

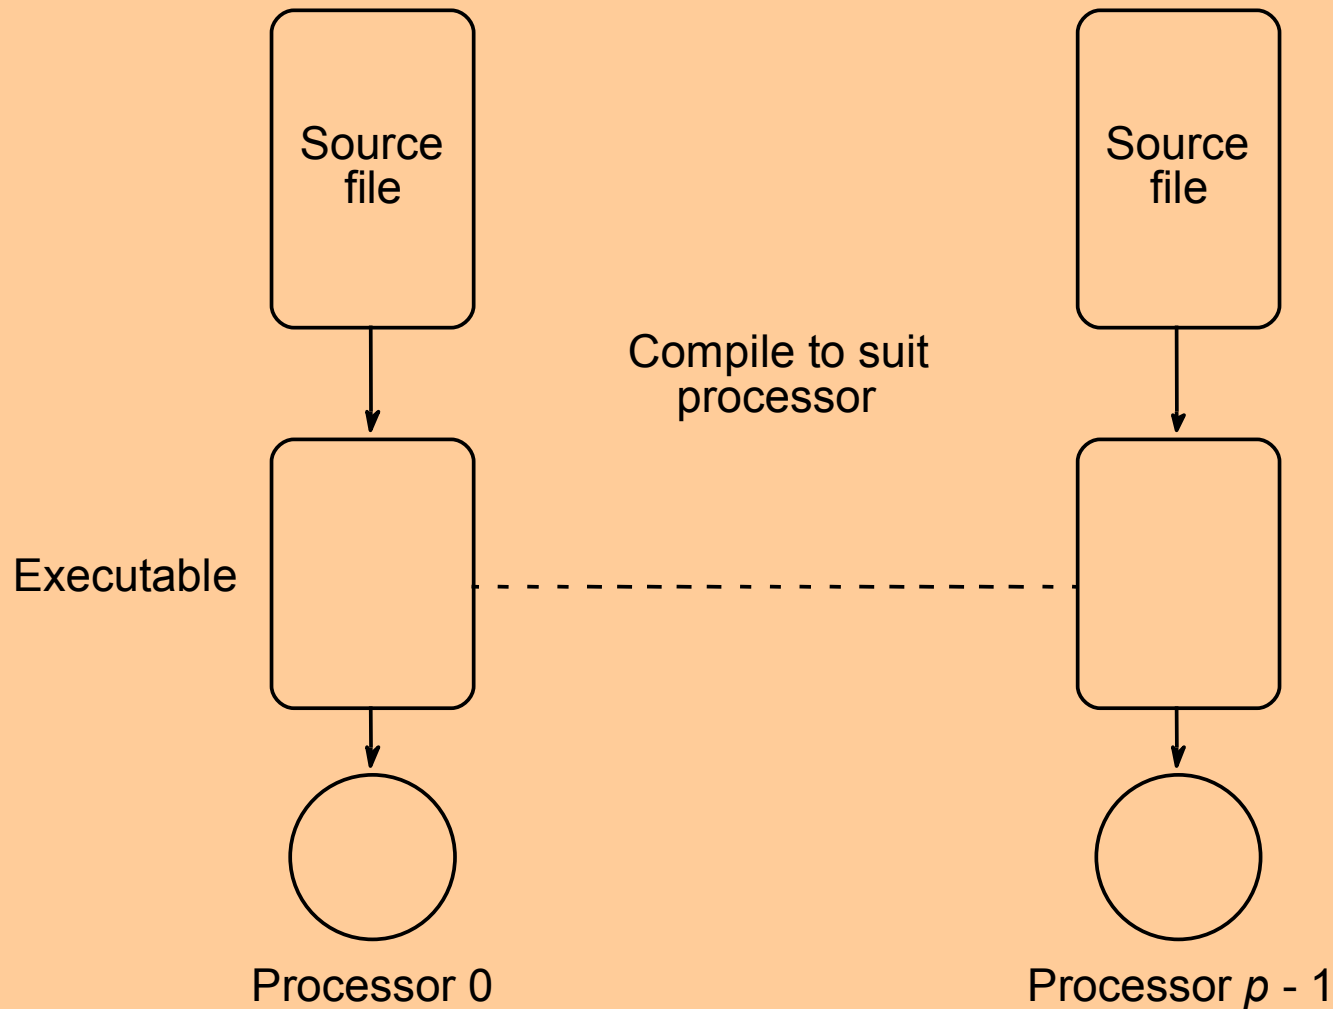
Parallel Programming with Message-Passing

Message-Passing Programming using User-level Message-Passing Libraries

Two primary mechanisms needed:

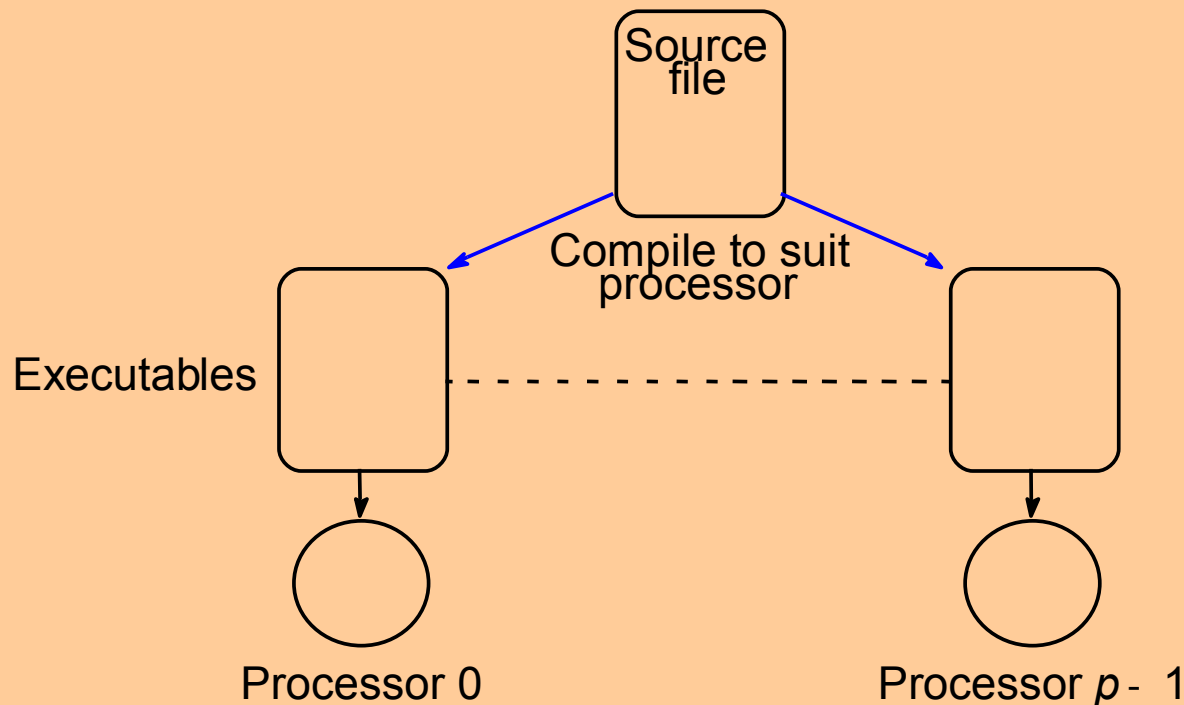
1. A method of creating separate processes for execution on different computers
2. A method of sending and receiving messages

Multiple program, multiple data (MPMD) model



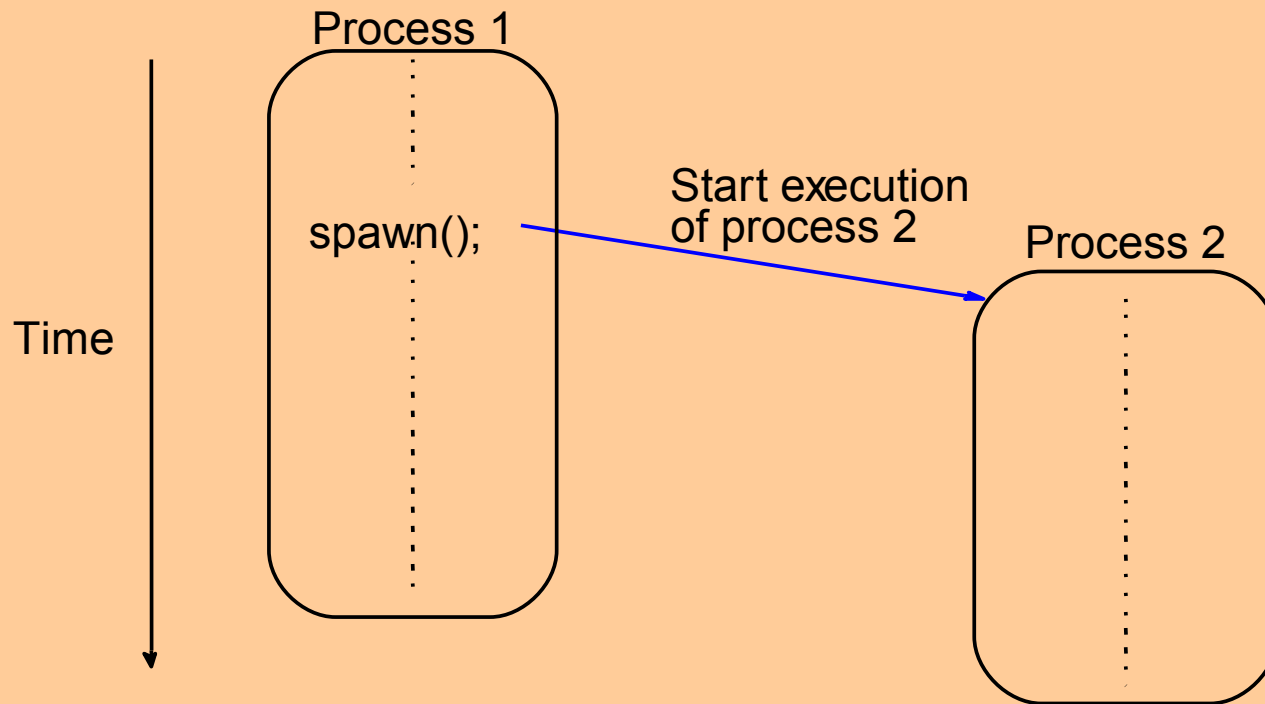
Single Program Multiple Data (SPMD) model

Different processes merged into one program. Control statements select different parts for each processor to execute. All executables started together - *static process creation*



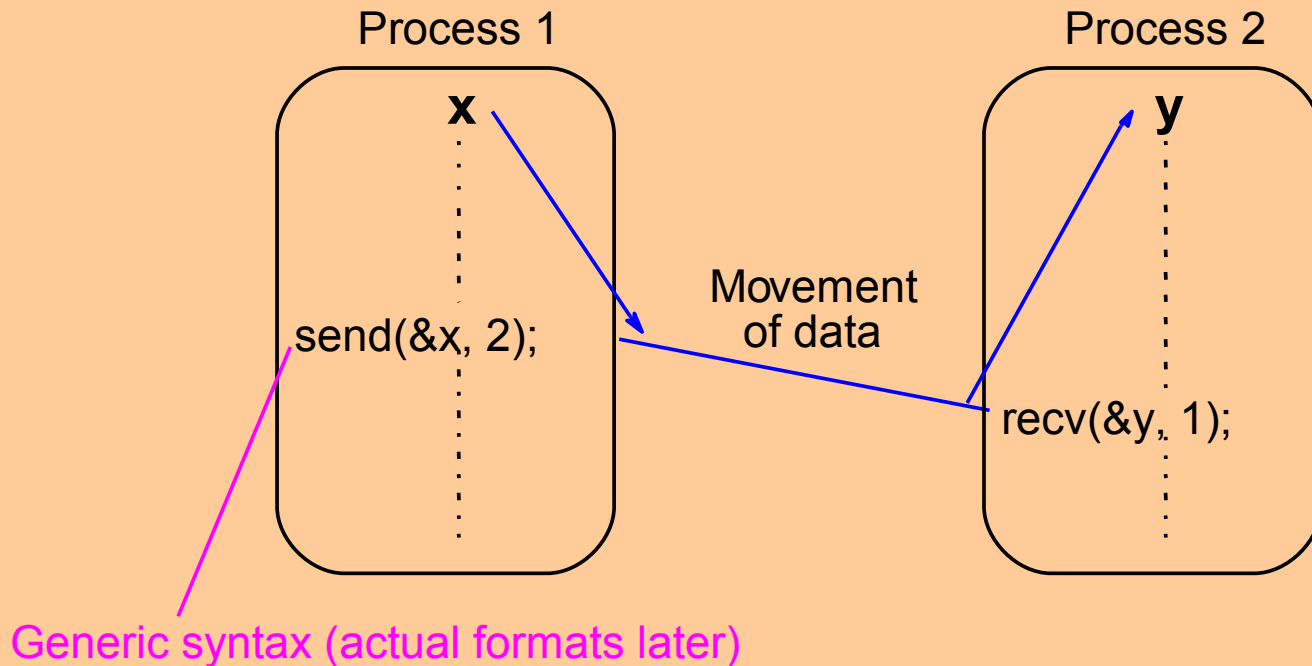
Multiple Program Multiple Data (MPMD) Model

Separate programs for each processor. One processor executes master process. Other processes started from within master process - *dynamic process creation*.



Basic “point-to-point” Send and Receive Routines

Passing a message between processes using
send() and recv() library calls:



Synchronous Message Passing

Routines that actually return when message transfer completed.

Synchronous send routine

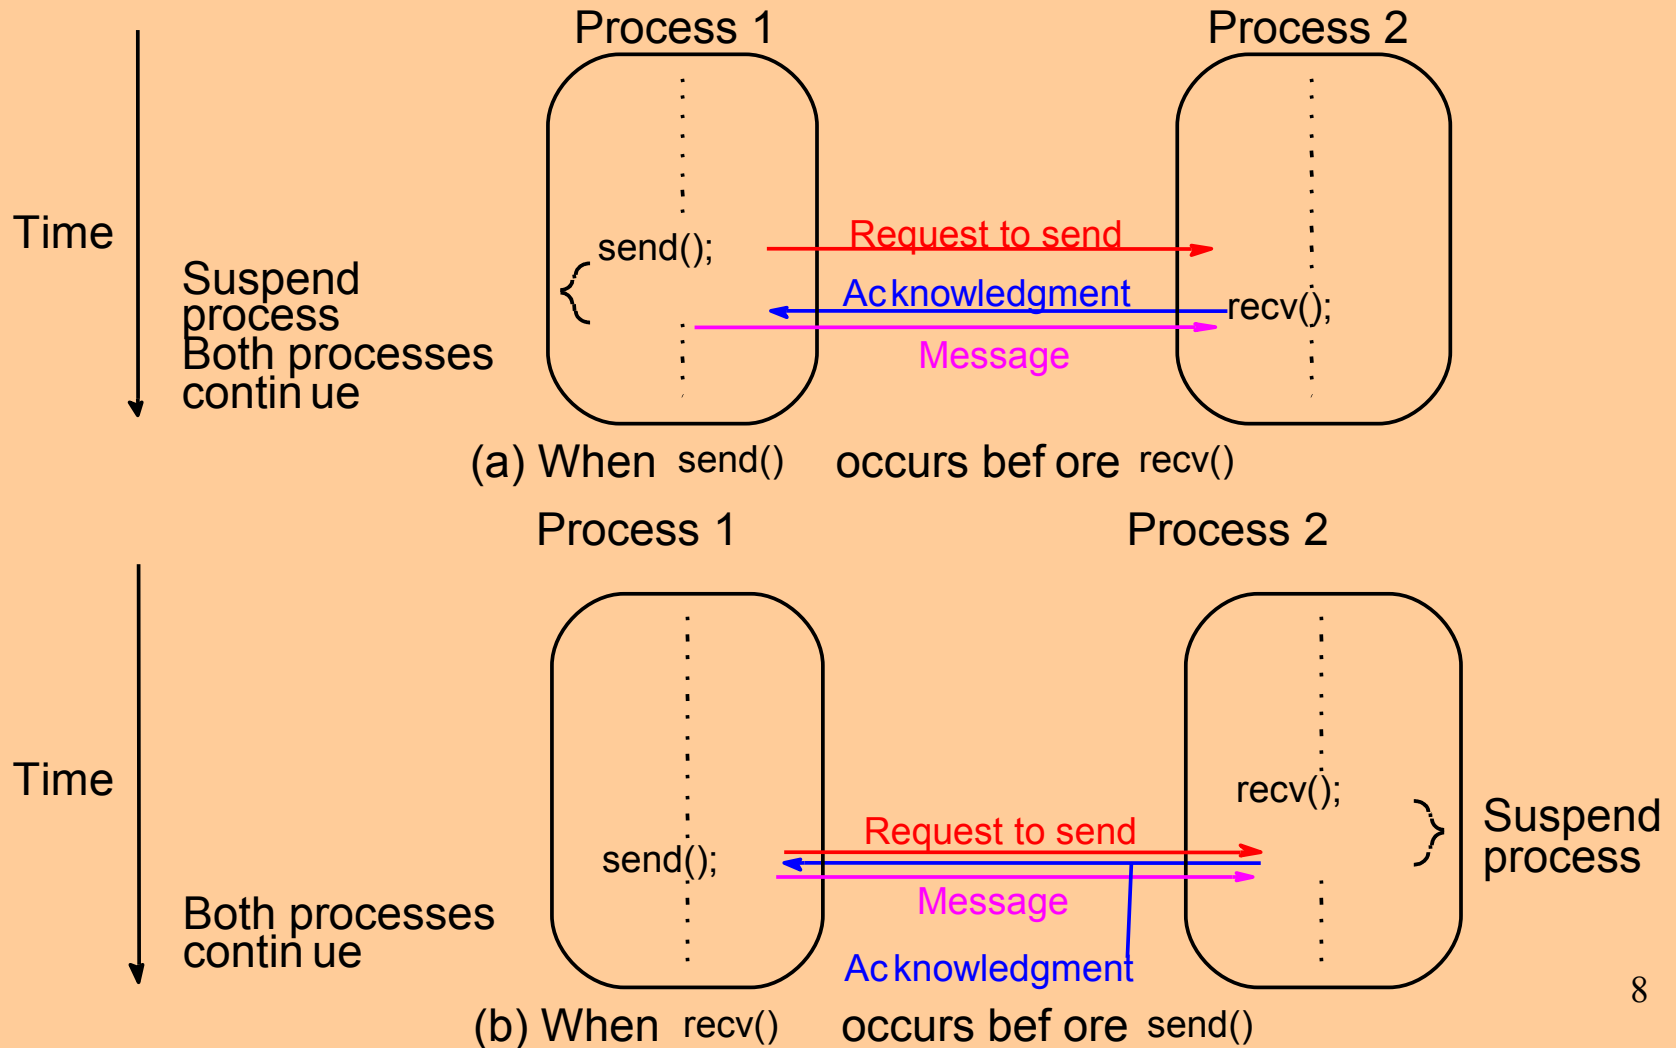
- Waits until complete message can be accepted by the receiving process before sending the message.

Synchronous receive routine

- Waits until the message it is expecting arrives.

Synchronous routines intrinsically perform two actions: They transfer data and they synchronize processes.

Synchronous send() and recv() using 3-way protocol



Asynchronous Message Passing

- Routines that do not wait for actions to complete before returning. Usually require local storage for messages.
- More than one version depending upon the actual semantics for returning.
- In general, they do not synchronize processes but allow processes to move forward sooner. Must be used with care.

MPI Definitions of Blocking and Non-Blocking

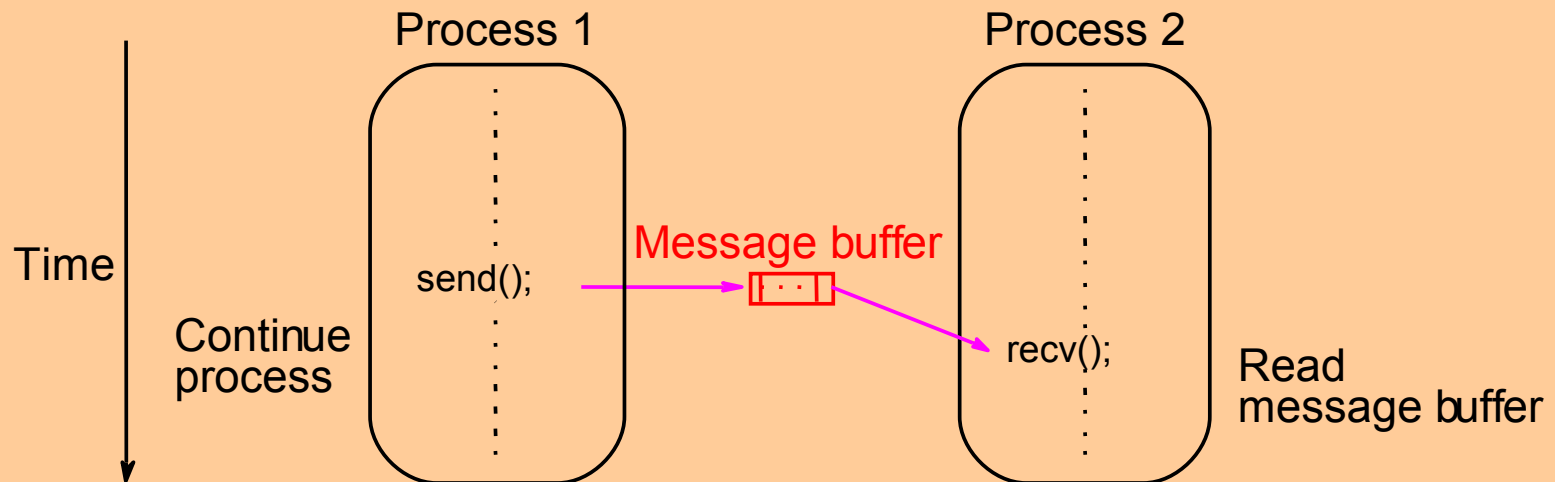
- **Blocking** - return after their local actions complete, though the message transfer may not have been completed.
- **Non-blocking** - return immediately.

Assumes that data storage used for transfer not modified by subsequent statements prior to being used for transfer, and it is left to the programmer to ensure this.

These terms may have different interpretations in other systems.

How message-passing routines return before message transfer completed

Message buffer needed between source and destination to hold message:



Asynchronous (blocking) routines changing to synchronous routines

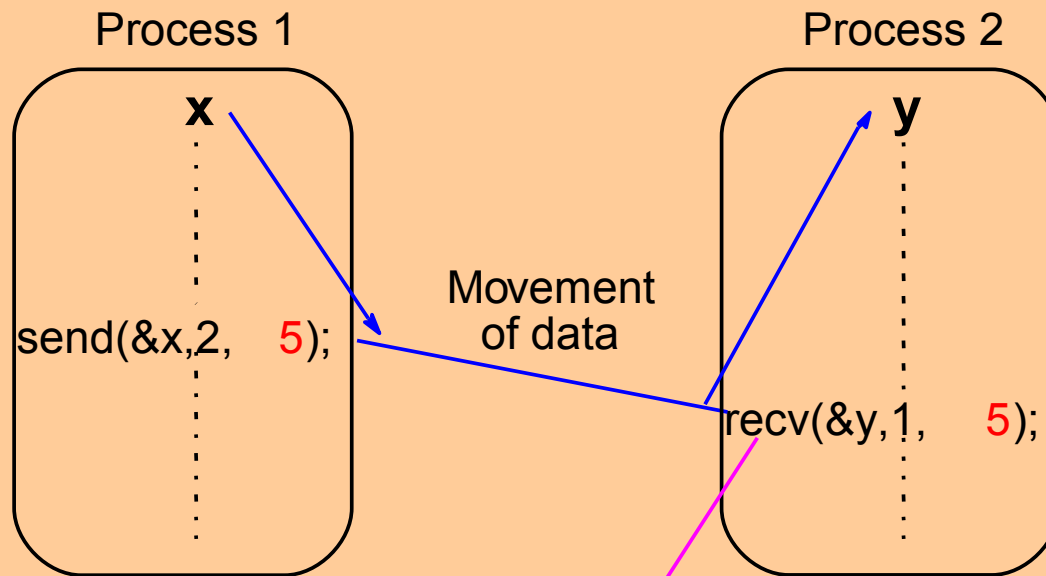
- Once local actions completed and message is safely on its way, sending process can continue with subsequent work.
- Buffers only of finite length and a point could be reached when send routine held up because all available buffer space exhausted.
- Then, send routine will wait until storage becomes re-available - i.e then routine behaves as a synchronous routine.

Message Tag

- Used to differentiate between different types of messages being sent.
- Message tag is carried within message.
- If special type matching is not required, a wild card message tag is used, so that the `recv()` will match with any `send()`.

Message Tag Example

To send a message, x , with message tag 5 from a source process, 1, to a destination process, 2, and assign to y :



Waits for a message from process 1 with a tag of 5

“Group” message passing routines (collective communication)


Have routines that send message(s) to a group of processes or receive message(s) from a group of processes

Higher efficiency than separate point-to-point routines although not absolutely necessary.

Broadcast

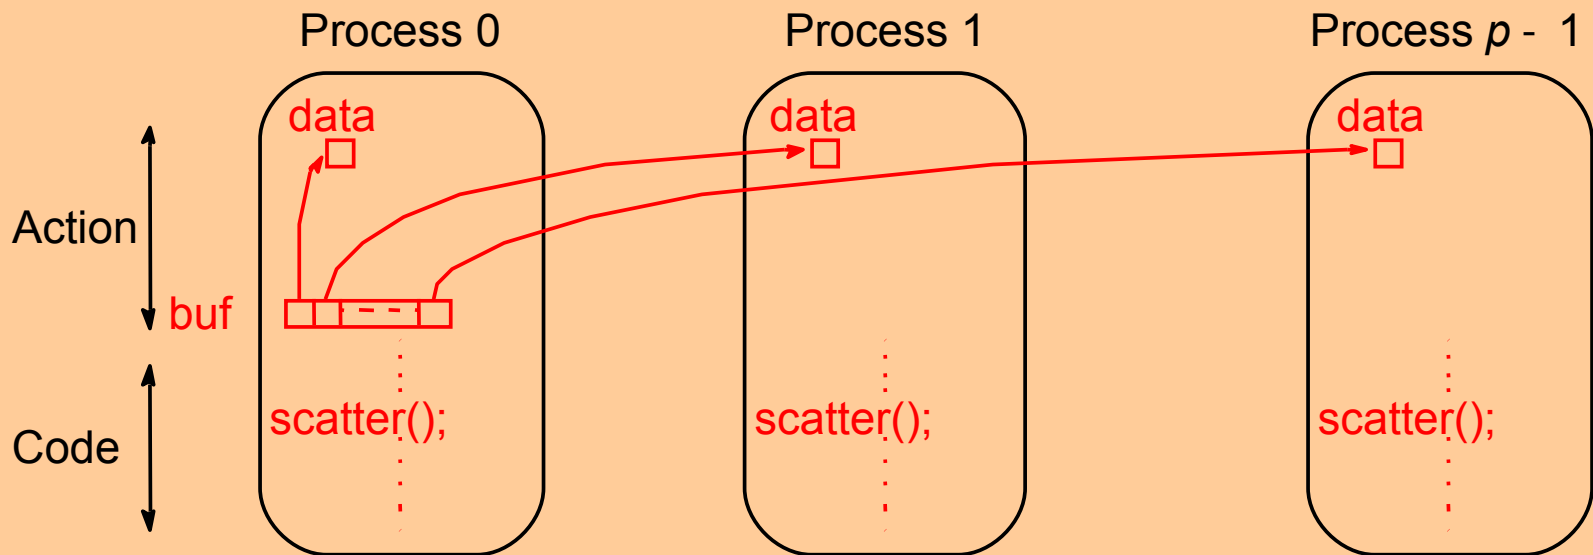
Sending same message to all processes concerned with problem.

Multicast - sending same message to defined group of processes.



Scatter


Sending each element of an array in root process to a separate process. Contents of i th location of array sent to i th process.



MPI form

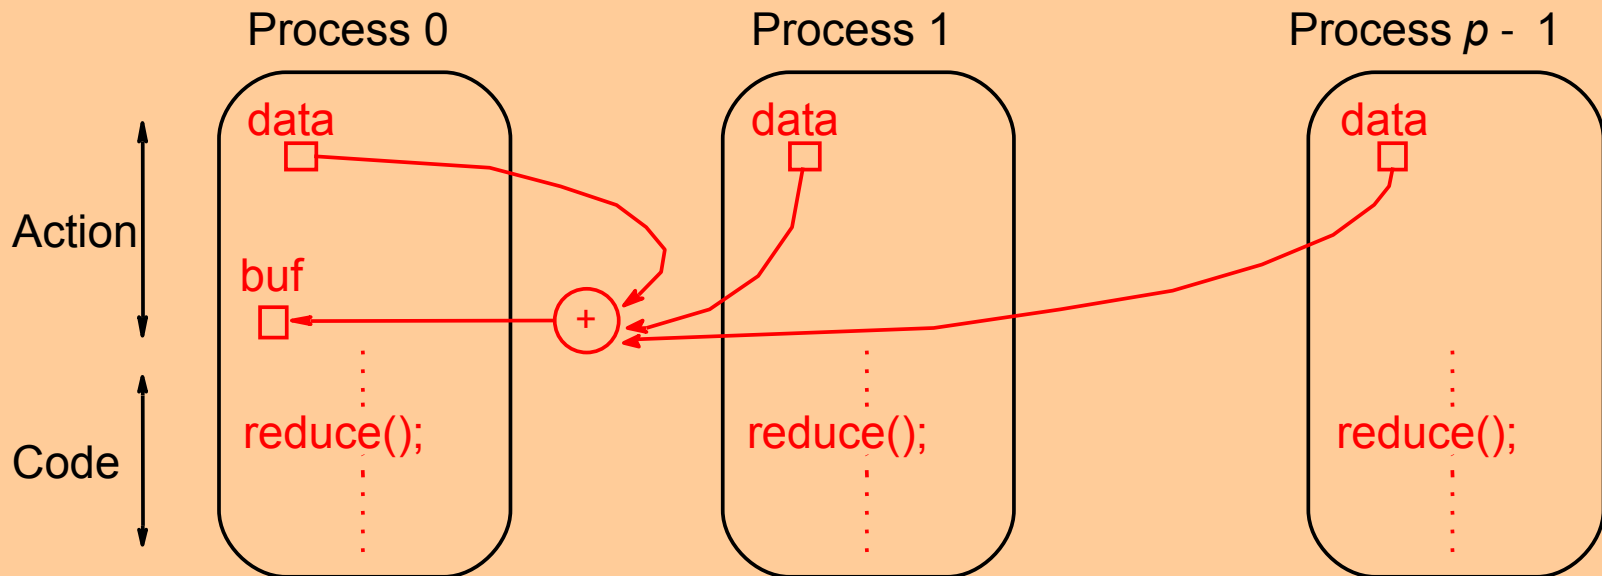
Gather

Having one process collect individual values from set of processes.



Reduce

Gather operation combined with specified arithmetic/logical operation.



MPI form

MPI

(Message Passing Interface)

- Message passing library standard developed by group of academics and industrial partners to foster more widespread use and portability.
- Defines routines, not implementation.
- Several free implementations exist.

MPI

Process Creation and Execution

- Purposely not defined - Will depend upon implementation.
- Only static process creation supported in MPI version 1. All processes must be defined prior to execution and started together.
- Originally SPMD model of computation.
- MPMD also possible with static creation - each program to be started together specified.

Communicators

- Defines scope of a communication operation.
- Processes have ranks associated with communicator.
- Initially, all processes enrolled in a “universe” called `MPI_COMM_WORLD`, and each process is given a unique rank, a number from 0 to $p - 1$, with p processes.
- Other communicators can be established for groups of processes.

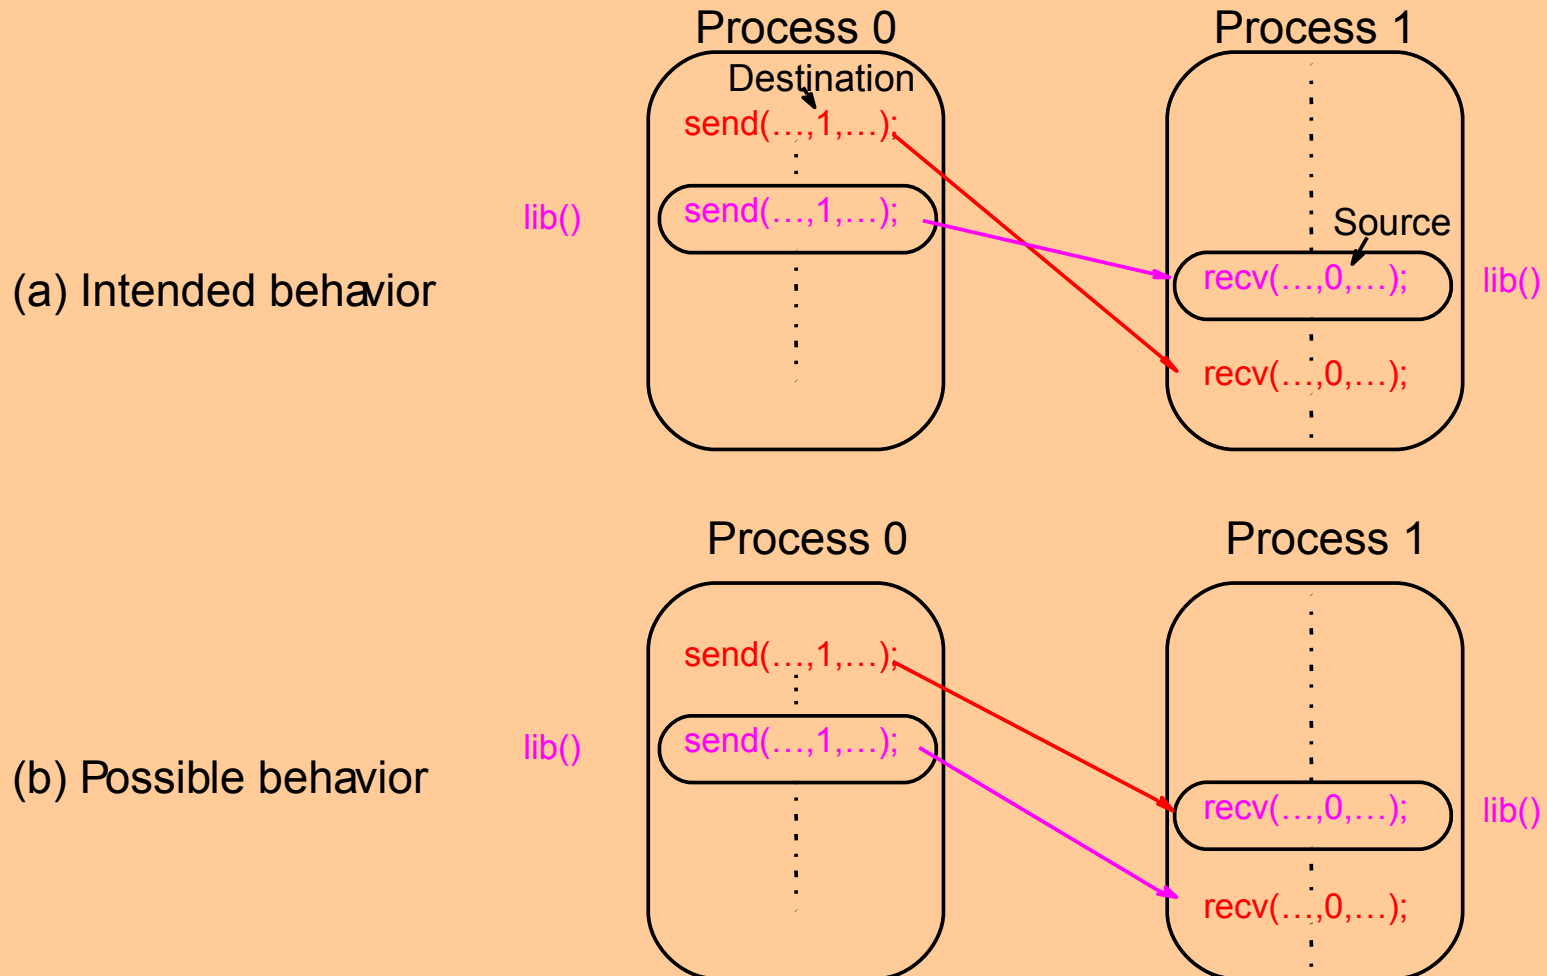
Using SPMD Computational Model

```
main (int argc, char *argv[])
{
    MPI_Init(&argc, &argv);
        .
        .
    MPI_Comm_rank(MPI_COMM_WORLD, &myrank); /*find process rank */

    if (myrank == 0)
        master();
    else
        slave();
        .
        .
    MPI_Finalize();
}
```

where master() and slave() are to be executed by master process and slave process, respectively.

Unsafe message passing - Example



MPI Solution

“Communicators”

- Defines a communication domain - a set of processes that are allowed to communicate between themselves.
- Communication domains of libraries can be separated from that of a user program.
- Used in all point-to-point and collective MPI message-passing communications.

Default Communicator

MPI_COMM_WORLD

- Exists as first communicator for all processes existing in the application.
- A set of MPI routines exists for forming communicators.
- Processes have a “rank” in a communicator.

MPI Point-to-Point Communication

- Uses send and receive routines with message tags (and communicator).
- Wild card message tags available

MPI Blocking Routines

- Return when “locally complete” - when location used to hold message can be used again or altered without affecting message being sent.
- Blocking send will send message and return - does not mean that message has been received, just that process free to move on without adversely affecting message.

Parameters of blocking send

MPI_Send(buf, count, datatype, dest, tag, comm)

Address of
send buffer

Number of items
to send

Datatype of
each item

Rank of destination
process

Message tag

Communicator

Parameters of blocking receive

MPI_Recv(buf, count, datatype, src, tag, comm, status)

Address of
receive buffer

Maximum number
of items to receive

Datatype of
each item

Rank of source
process

Message tag

Communicator

Status
after operation

Example

To send an integer x from process 0 to process 1,

```
MPI_Comm_rank(MPI_COMM_WORLD,&myrank); /* find rank */

if (myrank == 0) {
    int x;
    MPI_Send(&x, 1, MPI_INT, 1, msgtag, MPI_COMM_WORLD);
} else if (myrank == 1) {
    int x;
    MPI_Recv(&x, 1, MPI_INT,
0,msgtag,MPI_COMM_WORLD,status);
}
```

MPI Nonblocking Routines

- **Nonblocking send** - `MPI_Isend()` - will return “immediately” even before source location is safe to be altered.
- **Nonblocking receive** - `MPI_Irecv()` - will return even if no message to accept.

Nonblocking Routine Formats

`MPI_Isend(buf, count, datatype, dest, tag, comm, request)`

`MPI_Irecv(buf, count, datatype, source, tag, comm, request)`

Completion detected by `MPI_Wait()` and `MPI_Test()`.

`MPI_Wait()` waits until operation completed and returns then.

`MPI_Test()` returns with flag set indicating whether operation completed at that time.

Need to know whether particular operation completed.

Determined by accessing `request` parameter.

Example

To send an integer x from process 0 to process 1 and allow process 0 to continue,

```
MPI_Comm_rank(MPI_COMM_WORLD, &myrank); /* find rank */
if (myrank == 0) {
    int x;
    MPI_Isend(&x,1,MPI_INT, 1, msgtag, MPI_COMM_WORLD, req1);
    compute();
    MPI_Wait(req1, status);
} else if (myrank == 1) {
    int x;
    MPI_Recv(&x,1,MPI_INT,0,msgtag, MPI_COMM_WORLD, status);
}
```

Send Communication Modes

- **Standard Mode Send** - Not assumed that corresponding receive routine has started. Amount of buffering not defined by MPI. If buffering provided, send could complete before receive reached.
- **Buffered Mode** - Send may start and return before a matching receive. Necessary to specify buffer space via routine `MPI_Buffer_attach()`.
- **Synchronous Mode** - Send and receive can start before each other but can only complete together.
- **Ready Mode** - Send can only start if matching receive already reached, otherwise error. Use with care.

- Each of the four modes can be applied to both blocking and nonblocking send routines.
- Only the standard mode is available for the blocking and nonblocking receive routines.
- Any type of send routine can be used with any type of receive routine.

Collective Communication

Involves set of processes, defined by an intra-communicator.
Message tags not present. Principal collective operations:

- `MPI_Bcast()` - Broadcast from root to all other processes
- `MPI_Gather()` - Gather values for group of processes
- `MPI_Scatter()` - Scatters buffer in parts to group of processes
- `MPI_Alltoall()` - Sends data from all processes to all processes
- `MPI_Reduce()` - Combine values on all processes to single value
- `MPI_Reduce_scatter()` - Combine values and scatter results
- `MPI_Scan()` - Compute prefix reductions of data on processes

Example

To gather items from group of processes into process 0, using dynamically allocated memory in root process:

```
int data[10];          /*data to be gathered from processes*/
MPI_Comm_rank(MPI_COMM_WORLD, &myrank);      /* find rank */
if (myrank == 0) {
    MPI_Comm_size(MPI_COMM_WORLD, &grp_size); /*find group size*/
    buf = (int *)malloc(grp_size*10*sizeof (int)); /*allocate memory*/
}
MPI_Gather(data,10,MPI_INT,buf,grp_size*10,MPI_INT,0,MPI_COMM_WORLD) ;
```

`MPI_Gather()` gathers from all processes, including root.

Barrier routine

- A means of synchronizing processes by stopping each one until they all have reached a specific “barrier” call.

Sample MPI program

Evaluating Parallel Programs

Sequential execution time, t_s : Estimate by counting computational steps of best sequential algorithm.

Parallel execution time, t_p : In addition to number of computational steps, t_{comp} , need to estimate communication overhead, t_{comm} :

$$t_p = t_{\text{comp}} + t_{\text{comm}}$$

Computational Time

Count number of computational steps.

When more than one process executed simultaneously, count computational steps of most complex process.

Generally, function of n and p , i.e.

$$t_{\text{comp}} = f(n, p)$$

Often break down computation time into parts. Then

$$t_{\text{comp}} = t_{\text{comp1}} + t_{\text{comp2}} + t_{\text{comp3}} + \dots$$

Analysis usually done assuming that all processors are same and operating at same speed.

Communication Time


Many factors, including network structure and network contention. As a first approximation, use

$$t_{\text{comm}} = t_{\text{startup}} + nt_{\text{data}}$$

t_{startup} is startup time, essentially time to send a message with no data. Assumed to be constant.

t_{data} is transmission time to send one data word, also assumed constant, and there are n data words.

Idealized Communication Time



Final communication time, t_{comm}

Summation of communication times of all sequential messages from a process, i.e.

$$t_{\text{comm}} = t_{\text{comm1}} + t_{\text{comm2}} + t_{\text{comm3}} + \dots$$

Communication patterns of all processes assumed same and take place together so that only one process need be considered.

Both t_{startup} and t_{data} , measured in units of one computational step, so that can add t_{comp} and t_{comm} together to obtain parallel execution time, t_p .

Benchmark Factors

With t_s , t_{comp} , and t_{comm} , can establish speedup factor and computation/communication ratio for a particular algorithm/implementation:

$$\text{Speedup factor} = \frac{t_s}{t_p} = \frac{t_s}{t_{\text{comp}} + t_{\text{comm}}}$$

$$\text{Computation/communication ratio} = \frac{t_{\text{comp}}}{t_{\text{comm}}}$$

Both functions of number of processors, p , and number of data elements, n .

Factors give indication of scalability of parallel solution with increasing number of processors and problem size.

Computation/communication ratio will highlight effect of communication with increasing problem size and system size.

Debugging/Evaluating Parallel Programs Empirically

Visualization Tools

Programs can be watched as they are executed in a space-time diagram (or process-time diagram):

