ECE574: Cluster Computing – Homework 9
HPC Power / Energy / Exam Review

Due: Tuesday 25 April 2017, 3:30pm

1. **Background**

   - We will measure the energy/energy-delay of some simple code.
   - You may work in a group of two if you want.

2. **Setup**

   - For this assignment you will need to log into the "quadro" machine that has the NVIDIA GPU. First log into the Haswell-EP machine As a reminder, use the username handed out in class and ssh in like this
     `ssh -p 2131 username@weaver-lab.eece.maine.edu`
     From there, type "ssh quadro"
   - Download the code template from the webpage. You can do this directly via
     `wget http://web.eece.maine.edu/~vweaver/classes/ece574_2017s/ece574_hw09_code.tar.gz`
     to avoid the hassle of copying it back and forth.
   - Decompress the code
     `tar -xzvf ece574_hw09_code.tar.gz`
   - Run `make` to compile the code.

3. **CUDA vs C (2 pts)**

   - The code in question is a SAXPY code, that is it does a single-precision (32-bit floating point) vector multiply plus add, i.e. $y[i] = a * x[i] + y[i]$
   - The code provided is both a C version and a CUDA version.
   - By default a vector size of 1 million is being used. Both the C and CUDA version take a command line argument that specifies how many repetitions of SAXPY to run.
   - Use the time command to measure the wall-clock time needed for 1, 4, 16, 64, 256, 1024, and 2048 iterations for both the C and CUDA versions.
   - You should see results similar to this:
     1 repetition C=0.033s, cuda=0.141s
     2048 repetitions C=11.65s cuda= 3.184s

   (a) Why is the C version faster with only 1 repetition?

   (b) Why is the Cuda version faster for 2048 repetitions?

   (c) What is the crossover point where Cuda is faster than C?

   (d) How could you improve the performance of the C version?

4. **Power/Energy (4 points)**

   - Calculate the energy needed for SAXPY-C using the RAPL counters. Run the SAXPY results for 2048 repetitions, but measuring energy usage. To measure the CPU energy being used, use a command line like this:
     `perf stat -a -e cycles,power/energy-pkg/ ./saxpy_c 2048`

   - Calculate the power used by CUDA SAXPY. First get the CPU energy usage by using the above command (but for `saxpy` instead of `saxpy_c`).
   - Then add in the energy usage of the NVIDIA card. You can measure the power used by the GPU with the `nvidia-smi` tool. Try running it to look at the results. You can get the power results specifically with the following command:
     `nvidia-smi --query-gpu=utilization.gpu,power.draw --format=csv -lms 100`
     If you run that in one window while running SAXPY in another you will see that the CUDA SAXPY instantaneous power. Get the average power and multiply by time to get the energy used.

   (a) How much total CPU energy is consumed by the C implementation?

   (b) How much total CPU+GPU energy is consumed by the GPU implementation?

   (c) Which implementation would you choose if speed were most important? If Energy were most important? If Energy delay were most important?

5. **Reliability / Checkpointing (2 points)**

   (a) Why might a cluster located at an observatory at the top of Mauna Kea in Hawaii have a higher failure rate than an identical cluster located at UMaine?

   (b) List a benefit to using application-level checkpointing in your code.

   (c) List a downside to using application-level checkpointing in your code.

6. **Big Data / Hadoop (2 points)**

   (a) What two major operations are used by Hadoop?

   (b) What language is used when writing Hadoop code?

   (c) Name one benefit of a distributed filesystem (such as the hadoop HDFS filesystem) over a centralized filesystem such as NFS.

7. Submitting your work.

   - Send me a document (pdf, txt, docx) including your name/names, data asked for, and answers to the questions. Please e-mail your document to me.