

# Parallelization: Sieve of Eratosthenes

By Aaron Weeden, Shodor Education Foundation, Inc.  
Exercise 1 – Introduction to Parallel Programming on a Cluster

This exercise will take you through writing a small piece of C code, outfitting it with MPI and OpenMP directives, and running it on a cluster.

You will need to obtain an account from the Cluster Computing Group at Earlham. Send email to [ccg@cs.earlham.edu](mailto:ccg@cs.earlham.edu) to request an account.

You will make use of the **vi** text editor, which is provided by default on most Linux-based operating systems, such as the one used by Earlham’s cluster.

In this exercise, any line with a dollar sign (\$) in front of it is a command to be entered in a shell (a command line utility used by the operating system to interact with the user).

## Part I: Write, compile, and run a serial program

1. Log into al-salam:

```
$ ssh <yourusername>@cluster.earlham.edu  
$ ssh as0
```

2. Create a small “Hello, World” program in C:

- a. Open a new file called `hello.c` in `vi`:

```
$ vi hello.c
```

- b. Enter `vi`’s “insert mode” by pressing the **i** key.

- c. Write a small C code that will print “Hello, World!” on the screen:

```
#include <stdio.h>  
  
int main(int argc, char** argv)  
{  
    printf("Hello, World!\n");  
  
    return 0;  
}
```

- d. Press **Escape (esc)** to exit `vi`’s insert mode.
- e. Save the file and exit `vi` by entering **<Shift>-Z-Z**.

3. Compile the code with GNU's compiler. This will produce an executable file called `hello`:

```
$ gcc -o hello hello.c
$
```

If any errors are listed, make sure there are no typos in `hello.c` (go back through step 2).

4. Create a script to run the program on the cluster.
  - a. Open a new file called `hello.qsub` in `vi`:

```
$ vi hello.qsub
```
  - b. Enter insert mode (as you did in step 2b) and write a small Portable Batch System (PBS) script:

```
#PBS -q ec
#PBS -o hello.out
#PBS -e hello.err

cd $PBS_O_WORKDIR

./hello
```

Each line of this script tells the scheduler to do something:

```
#PBS -q ec says to use the "ec" queue.
#PBS -o hello.out says to save the output of standard out to a file
called hello.out rather than to print it on the terminal.
#PBS -e hello.err says to save the output of standard error to a file
called hello.err rather than to print it on the terminal.
cd $PBS_O_WORKDIR tells the scheduler to change directories to the
directory from which the job is submitted.
./hello says to run the hello executable.
```

- c. Save the file and exit `vi` (as you did in steps 2d and 2e).

5. Submit a job to the scheduler:

```
$ qsub hello.qsub
19098.as0.al-salam.loc
$
```

6. This will submit a job and output its job ID, 19098 in this example. Your job will now be waiting in the queue, running, or finished. You can monitor it at any time by entering `qstat 19098` (or whatever your Job ID is) in the shell.

You may see something like the following:

```
qstat 19098
qstat: Unknown Job Id 19098.as0.al-salam.loc
```

This means the job is complete.

If the job were instead still running, you would see something like the following table:

Job id	Name	User	Time Use	S	Queue
19098.as0	STDIN	amweeden06	0	R	ec

In this output, the S column is the status column. The letter under this column tells you the status of the job; Q means it is waiting in the queue and R means it is running.

7. Once the job is complete, show the contents of `hello.out` with the `cat` command:

```
$ cat hello.out
Hello, World!
$
```

8. `hello.err` should be empty if there were no errors in running the program. Show the contents of `hello.err` with the `cat` command:

```
$ cat hello.err
$
```

If this command returns just a prompt (`$`), then the file is empty and there were no errors. Otherwise, the errors will be listed.

## Part II. Outfit the program with MPI

9. We will now make a parallel version of the code using MPI. First we will tell the program to include the MPI library. We also tell the program that we are using MPI by putting `MPI_Init` at the top of `main` and `MPI_Finalize` at the bottom.

- a. Open `hello.c` and add the lines to the code as below:

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

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

```
{
    MPI_Init(&argc, &argv);

    printf("Hello, World!\n");

    MPI_Finalize();

    return 0;
}
```

b. Save and quit the file.

10. Compile the code with GNU's MPI compiler:

```
$ mpicc -o hello hello.c
$
```

If any errors are listed, make sure there are no typos in `hello.c` (go back through step 9).

11. Edit the PBS script to use the MPI run command:

a. Open the `hello.qsub` file in `vi`:

```
$ vi hello.qsub
```

b. Enter insert mode and change the last line to use `mpirun` as below:

```
#PBS -q ec
#PBS -o hello.out
#PBS -e hello.err

cd $PBS_O_WORKDIR

mpirun -np 2 ./hello
```

Here `-np 2` tells MPI to use 2 processes. Both processes will run the `hello` executable.

c. Save and quit the file.

12. Submit a job to the scheduler:

```
$ qsub hello.qsub
19099.as0.al-salam.loc
$
```

13. Monitor the job with `qstat`. Once it finishes, view the contents of standard out and standard error:

```
$ cat hello.out
```

```
Hello, World!  
Hello, World!  
$ cat hello.err  
$
```

What do you notice about hello.out this time?

14. Let's have the processes print some useful information. We will have them print their rank, the total number of processes, and the name of the processor on which they are running.

a. Open hello.c in vi:

```
$ vi hello.c
```

b. Add the following lines to hello.c:

```
#include <mpi.h>  
#include <stdio.h>  
  
int main(int argc, char** argv)  
{  
    int rank = 0;  
    int size = 0;  
    int len = 0;  
    char name[MPI_MAX_PROCESSOR_NAME];  
  
    MPI_Init(&argc, &argv);  
  
    MPI_Comm_rank(MPI_COMM_WORLD, &rank);  
    MPI_Comm_size(MPI_COMM_WORLD, &size);  
    MPI_Get_processor_name(name, &len);  
  
    printf("Hello, World from rank %d of %d on %s\n",  
           rank, size, name);  
  
    MPI_Finalize();  
  
    return 0;  
}
```

rank will be the rank of the process, size the total number of processes, and name the name of the processor on which the process is running.

c. Save and quit the file.

15. Compile the code with GNU's MPI compiler:

```
$ mpicc -o hello hello.c
```

```
$
```

16. We don't need to change the PBS script because we will be using the same `mpirun -np 2 ./hello` command to execute the program. We will expect to see something different in `hello.out`, however. Let's submit the job and see what we get:

```
$ qsub hello.qsub
19100.as0.al-salam.loc
$
```

17. Monitor the job with `qstat`, and once it is finished check the contents of `hello.out` and `hello.err`:

```
$ cat hello.out
Hello, World from rank 0 of 2 on as1.al-salam.loc
Hello, World from rank 1 of 2 on as1.al-salam.loc
$ cat hello.err
$
```

On which processor did Rank 0 run for you? How about Rank 1?

18. Let's try running across multiple nodes instead of just one node (`as1.al-salam.loc` in the example above). Edit the `hello.qsub` file to include the following lines:

```
#PBS -q ec
#PBS -o hello.out
#PBS -e hello.err
#PBS -l nodes=2:ppn=1
```

```
cd $PBS_O_WORKDIR
```

```
mpirun -np 2 ./hello
```

The line that we added, `#PBS -l nodes=2:ppn=1`, says to run the job on 2 nodes with 1 process per node.

19. Submit a job:

```
$ qsub hello.qsub
19104.as0.al-salam.loc
$
```

20. Monitor the job with `qstat` until it finishes, then output the contents of `hello.out` and `hello.err`:

```
$ cat hello.out
```

```
Hello, World from rank 0 of 2 on as2.al-salam.loc
Hello, World from rank 1 of 2 on as1.al-salam.loc
$ cat hello.err
$
```

Now on which processor did Rank 0 run for you? Rank 1?

### Part III. Outfit the program with OpenMP

1. We first need to tell the program to include the OpenMP library. Open `hello.c` in `vi` and add a line to the top:

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

2. Compile the code with OpenMP support through the GNU compiler by using the `-fopenmp` option:

```
$ mpicc -fopenmp -o hello hello.c
$
```

3. OpenMP does not require any special run command or arguments. We may wish to tell the program how many OpenMP threads over which to parallelize, however. Open `hello.qsub` and add a line before the `mpirun` command:

```
cd $PBS_O_WORKDIR

export OMP_NUM_THREADS=2

mpirun -np 2 ./hello
```

This line tells the program to spawn 2 OpenMP threads per process when it executes an OpenMP parallel section.

4. Submit a job with `qsub`, monitor it with `qstat` until it finishes, and then view the contents of `hello.out` and `hello.err`:

```
$ qsub hello.qsub
19105.as0.al-salam.loc
$ qstat 19105
qstat: Unknown Job Id 19105.as0.al-salam.loc
$ cat hello.out
Hello, World from rank 0 of 2 on as2.al-salam.loc
Hello, World from rank 1 of 2 on as1.al-salam.loc
$ cat hello.err
```

Parallelization: Sieve of Eratosthenes

Aaron Weeden - Shodor Education Foundation, Inc.

Exercise 1

Page 7

\$

What do you notice about the output? You might expect to see 4 “Hello, World”s because the program is supposed to spawn 2 OpenMP threads per process. However, OpenMP will not spawn any threads unless it is explicitly told to do so by marking a parallel section with an OpenMP **pragma**, hence we still only get 2 “Hello, World”s.

5. Let’s mark the `printf` as part of a parallel section so each thread will print the rank of the process, the total number of processes, the thread number, the total number of threads, and the processor on which it is running. Open `hello.c` in `vi` and make the following change to the `printf`:

```
#pragma omp parallel
{
    printf("Hello, World from rank %d of %d, thread %d
of %d on %s\n", rank, size, omp_get_thread_num(),
omp_get_num_threads(), name);
}
```

Note that we have now surrounded the `printf` by `#pragma omp parallel` followed by curly braces. This indicates that the `printf` is part of a parallel section that will be executed by multiple OpenMP threads.

Note also that we added the functions `omp_get_thread_num()` and `omp_get_num_threads()`. These will return the thread number of the thread and the total number of threads, respectively.

6. Compile the code and submit a job. Monitor it with `qstat` until it finishes, then view the contents of `hello.out` and `hello.err`:

```
$ mpicc -fopenmp -o hello hello.c
$ qsub hello.qsub
19106.as0.al-salam.loc
$ qstat 19106
qstat: Unknown Job Id 19106.as0.al-salam.loc
$ cat hello.out
Hello, World from rank 0 of 2, thread 1 of 2 on
as2.al-salam.loc
Hello, World from rank 0 of 2, thread 0 of 2 on
as2.al-salam.loc
Hello, World from rank 1 of 2, thread 0 of 2 on
as1.al-salam.loc
```

```
Hello, World from rank 1 of 2, thread 1 of 2 on  
as1.al-salam.loc  
$ cat hello.err  
$
```

What do you notice about the output now?

This completes the exercise.