

Figure 1: MPI messages.

1 MPI Hands-On - Sending and Receiving Messages I

Questions:

- To whom is data sent?
- What is sent?
- How does the receiver identify it?

1.1 Current Message-Passing

Message = data + envelope




Figure 2: Data+Envelope.

- MPI Data; Arguments
 - **startbuf** (starting location of data)
 - **count** (number of elements)

- * receive count \geq send count
- **datatype** (basic or derived)
 - * receiver datatype = send datatype (unless MPI_PACKED)
 - * Elementary (all C and FORTRAN types). Specifications of elementary datatypes allows heterogeneous communication.
 - * MPI basic datatypes for C:

MPI Datatype	C Datatype
MPI_CHAR	signed char
MPI_SHORT	signed short int
MPI_INT	signed int
MPI_LONG	signed long int
MPI_UNSIGNED_CHAR	unsigned char
MPI_UNSIGNED_SHORT	unsigned short int
MPI_UNSIGNED	unsigned int
MPI_UNSIGNED_LONG	unsigned long int
MPI_FLOAT	float
MPI_DOUBLE	double
MPI_LONG_DOUBLE	long double
MPI_BYTE	
MPI_PACKED	

Figure 3: MPI basic datatypes for C.

- MPI Envelope; Arguments
 - **destination or source**
 - * rank in a communicator
 - * receive = sender or MPI_ANY_SOURCE
 - **tag**
 - * integer chosen by programmer
 - * receive = sender or MPI_ANY_TAG (wild cards allowed)
 - **communicator**
 - * defines communication "space"
 - * group + context
 - * receive = send
 - Collective operations typically operated on all processes.

- All communication (not just collective operations) takes place in groups.
- A context partitions the communication space. A message sent in one context cannot be received in another context. Contexts are managed by the system.
- A group and a context are combined in a communicator.
- Source/destination in send/receive operations refer to rank in group associated with a given communicator.

1.2 The Buffer

Sending and receiving only a contiguous array of bytes. Specified in MPI by *starting address*, *datatype*, and *count*

- hides the real data structure from hardware which might be able to handle it directly.
- requires pre-packing dispersed data
 - rows of a matrix stored columnwise.
 - general collections of structures.
- prevents communications between machines with different representations (even lengths) for same data type

1.3 MPI Basic Send/Receive

Thus the basic send (blocking!!) has become:

```
MPI_Send( start, count, datatype, dest, tag, comm )
```

and the receive (blocking!!):

```
MPI_Recv(start, count, datatype, source, tag, comm, status)
```

The source, tag, and count of the message actually received can be retrieved from *status*.

```
MPI_Status status;
MPI_Recv( ..., &status );
... status.MPI_TAG; ... status.MPI_SOURCE;
MPI_Get_count( &status, datatype, &count );
```

MPI_Get_count may be used to determine how much data of a particular type was received.

Two simple collective operations (just to introduce!):

```
MPI_Bcast(start, count, datatype, root, comm)
MPI_Reduce(start, result, count, datatype,
           operation, root, comm)
```

1.4 Exercises/Examples

1. An example for communication world [code1](#).

```
#include <stdio.h>
#include <mpi.h>
int main(int argc, char **argv)
{
    int my_rank, numprocs;
    MPI_Init(&argc,&argv);
    MPI_Comm_size(MPI_COMM_WORLD,&numprocs);
    MPI_Comm_rank(MPI_COMM_WORLD,&id);
    /*printf("Hello! It is processor %d.\n", id);*/
    if (my_rank == 0)
    {
        printf("Hello! It is processor 0. There are %d processors in this
               communication world.\n", numprocs);
        printf ("I am process %i out of %i: Hello world!\n",my_rank, size);
    }
    else
    {
        printf("I am process %i out of %i: Hello world!\n", my_rank, size);
    }
    MPI_Finalize();
    return 0;
}
```

2. Write a program to send/receive and print out your name and age to each processors. Hints:

```
char* my_name = "Cem Ozdogan";
MPI_Send(&my_name, 11, MPI_CHAR, dest, 2, MPI_COMM_WORLD);
MPI_Recv(&recv_my_name, 11, MPI_CHAR, 0, 2, MPI_COMM_WORLD, &status);
```