COMP 5704: Parallel Algorithms and Applications in Data Science

Course Outline

Multiprocessor architectures from an application programmer's perspective: programming models, processor clusters, multi-core processors, GPUs, algorithmic paradigms, efficient parallel problem solving, scalability and portability. Projects on high performance computing in Data Science, incl. data analytics, bioinformatics, simulations. Programming experience on parallel processing equipment.

Course Material

- Lectures
- Reading Material
- Resources

Student Projects

- Projects & Presentations
- Steps & Deadlines
- Project Template

Student Evaluation

Student Projects: Every student is expected to select a topic of interest for a parallel programming project. A project includes (a) a literature review, (b) research and programming work, and (c) a final paper.

Student Presentations: Every student is expected to give a class presentation related to his/her project.

There are no further assignments or exams. The course evaluation is based on the project (70%: 20% for literature review, 30% for the quality of the research or programming work, 20% for the final paper organization and style) and presentation (30%: 15% for the quality of the slides, 15% for the quality of the presentation).