COMP 705: Advanced Parallel Computing HW 7: Accelerating Code using Hybrid Programming: MPI+OpenMPI and MPI+GPU

Mary Thomas

Department of Computer Science
Computational Science Research Center (CSRC)
San Diego State University (SDSU)

Due: 11/20/17
Last Updated: 11/06/17

Table of Contents

"Wave" Generator Using Matrix-Matrix Multiplication of Airy Disk Function

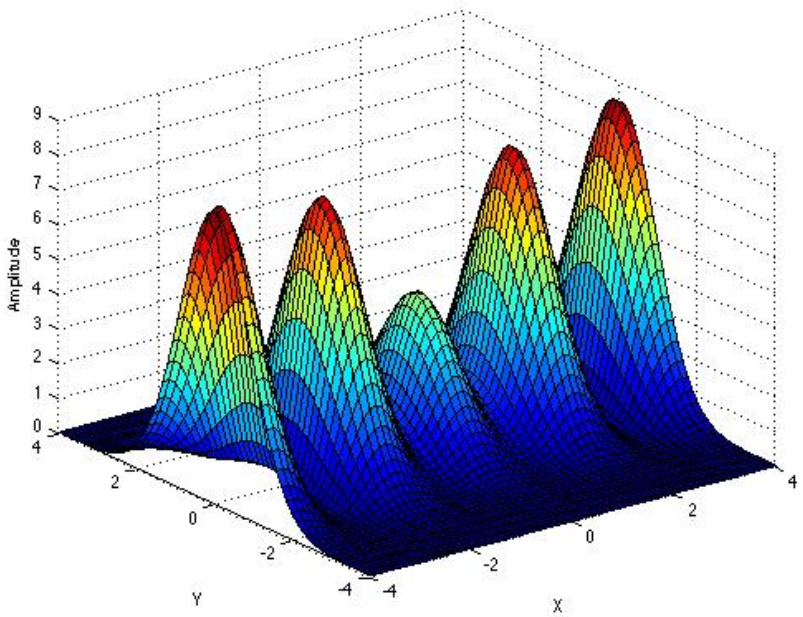

"Wave" Generator Using Matrix-Matrix Multiplication of an Airy Disk Function

See topic notes: https://edoras.sdsu.edu/~mthomas/f17.705/homework/WaveCalcsUsingMat-Mat-Mult.pdf

Matrix-Matrix Multiplication Using Hybrid MPI+OpenMP, and MPI+GPU

- In this assignment you will use an MPI 2D matrix-matrix multiplication that computes a wave-like function based on an airy disk pattern.
- You will accelerate the model using hybrid MPI+OpenMP, and MPI+GPU approaches.
- See topic notes on hybrid programming: `https://edoras.sdsu.edu/~mthomas/f17.705/topics/hybrid`

Prob 6A: Hybrid MPI+OpenMP Matrix-Matrix Multiplication

- Process command line args for relevant variables
- Test performance as a function of matrix size and number of threads.
- Create batch script to run on single CPU nodes, and multiple OpenMP threads.
- Note: See class topic notes with hints on running MPI+OpenMP jobs on tuckoo:
  `https://edoras.sdsu.edu/~mthomas/f17.705/topics/hybrid`

Prob 6B: Parallel 3D Jacobi Solver Using Hybrid MPI+GPU

- Process command line args for relevant variables
- Use CUDA device functions to get device information such as number of devices on the node, max #threads, number of blocks, etc.
- Create batch script to run on multiple CPU nodes, and multiple gpu devices on each node (try 2).
- Test performance as a function of matrix size and number of GPU cores.
- Note: See class topic notes with hints on running MPI+OpenMP jobs on tuckoo:
  `https://edoras.sdsu.edu/~mthomas/f17.705/topics/hybrid`

Notes

- How do you compare performance between serial, mpi, mpi+openmp, mpi+gpu?
- Characterize the performance of the model as a function of problem size, MPI cores, and OpenMPI threads or GPU cores.
- Define a metric for comparing the serial, MPI, and hybrid models:
- How do you directly compare different methods using metric like GFLOPS vs number of points (ProbSize)?
- does code scale with MPI nodes only? OpenMPI threads? GPU cores only?