

Exercises: Message-Passing Programming

David Henty

1 Hello World

1. Write an MPI program which prints the message “Hello World”. Compile and run on one process
2. Run on several processes in parallel.
3. Modify your program so that each process prints out both its rank and the total number of processes P that the code is running on, ie the size of `MPI_COMM_WORLD`.
4. Modify your program so that only the master process (ie rank 0) prints out a message (very useful when you run with hundreds of processes).
5. What happens if you omit the final MPI procedure call in your program?

2 Parallel calculation of π

An approximation to the value π can be obtained from the following expression

$$\frac{\pi}{4} = \int_0^1 \frac{dx}{1+x^2} \approx \frac{1}{N} \sum_{i=1}^N \frac{1}{1 + \left(\frac{i-\frac{1}{2}}{N}\right)^2}$$

where the answer becomes more accurate with increasing N . Iterations over i are independent so the calculation can be parallelised.

For the following exercises you should set $N = 840$. This number is divisible by 2, 3, 4, 5, 6, 7 and 8 which is convenient when you parallelise the calculation!

1. Modify your Hello World program so that each process independently computes the value of π and prints it to the screen. Check that the values are correct (each process should print the same value).
2. Now arrange for different processes to do the computation for different ranges of i . For example, on two processes: rank 0 would do $i = 1, 2, \dots, \frac{N}{2}$; rank 1 would do $i = \frac{N}{2} + 1, \frac{N}{2} + 2, \dots, N$. Print the partial sums to the screen and check the values are correct by adding them up by hand.
3. Now we want to accumulate these partial sums by sending them to the master (rank 0) to add up:
 - all processes (except the master) send their partial sum to the master
 - the master receives the values from all the other processes, adding them to its own partial sum

You should use the MPI routines `MPI_Ssend` and `MPI_Recv`.

4. Use the function `MPI_Wtime` (see below) to record the time it takes to perform the calculation. For a given value of N , does the time decrease as you increase the number of processes? Note that to ensure that the calculation takes a sensible amount of time (eg more than a second) you will probably have to perform the calculation of π several thousands of times.
5. Ensure your program works correctly if N is not an exact multiple of the number of processes P .

Size (bytes)	# Iterations	Total time (secs)	Time per message	Bandwidth (Mb/s)

Table 1: Ping-Pong Results

2.1 Timing MPI Programs

The `MPI_Wtime()` routine returns a double-precision floating-point number which represents elapsed wall-clock time in seconds. The timer has no defined starting-point, so in order to time a piece of code, two calls are needed and the difference should be taken between them.

3 Ping Pong

- Write a program in which two processes (say rank 0 and rank 1) repeatedly pass a message back and forth. Use the synchronous mode `MPI_Ssend` to send the data. You should write your program so that it operates correctly even when run on more than two processes, ie processes with rank greater than one should simply do nothing. For simplicity, use a message that is an array of integers. Remember that this is like a game of table-tennis:
 - rank 0 should send a message to rank 1
 - rank 1 should receive this message then send *the same data* back to rank 0
 - rank 0 should receive the message from rank 1 and then return it
 - etc. etc.
- Insert timing calls to measure the time taken by all the communications. You will need to time many ping-pong iterations to get a reasonable elapsed time, especially for small message lengths.
- Investigate how the time taken varies with the size of the message. You should fill in your results in Table 1. What is the asymptotic bandwidth for large messages?
- Plot a graph of time against message size to determine the latency (ie the time taken for a message of zero length). A simple way to produce plots is using the `axx` command. We can provide a sample file which you can view by typing `axx testplot.ax`. Clicking on the graph closes the window.

The bandwidth and latency are key characteristics of any parallel machine, so it is always instructive to run this ping-pong code on any new computers you may get access to.

3.1 Extra exercises

- How do the ping-pong bandwidth and latency figures vary when you use buffered or standard modes (`MPI_Bsend` and `MPI_Send`)?

Note: to send large messages with buffered sends you may have to supply MPI with additional buffer space using `MPI_Buffer_attach`.
- Write a program in which the master process sends the same message to all the other processes in `MPI_COMM_WORLD` and then receives the message back from all of them. How does the time taken vary with the size of the messages and with the number of processes?




Figure 1: Global sum of four variables

4 Rotating information around a ring

Consider a set of processes arranged in a ring as shown in Figure 1. A simple way to perform a global sum of data stored on each process (a parallel reduction operation) is to rotate each piece of data all the way round the ring. At each iteration, a process receives some data from the left, adds the value to its running total, then passes the data it has just received on to the right.

Figure 1 illustrates how this works for four processes (ranks 0, 1, 2 and 3) who hold values A , B , C and D respectively. The running total on each process is shown in the square box, and the data being sent between processes is shown next to the arrow. After three steps ($P - 1$ steps in general for P processes) each process has computed the global sum $A + B + C + D$.

1. Write a program that performs a global sum using this simple ring method. Each process needs to know the ranks of its two neighbours in the ring, which stay constant throughout the program. You should use synchronous sends (`MPI_Ssend`) and avoid deadlock by using non-blocking forms for *either* the send (`MPI_Issend`) *or* the receive (`MPI_Irecv`). Remember that you cannot assume that a non-blocking operation has completed until you have issued an explicit wait.

We need to initialise the local variables on each process with some process-dependent value. For simplicity, we will just use the value of *rank*, ie in Figure 1 this would mean $A = 0$, $B = 1$, $C = 2$ and $D = 3$. You should check that every process computes the sum correctly (eg print the final value to the screen), which in this case is $P(P - 1)/2$.

2. Your program should compute the correct global sum for any set of input values. If you initialise the local values to $(rank + 1)^2$, do you get the correct result $P(P + 1)(2P + 1)/6$?

4.1 Extra exercises

1. Measure the time taken for a global sum and investigate how it varies with increasing P . Plot a graph of time against P — does the ring method scale as you would expect?
2. Using these timings, estimate how long it takes to send each individual message between processes. How does this result compare with the latency figures from the ping-pong exercise?

3. The `MPI_Sendrecv` call is designed to avoid deadlock by combining the separate send and receive operations into a single routine. Write a new version of the global sum using this routine and compare the time taken with the previous implementation. Which is faster?
4. Investigate the time taken when you use standard and buffered sends rather than synchronous mode (using `MPI_Bsend` you do not even need to use the non-blocking form as it is guaranteed to be asynchronous). Which is the fastest? By comparing to the time taken by the combined send and receive operation, can you guess how `MPI_Sendrecv` is actually being implemented?

5 Collective communications

1. Re-write the ring example using an MPI reduction operation to perform the global sum.
2. How does the execution time vary with P and how does this compare to your own implementation which used the ring method?

5.1 Extra exercises

1. Use a different collective operation to perform a scan, ie a partial sum across different numbers of processes. How does the scaling with P compare to your own implementation which used a non-periodic 1D Cartesian topology?
2. Print out the results of the partial sum on each process. Now see if you can ensure that the output appears on the screen in the correct order (first rank 0, then rank 1 etc).
3. Modify the program that performed a global sum on a compound datatype so that it uses an MPI collective routine. You will have to register your own reduction operation so that the MPI library knows what calculation you want to be performed. Remember that addition is not a pre-defined operation on your compound type; it still has to be defined even in the native language.

6 Derived Datatypes

We will extend exercise 4 to perform a global sum of a non-basic datatype. A simple example of this is a compound type containing both an integer and a double-precision floating-point number. Such a compound type can be declared as a structure in C, or a derived type in Fortran, as follows:

```

struct compound
{
    int    ival;
    double dval;
};

struct compound x,y;

x.ival = 1;
y.dval = 9.0;

type compound
integer      :: ival
double precision :: dval

end type compound

type(compound) :: x,y

x%ival = 1
y%dval = 9.0

```

If you are unfamiliar with using derived types in Fortran then I recommend that you go straight to exercise number 2 which deals with defining MPI datatypes to map onto subsections of arrays. This is, in fact, the most common use of derived types in scientific applications of MPI.




Figure 2: Diagrammatic representation of the *mcols* and *mrows* matrix subsections

1. Modify the ring exercise so that it uses an `MPI_Type_struct` derived datatype to pass round the above compound type, and computes separate integer and floating-point totals. You will need to use `MPI_Address` to obtain the displacements of `ival` and `dval`. Initialise the integer part to `rank` and the floating-point part to $(rank + 1)^2$ and check that you obtain the correct results.
2. Modify your existing ping-pong code to exchange $N \times N$ square matrices between the processes (`int msg[N][N]` in C or `INTEGER MSG(N,N)` in Fortran). Initialise the matrix elements to be equal to `rank` so that you can keep track of the messages. Define `MPI_Type_contiguous` and `MPI_Type_vector` derived types to represent $N \times M$ (type *mcols*) and $M \times N$ (type *mrows*) subsections of the matrix, where $M \leq N$. Which datatype to use for each subsection depends on whether you are using C or Fortran. You may find it helpful to refer to Figure 2 to clarify this, where I draw the arrays in the standard fashion for matrices (indices start from one rather than zero, first index goes down, second index goes across).

Set $N = 10$ and $M = 3$ and exchange columns 4, 5 and 6 of the matrix using the *mcols* type. Print the entire matrix on each process after every stage of the ping-pong to check that for correctness. Now modify your program to exchange rows 4, 5 and 6 using the *mrows* type.

6.1 Extra exercises

1. For the compound type, print out the values of the displacements. Do you understand the results?
2. Modify your ping-pong code for matrix subsections so that you send type *mcols* and receive type *mrows*. Do things function as you would expect?

7 Rotating information using a Cartesian topology

For a 1D arrangement of processes it may seem a lot of effort to use a Cartesian topology rather than simply managing the processes by hand, eg calculating the ranks of the nearest neighbours. It is, however, well worth learning how to use topologies as these book-keeping calculations can become very tedious to do by hand in two and three dimensions.

For simplicity we will only actually use the routines in one dimension. Even for this simple case I hope that the exercise shows how easy it is to change things like the boundary conditions when using topologies.

1. Re-write the passing-around-a-ring exercise so that it uses a one-dimensional Cartesian topology, computing the ranks of the nearest neighbours using `MPI_Cart_shift`. Remember to set the periodicity of the boundary conditions appropriately for a ring rather than a line.

2. Alter the boundary conditions to change the topology from a ring to a line, and re-run your program. Be sure to run using both of the initial values, ie *rank* and $(rank + 1)^2$. Do you understand the output? What reduction operation is now being implemented? What are the the neighbouring ranks for the two processes at the extreme ends of the line?

7.1 Extra exercises

1. Measure the time taken for the global sum in both periodic and non-periodic topologies, and investigate how it varies with P .
2. Extend the one-dimensional ring topology to a two-dimensional cylinder (periodic in one direction, non-periodic in the other). Perform two separate reduction operations, one in each of the two dimensions of the cylinder.

7.2 Global Summation Using a Hypercube Algorithm

Although you should always perform global summations by using `MPI_Reduce` or `MPI_Allreduce` with `MPI_Op=MPI_SUM`, it is an interesting exercise to program your own version using a more efficient algorithm than the previous naive “message-round-a-ring” approach.

A more efficient method, at least for a number of processes that is a power of two, is to imagine that the processes are arranged in a cube. The coordinates of the processes in the cube are taken from the binary representation of the rank, therefore ensuring that exactly one process sits at each vertex of the cube. Processes operate in pairs, swapping partial sums between neighbouring processes in each dimension in turn. Figure 3 illustrates how this works in three dimensions (ie $2^3 = 8$ processes).




Figure 3: Communications pattern for global sum on 8 processes

An elegant way to program this is to construct a periodic cartesian topology of the appropriate dimension and compute neighbours using `MPI_Cart_shift`. When each process swaps data with its neighbour you must ensure that your program does not deadlock. This can be done by a variety of methods, including a ping-pong approach where each pair of processes agrees in advance who will do a send followed by a receive, and who will do a receive followed by a send.

How do you expect the time to scale with the number of processes, and how does this compare to the measured time taken by `MPI_Allreduce`?