CPS 5401 Fall 2015 Name _____
Homework 3
Due Tuesday, November 3

Architecture and OpenMP

To turn in your assignment, please upload your program files and your written answers to your SVN repository in a directory named <username-hw3>, where <username> is your Miner username.

1. Construct the roofline model for a Stampede supercomputer multicore node (without the GPU or Xeon Phi accelerator). Use the machine specifications to obtain the horizontal roofline. Use the STREAM benchmark to obtain the memory bandwidth roofline.

2. Write an OpenMP parallel version of matrix-vector multiplication. Be sure to use a scalar variable to accumulate the sum inside the loop so as to avoid false sharing. Do a few runs with increasing dimensions until you achieve a stable FLOPS rate that is not dominated by overhead. Determine the arithmetic intensity of your code and place it on the roofline model from question 1. Compare your measured performance with the roofline model prediction. You may need to limit how many threads you use to the number that saturates the memory bandwidth of the node.

3. You have been provided with a Fortran90 implementation of naïve matrix-matrix multiply in the file matmul.f90. Run this code, along with variants you program, to fill in the table below with the runtimes in seconds. The ijk, kij, etc. variants refer to the orderings of the I, j, and k loops in the code. The dgemm variant refers to calling the Intel Math Kernel Library (MKL) dgemm library routine. Refer to the Stampede User Guide for instructions on how to use MKL. Explain the differences in performance that you observe in terms of cache effects.

Order	**ijk**	**jki**	**kij**	**kji**	**dgemm**
750					
1500					

4. You have been provided with an OpenMP code compute_pi.c that uses a Monte Carlo method to compute π. However, the code produces incorrect results. Find and fix the problem. Be sure to avoid unnecessary false sharing and serialization in your solution so that it runs efficiently.