

# 1 Bases en MPI

```

Hello world en MPI
1 #include <stdio.h>
2 #include <stdlib.h>
3 #include <mpi.h>
4 int main( int argc, char** argv ) {
5     int size, rank;
6     MPI_Init(&argc, &argv);
7     MPI_Comm_size( MPI_COMM_WORLD, &size);
8     MPI_Comm_rank( MPI_COMM_WORLD, &rank);
9
10    printf("I am rank %d in %d\n", rank, size);
11
12    MPI_Finalize();
13    return EXIT_SUCCESS;
14 }

```

```

Exécution réelle
1 # Compilation (gcc avec les bonnes options)
2 $ mpicc -o helloworld helloworld.c
3 # gcc -L/my/mpi/lib -lmpi -I/my/mpi/include ..
4
5 # Exécution
6 $ mpirun -np 4 -hostfile machines ./helloworld
7 I am rank 3 in 4
8 I am rank 2 in 4
9 I am rank 1 in 4
10 I am rank 0 in 4

```

```

Fichier machines
1 host-0
2 host-1
3 host-2
4 ...

```

```

Exécution simulée avec SMPI
1 # Compilation (lance gcc avec les bonnes options)
2 $ smpicc -o helloworld helloworld.c
3
4 # Exécution
5 $ smpirun -np 4 -hostfile machines -platform ./cluster.xml ./helloworld
6 [0.000000] [xbt_cfg/INFO] Configuration change: Set 'smpi/privatization' to '1'
7 [0.000000] [xbt_cfg/INFO] Configuration change: Set 'smpi/np' to '4'
8 [0.000000] [xbt_cfg/INFO] Configuration change: Set 'smpi/hostfile' to 'cluster_hostfile.txt'
9 [0.000000] [xbt_cfg/INFO] Configuration change: Set 'surf/precision' to '1e-9'
10 [0.000000] [xbt_cfg/INFO] Configuration change: Set 'network/model' to 'SMPI'
11 [0.000000] [xbt_cfg/INFO] Configuration change: Set 'smpi/tmpdir' to '/tmp'
12 [0.000000] [smpi_config/INFO] You did not set the power of the host running the simulation.
13 The timings will certainly not be accurate. Use the option "--cfg=smpi/host-speed:<flops>" to set its
14 Check https://simgrid.org/doc/latest/Configuring_SimGrid.html#automatic-benchmarking-of-smpi-code for m
15 I am rank 0 in 4
16 I am rank 1 in 4
17 I am rank 2 in 4
18 I am rank 3 in 4

```

## 2 Communications point-à-point bloquantes

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

```
int MPI_Recv( void *buf, int count, MPI_Datatype datatype, int orig, int tag, MPI_Comm comm, MPI_Status *status )
```

```

Exemple de ping-pong en MPI
1 #include <stdio.h>
2 #include <stdlib.h>
3 #include <mpi.h>
4 int main( int argc, char** argv ) {
5     int rank;
6     int token = 42;
7     MPI_Status status;
8     MPI_Init( &argc, &argv );
9     MPI_Comm_rank( MPI_COMM_WORLD, &rank );
10    if( 0 == rank ) {
11        MPI_Send( &token, 1, MPI_INT, 1, 0, MPI_COMM_WORLD );
12        MPI_Recv( &token, 1, MPI_INT, 1, 0, MPI_COMM_WORLD, &status );
13    } else if ( 1 == rank ) {
14        MPI_Recv( &token, 1, MPI_INT, 0, 0, MPI_COMM_WORLD, &status );
15        MPI_Send( &token, 1, MPI_INT, 0, 0, MPI_COMM_WORLD );
16    }
17    MPI_Finalize();
18    return EXIT_SUCCESS;
19 }

```

```

1  include <stdio.h>
2  include <stdlib.h>
3  include <mpi.h>
4
5  #define N (1024 * 1024 * 1)
6
7  int main(int argc, char *argv[])
8
9  {
10     int size, rank;
11     struct timeval start, end;
12     MPI_Init(&argc, &argv);
13
14     MPI_Comm_rank(MPI_COMM_WORLD, &rank);
15     MPI_Comm_size(MPI_COMM_WORLD, &size);
16
17     char *buf = malloc(sizeof(char) * N); //10MiB
18
19     // Communicate along the ring
20     if (rank == 0) {
21         gettimeofday(&start, NULL);
22         printf("Rank %d: sending the message rank %d\n", rank, 1);
23         MPI_Send(buf, N, MPI_BYTE, 1, 1, MPI_COMM_WORLD);
24         MPI_Recv(buf, N, MPI_BYTE, size-1, 1, MPI_COMM_WORLD, MPI_STATUS_IGNORE);
25         printf("Rank %d: received the message from rank %d\n", rank, size-1);
26         gettimeofday(&end, NULL);
27         printf("%f\n", (end.tv_sec*1000000.0 + end.tv_usec -
28             start.tv_sec*1000000.0 - start.tv_usec) / 1000000.0);
29     } else {
30         MPI_Recv(buf, N, MPI_BYTE, rank-1, 1, MPI_COMM_WORLD, MPI_STATUS_IGNORE);
31         int dest = (rank+1)%size;
32         printf("Rank %d: sending the message to rank %d\n", rank, dest);
33         MPI_Send(buf, N, MPI_BYTE, dest, 1, MPI_COMM_WORLD);
34     }
35
36     MPI_Finalize();
37     return 0;
38 }

```

### 3 Communications point-à-point asynchrones

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

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

```
int MPI_Wait( MPI_Request* request, MPI_Status* status )
int MPI_Waitall(int count, MPI_Request* reqs, MPI_Status *stats)
int MPI_Waitany(int count, MPI_Request* reqs, int *index, MPI_Status *stat)
```

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

### 4 Quelques sites web utiles

- <https://theartofhpc.com/>: Plusieurs livres très complets sur le sujet. MPI est dans le tome 2; Makefile, cmake et git sont dans le tome 4.
- [https://web.corral.tacc.utexas.edu/CompEdu/pdf/pcse/mpi\\_c\\_course.pdf](https://web.corral.tacc.utexas.edu/CompEdu/pdf/pcse/mpi_c_course.pdf) 400 slides du même auteur: Cours de MPI en C.
- <http://condor.cc.ku.edu/~grobe/docs/intro-MPI-C.shtml> Une autre introduction au MPI en C.