Elective course PHY690W

Course Title: High Performance Computing and Machine Learning

Instructor: Mahendra K. Verma (PHY)

About the course: A significant part of the course would be on parallel programming for scientific applications. The idea is to enable students to write and analyse parallel programs. The focus would be on parallel programming tools, specially MPI and OpenMP. We will start with parallel Python and then go to parallel programming in C++. Hence, a good knowledge of these two computer languages is mandatory. We will also briefly cover parallel programming on GPUs.

In addition, we will provide introduction to machine learning. Due to lack of time, we will cover only a small fraction of topics in this field (regression and basics of neural networks).

Prerequisite: Good knowledge in computer programming in Python and C++. Good background in partial differential equations (PDEs), specifically flow equations and Schrödinger equation.

Who can take the course: Ph. D., Masters, and Advanced UG students (4th and 5th year). I will restrict the number of students to 25 because it is an intense course with projects.

Course Contents:

1. Introduction to HPC and scientific computing. Overview of major applications
2. Supercomputing architecture; multicores; shared memory; switch etc.
3. Review of Python programming
4. Simple tricks for speeding up computer codes, e.g., loop restructuring, vectorisation
5. Parallel Python; Introduction to Message Passing Interface (MPI)
6. Solving PDEs parallels in Python
7. Moving to parallel programming in C++
8. Programming in OpenMP
9. Case study on some major applications
10. Managing large data; visualisation tools.
11. Parallel programming on GPUs; using Nvidia
12. Basics of Machine learning (ML). Role of parallel programming in ML.
13. Using ML tools on GPUs

Evaluation: We will follow project-based evaluation. The projects will be quite difficult; good programming background is required for the projects. There will be regular quizzes and assignments in the first part of the course. There will be points for class participation and group work. A rough division of marking: Projects (50%); Quizzes + HW (35%); Attendance and class participation (15%).

References

1. P. S. Pacheco, An Introduction to Parallel Programming, Elsevier (2011)
2. M. Quinn, Parallel Programming in C and OpenMP, McCraw Hill Education (India) (2003)
3. A. Grama, A. Gupta, G. Karypis, and V. Kumar, Introduction to Parallel Computing, Pearson (2007)
4. Research papers.