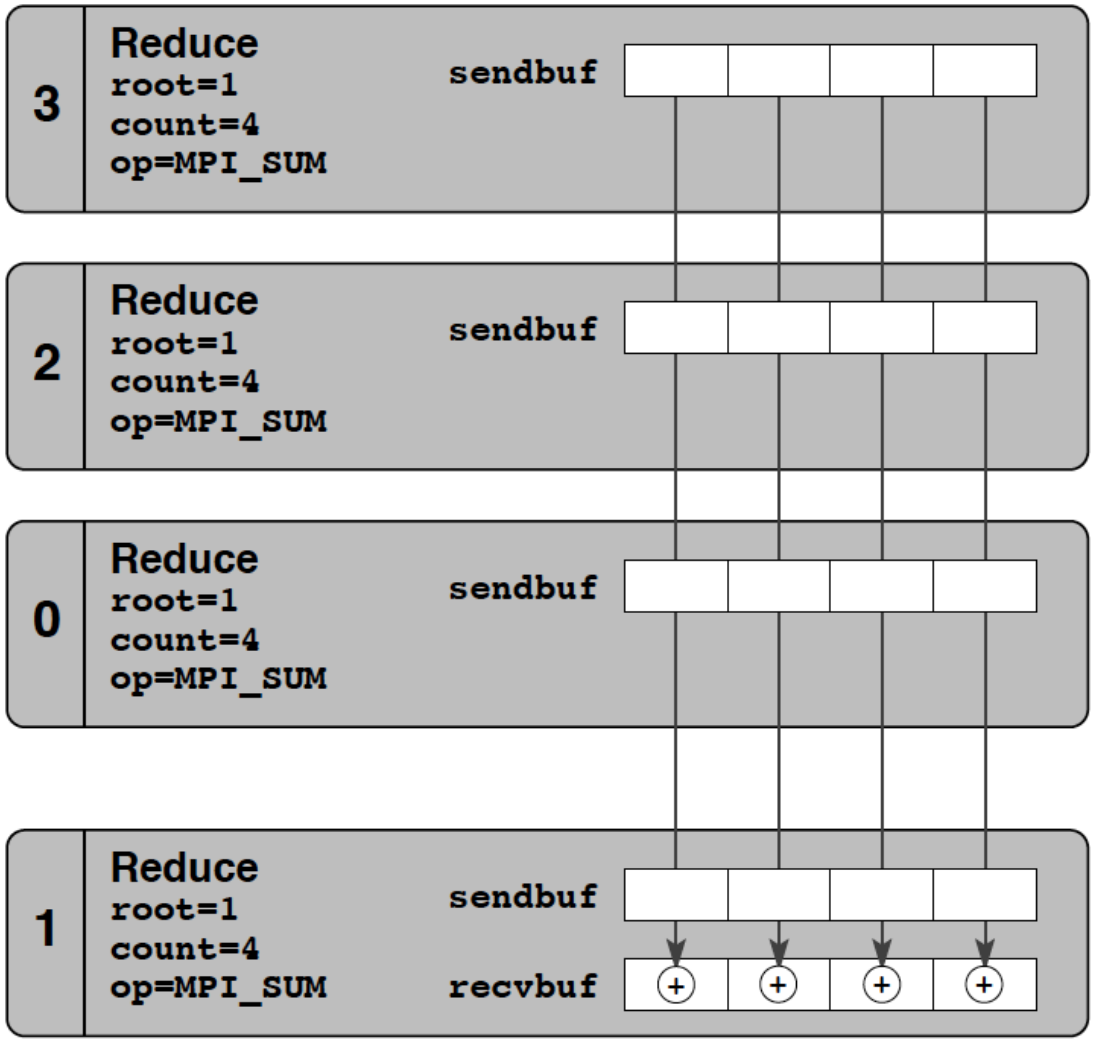


Figure 9.4: A reduction on an array of length `count` (a sum in this example) is performed by `MPI_Reduce()`. Every process must provide a send buffer. The receive buffer argument is only used on the root process. The local copy on root can be prevented by specifying `MPI_IN_PLACE` instead of a send buffer address.

Rewrite of the "numerical integration" example

```
int rank, size;
double a=0.0, b=1.0, mya, myb, psum, res=0.;
MPI_Status status;

MPI_Init (&nargs, &args);
MPI_Comm_size (MPI_COMM_WORLD, &size);
MPI_Comm_rank (MPI_COMM_WORLD, &rank);

mya = a + rank*(b-a)/size;
myb = mya + (b-a)/size;
psum = integrate(mya,myb); // Each integrates over its "subdomain"

MPI_Reduce(&psum,&res,1,MPI_DOUBLE,MPI_SUM,0,MPI_COMM_WORLD);
if (rank==0)
  printf("Result: %g\n", res);

MPI_Finalize();
```

MPI_Allreduce

```
int MPI_Allreduce(const void *sendbuf, void *recvbuf, int count,
                  MPI_Datatype datatype, MPI_Op op, MPI_Comm comm)
```

The result of the reduction operation (`recvbuf`) is available on all MPI ranks.

Nonblocking point-to-point communication

Point-to-point communication can be performed with nonblocking semantics.

- A nonblocking point-to-point call merely **initiates** a message transmission and returns very quickly to the user code.
- In an efficient implementation, waiting for data to arrive and the actual data transfer occur in the background, leaving resources free for computation.
- In other words, nonblocking MPI is a way in which communication may be overlapped with computation if implemented efficiently.
- The message buffer **must not be used** as long as the user program has not been notified that it is safe to do so (which can be checked by suitable MPI calls).
- Nonblocking and blocking MPI calls are mutually compatible, i.e., a message sent via a blocking send can be matched by a nonblocking receive.

MPI_Isend

```
int MPI_Isend(const void *buf, int count, MPI_Datatype datatype,
              int dest, int tag,
              MPI_Comm comm, MPI_Request *request)
```

- As opposed to the blocking send (MPI_Send()), MPI_Isend() has an additional output argument, the *request* handle (of type `struct MPI_Request`).
- It serves as an identifier by which the program can later refer to the "pending" communication request.

MPI_Irecv

```
int MPI_Irecv(void *buf, int count, MPI_Datatype datatype,
              int source, int tag,
              MPI_Comm comm, MPI_Request * request)
```

- The `status` object known from MPI_Recv() is **not** needed for MPI_Irecv.

MPI_Test & MPI_Wait

Checking a pending communication for completion can be done via the MPI_Test() and MPI_Wait() functions. The former only tests for completion and returns a flag, while the latter blocks until the buffer can be used.

```
int MPI_Test(MPI_Request *request, int *flag, MPI_Status *status)
```

```
int MPI_Wait(MPI_Request *request, MPI_Status *status)
```

The status object contains useful information only if the pending communication is a completed receive (i.e., in the case of MPI_Test() the value of flag must be true).

MPI_Waitall

In the case of multiple non-blocking communication operations (multiple requests pending), it is more convenient to use the MPI_Waitall function:

```
int MPI_Waitall(int count, MPI_Request array_of_requests[],
                MPI_Status array_of_statuses[])
```

The "numerical integration" example, yet again

```
// ...
double *tmp;
MPI_Request *request_array;
MPI_Status *status_array;

if (rank==0) {
  // allocate arrays of tmp, request_array & status_array...
  for (i=1; i<size; i++)
    MPI_Irecv(&tmp[i-1],1,MPI_DOUBLE,i,0,MPI_COMM_WORLD,&request_array[i-1]);
}

mya = a + rank*(b-a)/size;
myb = mya + (b-a)/size;
psum = integrate(mya,myb); // Each integrates over its "subdomain"

if (rank==0) {
  double res = psum;
  MPI_Waitall (size-1,request_array,status_array);
  for (i=1; i<size; i++)
    res += tmp[i-1];
  printf("Result: %g\n", res);
}
else {
  MPI_Send(&psum,1,MPI_DOUBLE,0,0,MPI_COMM_WORLD);
}
```

Benefits of nonblocking communication

- Nonblocking communication provides an obvious way to overlap communication, i.e., overhead, with useful work.
- The possible performance advantage, however, depends on many factors, and may even be nonexistent.
- But even if there is no real overlap, multiple outstanding nonblocking requests may improve performance because the MPI library can decide which of them gets serviced first.
- Nonblocking communication helps to avoid deadlock.

A short summary

	Point-to-point	Collective

Blocking	`MPI_Send()`
	`MPI_Ssend()`
	`MPI_Bsend()`
	`MPI_Recv()`
		`MPI_Barrier()`
		`MPI_Bcast()`
		`MPI_Scatter()`/
		`MPI_Gather()`
		`MPI_Reduce()`
		`MPI_Reduce_scatter()`
		`MPI_Allreduce()`

Nonblocking	`MPI_Isend()`
	`MPI_Irecv()`
	`MPI_Wait()`/`MPI_Test()`
	`MPI_Waitany()`/
	`MPI_Testany()`
	`MPI_Waitsome()`/
	`MPI_Testsome()`
	`MPI_Waitall()`/
	`MPI_Testall()`
		N/A

Table 9.3: MPI's communication modes and a nonexhaustive overview of the corresponding subroutines.

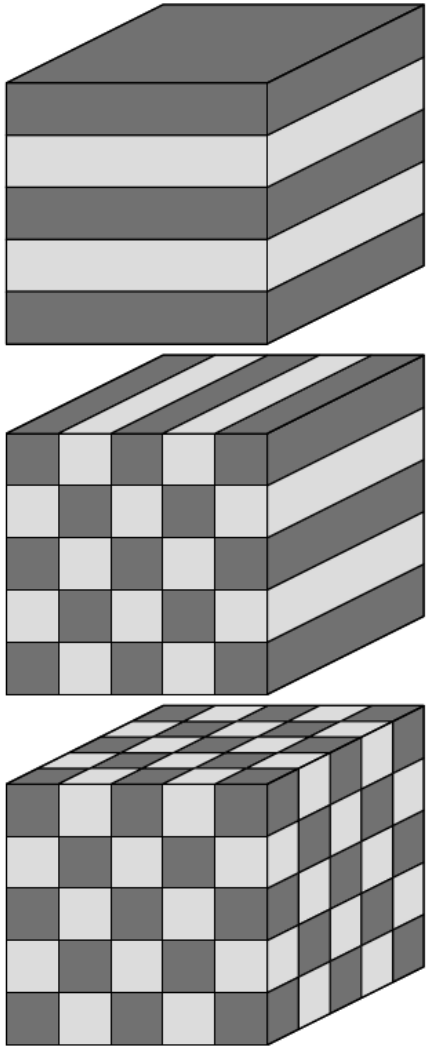
Reducing communication overhead

Domain decomposition

Domain decomposition is one of the most important strategies for data parallelism:

- We want to understand the impact of domain decomposition on communication overhead.
- We want to avoid “wrong” decompositions.

Example: decomposing a 3D domain



"Slabs"

$$c_{1d}(L,N) = L \cdot L \cdot w \cdot 2 = 2wL^2$$

"Poles"

$$c_{2d}(L,N) = L \cdot \frac{L}{\sqrt{N}} \cdot w \cdot (2+2) = 4wL^2N^{-1/2}$$

"Cubes"

$$c_{3d}(L,N) = \frac{L}{\sqrt[3]{N}} \cdot \frac{L}{\sqrt[3]{N}} \cdot w \cdot (2+2+2) = 6wL^2N^{-2/3}$$

Figure 10.9: 3D domain decomposition of a cubic domain of size L^3 (strong scaling) and periodic boundary conditions: Per-process communication volume $c(L,N)$ for a single-site data volume w (in bytes) on N processes when cutting in one (top), two (middle), or all three (bottom) dimensions.

The amount of communication depends on decomposition

Example: The cubic domain is of size L^3, which has periodic boundary conditions, is decomposed into N subdomains:

- The domain cuts can be performed in one, two, or all three dimensions;
- The total number of elements that must be communicated by one process with its neighbors, $c(L, N)$, can be found as simple formulas;

Decomposing as "slabs"

- 1D cuts

$$c_{1\mathrm{d}}(L, N) = L \cdot L \cdot w \cdot 2 = 2wL^2$$

- Should not be used in general because they incur a much larger and, more importantly, N-independent overhead (doesn't decrease when N grows).
- A constant cost of communication per subdomain will greatly harm strong scalability, because performance saturates at a lower level, determined by the message size (i.e., the slab area) instead of the latency.

Decomposing as "poles"

- 2D cuts

$$c_{2\mathrm{d}}(L, N) = L \cdot \frac{L}{\sqrt{N}} \cdot w \cdot (2 + 2) = 4w \frac{L^2}{\sqrt{N}}$$

- The volume of communication per process decreases with increasing N;
- However, the surface-to-volume ratio will grow with N.

Decomposing as "cubes"

- 3D cuts

$$c_{3\mathrm{d}}(L,N) = \frac{L}{\sqrt[3]{N}} \cdot \frac{L}{\sqrt[3]{N}} \cdot w \cdot (2+2+2) = 6w\frac{L^2}{\left(\sqrt[3]{N}\right)^2}$$

- Smaller communication volume per process than 2D decomposition (as long as $N > \left(\frac{3}{2}\right)^6$);
- Still has the same problems as with 2D decomposition;
- If latency dominates the overhead, "optimal" 3D decomposition may even be counterproductive because of the six neighbors for each domain.
- **Another factor for consideration:** cost of message packing & unpacking

Aggregating messages

- If a parallel algorithm requires transmission of a lot of small messages between processes, communication becomes latency-bound because each message incurs latency.
- Hence, small messages should be aggregated into contiguous buffers and sent in larger chunks so that the "latency penalty" is only paid once.
- This advantage pertains to point-to-point and collective communication alike.

Example of message aggregation

- A typical case for message aggregation comes up when separate, i.e., noncontiguous data items must be transferred between processes.
- Example: x-direction communication related to 2D domain decomposition—need to send and receive "columns".

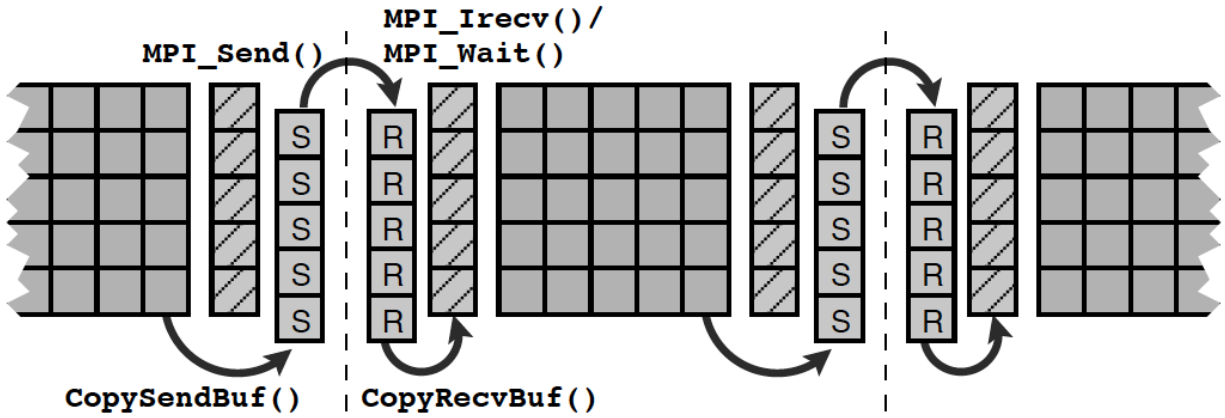


Figure 9.9: Halo communication for the Jacobi solver (illustrated in two dimensions here) along one of the coordinate directions. Hatched cells are ghost layers, and cells labeled "R" ("S") belong to the intermediate receive (send) buffer. The latter is being reused for all other directions. Note that halos are always provided for the grid that gets read (not written) in the upcoming sweep. Fixed boundary cells are omitted for clarity.

Asynchronous communication

- A further chance for increasing efficiency of parallel programs is overlapping communication and computation.
- Nonblocking point-to-point communication offers an opportunity for achieving this.
- Example of 2D domain decomposition for 2D Jacobi computation:
  1. Compute only the rows `phi_new[1][...]`, `phi_new[YMAX-2][...]` and columns `phi_new[...][1]`, `phi_new[...][XMAX-2]`
  2. Initiate sending/receiving of all the messages by `MPI_Isend` and `MPI_Irecv`
  3. Compute the "interior" of the interior points (hopefully MPI messages are being transferred at the same time)
  4. Finish all MPI calls by `MPI_Wait`

Difference between nonblocking and asynchornous communciation

Very important: One must strictly differentiate between non-blocking and truly asynchronous communication. Nonblocking semantics, according to the MPI standard, merely implies that the message buffer cannot be used after the call has returned from the MPI library; **In other words**, it is entirely up to the implementation whether data transfer, i.e., MPI progress, takes place while user code is being executed outside MPI.

Exercises

Exercise 1

What will be the result if executed using 8 MPI processes?

```
#include <mpi.h>
#include <stdio.h>
int main (int nargs, char **args)
{
  int own_value, in_value, out_value, i;
  int rank, size;
  int send_to, recv_from;
  MPI_Status recv_status;
  MPI_Request recv_req;

  MPI_Init (&nargs, &args);
  MPI_Comm_size (MPI_COMM_WORLD, &size);
  MPI_Comm_rank (MPI_COMM_WORLD, &rank);

  own_value = rank;
  recv_from = (rank+2)%size;
  send_to = (rank-2+size)%size;
  out_value = own_value;

  for (i=0; i<(size/2)-1; i++) {
    MPI_Irecv(&in_value, 1, MPI_INT, recv_from, 0, MPI_COMM_WORLD, &recv_req);
    MPI_Send (&out_value, 1, MPI_INT, send_to, 0, MPI_COMM_WORLD);
    MPI_Wait (&recv_req, &recv_status);
    own_value += in_value;
    out_value = in_value;
  }

  printf("On rank <%d>, own_value=%d\n",rank,own_value);
  MPI_Finalize ();
  return 0;
}
```

Exercise 2

Implement an MPI parallelized function that has the purpose of computing a matrix-vector multiplication $y = Ax$:

```
MatVecProd_MPI (int n, double **A, double *x, double *y)
```

This function is to be called by all the MPI processes (of `MPI_COMM_WORLD`), where the arrays `A`, `x` and `y` are only pre-allocated with numerical values on MPI process 0. Specifically, on MPI process 0, the 2D array `A` is of dimension $n \times n$, and the 1D arrays `x` and `y` are both of length n. All the other MPI processes have `NULL` as the input arguments `A`, `x` and `y` when calling the `MatVecProd_MPI` function.

Moreover,

Exercise 3

```
t = 0.0;
while (t < T) {
  t += dt;

  /* compute all the interior points */
  for (i=1; i<=M; i++)
    up[i] = 2*u[i]-um[i]
           +((dt*dt)/(dx*dx))*(u[i-1]-2*u[i]+u[i+1]);

  up[0] = value_of_left_BC(t);   // left boundary point
  up[M+1] = value_of_rigt_BC(t); // right boundary point

  /* preparation for next time step: array shuffle */
  tmp = um;
  um = u;
  u = up;
  up = tmp;
}
```

Write an MPI parallelized implementation of the above code, based on 1D domain decomposition.

Exercise 4

Improve the MPI implementation from the previous exercise, so that it has the possibility of hiding the communication overhead by using non-blocking MPI commands for sending and receiving messages.

Suggested solution for Exercise 4 of Module 2

```
#pragma omp parallel
{
  // declaration of thread-private variables
  int i, istart, iend, num_threads, thread_id;

  num_threads = omp_get_num_threads();
  thread_id = omp_get_thread_num();

  // calculate start and end indices for each thread
  istart = thread_id*(n-1)/num_threads; iend = (thread_id+1)*(n-1)/num_threads;

  d[iend] = d[0];
  for (i=istart; i<iend; i++) {
    a[i] = b[i] + c[i];
    d[iend] += a[i]*c[i];
  }
#pragma omp barrier

#pragma omp single
{
  for (i=1; i<num_threads; i++)
     d[(i+1)*(n-1)/num_threads] += d[i*(n-1)/num_threads] - d[0];
}

  for (i=istart; i<iend-1; i++)  // very important for each thread to stop at i=iend-2
     d[i+1] = d[i] + a[i]*c[i];
}  // end of parallel region
```