

ECE574: Cluster Computing – Homework 6

MPI

Due: Friday, 19 March 2021, 5:00pm

1. Background

- In this homework we will take the sobel code from earlier homeworks and parallelize it using MPI.

2. Setup

- For this assignment, log into the same Haswell-EP machine we used in previous homeworks. As a reminder, use the username handed out in class and ssh in like this

```
ssh -p 2131 username@weaver-lab.eece.maine.edu
```
- Download the code template from the webpage. You can do this directly via

```
wget http://web.eece.maine.edu/~vweaver/classes/ece574/ece574_hw6_code.tar.gz
```

to avoid the hassle of copying it back and forth.
- Decompress the code

```
tar -xzvf ece574_hw06_code.tar.gz
```
- Run make to compile the code.
- You may use your own code from a previous assignment as a basis for this assignment. (Alternately some really poorly-optimized sample code is provided).

3. Coarse-grained Code (10 points)

Use MPI to parallelize your code. Use the sample code, or you might want to use one of your previous assignments as a basis.

If not using the sample code you will need to make sure your code includes `mpi.h` and that it calls `MPI_init()` at the beginning and `MPI_Finalize()` at the end.

Edit the file `sobel_coarse.c`

Be sure to comment your code!

A suggested first (coarse) implementation is this:

- (a) Get the rank and size parameters.
- (b) Be sure to only load the jpeg in rank 0.
- (c) You will need to send the image parameters (`image.x`, `image.y`, `image.depth`) to all the other ranks so they know how big to allocate `new_image`, `sobel_x`, and `sobel_y`. MPI is optimized for sending arrays of same sized data, so sending an array of 3 INTS might be your best bet.
- (d) You will also need to `malloc()` `image.pixels` in the non rank-0 threads. (Because usually it's `load_jpeg()` that does that).
- (e) Use `MPI_Bcast()` to broadcast the entire image data from rank0 to all the other ranks. You want to broadcast `image.pixels`, not the entire `image` struct (remember, in MPI you can't send structs, just arrays).

- (f) Modify `generic_convolve()` so it takes a range of `y` to operate on. Then calculate this `y` range based on the rank and size parameters. (For example, if we detect size of 4, then we are running on 4 nodes and each rank should get 1/4 of the results)
- (g) Be sure to run both `sobel_x` and `sobel_y` on the subset for that rank.
- (h) Use `MPI_Gather()` to get the results and combine them into the result in rank 0. NOTE: `MPI_Gather()` will gather from the `*start*` of each buffer and put it in the proper place in the result. So you have to modify the convolve routine to store the output starting at offset 0, rather than at offset `ystart`. If you forget this, the bottom $((N-1)/N)$ th of your image will be blank.
- (i) On rank 0 alone, run `combine`.
- (j) On rank 0 alone, write the output to a file.
- (k) Also be sure the code calls `MPI_wtime()` to get the wallclock times for Load, Convolve, Combine, and store much like we did with PAPI in the OpenMP code. You only need to record and print these from rank 0.

Run on the Haswell-EP machine for 1, 2, 4, 8 and 16 threads and report the results, as well as reporting the speedup and parallel efficiency for the total time.

Run your code with:

```
sbatch -n X time_coarse.sh
```

where you replace `X` with the number of cores to use.

The md5sum of the expected output is `7a17b02fe7e4e676b575f6f66ba4fa01`.

If (for fun) you want a bigger image to test with, try `/opt/ece574/jan_15_2017_high_res.jpg`

4. Some Debugging Hints

- If you have puzzling results, debug at each step of the way.
- Start by testing the `N=1` case, then `N=2` case
- You can verify the initial `MPI_Bcast()` is working by adding code something like

```
if (rank==1) {
    store_jpeg("out1.jpg", &image);
}
```

and verify that the output image is same as the input image.

- One you verify the broadcast works, do the same for the separate `sobelx` and `sobely` results.
- After verifying those looks good, do the same for after the gather. Make sure gather range is the same as your convolve ranges.
- Other things to watch for:
 - If you get a diagonal pattern in the output, be sure you are gathering in even multiples of `xsize`
 - Be sure your limits are set properly. Print the limits out and verify they are being set properly.

5. Submitting your work.

- Be sure to edit the README to include your name, as well as the timing results, and any notes you want to add about your something cool.
- Run `make submit` and it should create a file called `hw06_submit.tar.gz`. E-mail this file to me.
- e-mail the file to me by the homework deadline.