CSC367: Parallel Programming
Winter 2023

Instructor

Massimiliano Meneghin
max.meneghin@utoronto.ca

Course Overview

This course provides discussions on parallel computing and its applications. A combination of lectures, assignments, labs, and the course project will expose students to parallel programming models such as OpenMP, Pthreads, MPI, and CUDA, as well as parallel algorithms. Students will learn how to write fast code for various scientific computations to execute on modern high-performance architectures.

Required Prerequisites:

CSC258: Computer Organization, CSC209: Software Tools and Systems Programming

Recommended Prerequisites:

CSC369: Operating Systems

Sections:

Lectures:

- Wednesdays and Fridays, 15:00-16:00 PM.
- Location: MP 102

Instructor:

- Massimiliano Meneghin
- max.meneghin@utoronto.ca

Office hours**:**

- Fridays 16:00-17:00 PM,
- Location: Instructor Office BA2283

Tutorial/Labs**:**

- Mondays, 15:00-16:00 PM. *This is the default lab date, however, if Scinet is down, we might swap a lab day with a lecture day, and the lab locations might also change for that day! Check Quercus announcements frequently to not miss updates!*
- Location: BA3175 BA3185 BA3195

Discussions:

Piazza for questions and discussions. You will be automatically added to the class Piazza.

Quercus for all course material along with important announcements. You should already be in Quercus.

Important dates:

Item	Release date	Due date
Assignment one	Tuesday, Jan. 17	Tuesday, Jan. 31
Assignment two	Tuesday, Jan. 31	Tuesday, Feb. 14
Assignment three	Tuesday, Feb. 14	Tuesday, Feb. 28
Project	Tuesday, Feb. 28	Phase 1: Tuesday, Mar. 14 Phase 2: Tuesday, Mar. 21
Assignment four	Tuesday, Mar. 21	Tuesday, Apr. 4

Item	Date
In-tutorial exam	Monday, Mar 6

Class Evaluation:

The final grade is based on one in-tutorial test, the course project, the assignments, and labs. There will be 4 programming assignments along with a programming project. There are approximately 8 labs.

%	Item
Assignments	40%
Project	30%
Labs	10%
In-tutorial test	20%

Late Policy:

All labs are due at 10 PM on the due date. Late lab submissions get 0 marks.

Assignments and the project are due at 10:00 PM on the due date. Check the handouts for final due dates. For assignments and the project each student will have 10 “grace” tokens. Each grace token is worth a 2-hour extension without any penalty. The 10 tokens are for the entire term, not per assignment/project. After you use all 10 tokens, any late submissions will get a mark 0. For each 2-hours of being late, a token gets deducted from *both* partners. Since a token gets deducted from both partners, if you repartner with someone else for a later assignment, you should be aware that you can only use up to the minimum between the tokens each of you has left.

Schedule (tentative):

Week	Topics
1	Introduction and overview of memory hierarchies
2	Performance modeling and shared memory
3	Shared memory parallelism and Pthreads
4	Introduction to OpenMP
5	Distributed memory programming and MPI
6	Sources of parallelism and locality
7	Reading week!
8	Performance measurement and parallel algorithms
9	Parallel algorithms continued and GPUs
10	Introduction to GPUs
11	CUDA programming
12	CUDA programming
13	Domain-specific languages

Important policies:

- You have to read the announcements posted in Quercus.
- Violation of Scinet policies will lead to failing the course. You are responsible to read the *Introduction to Scinet* documentation (released during the first week of class) and carefully follow all rules.
- Please do not send emails directly to the TAs. They will not be replied. Emails to the instructor will be replied within 48 hours or 72 hours if it's a weekend but these emails have to be limited to urgent and personal enquiries.
- Piazza should be used to answer and discuss questions about assignments, the project, and labs. The instructor and the TAs will reply to Piazza questions within 48 hours of the post time or 72 hours if it's a weekend.

- Except for the GPU related labs and assignments, all your code and solutions will be graded on the Scinet cluster. We will post announcements if Scinet experiences downtimes. If needed we will ask you to use other computing resources for Scinet downtimes. Carefully follow all announcements!
- Your submissions and code should (1) follow submission instructions carefully to avoid penalties; (2) compile! If your code doesn't compile you will get 0 marks.

Academic Integrity: Plagiarism

- Plagiarism is an extremely serious academic offence with severe penalties.
- Do not post your code publically! Your repositories have to be private! Even after the course is over.
- The recording of course lectures should not be distributed outside of the class. Doing so will violate the plagiarism policy.
- Read the licences in all lab, assignment, and project handouts carefully.
- Do not search for solutions - we will use technology to detect this!
- You should not use code fragments from public places, even if it is open source.
- You and your partner are both responsible for any detected plagiarism.
- Do not give your code to anyone!
- Code or answers to assignment questions should not be posted in Piazza.

Strongly Recommended Readings

- Introduction to Parallel Computing - A. Grama, A. Gupta, G. Karypis, V. Kumar
- CUDA by example: An introduction to General-Purpose GPU programming – J. Sanders, E. Kandroot
- Computer Architecture: A Quantitative Approach- David A. Patterson and John L. Hennessy

Recommended Additional Reading

- An Introduction to Parallel Programming - Peter Pacheco
- Parallel programming in C with MPI and OpenMP - Michael J. Quinn
- The Art of Multiprocessor Programming - Maurice Herlihy and Nir Shavit
- D.H. Bailey, R.F. Lucas and S. Williams 2010. *Performance Tuning of Scientific Applications*, CRC Press.
- G. Hager and G. Wellein 2010. *Introduction to High Performance Computing for Scientists and Engineers*, CRC Press.
- Lin, C., and Snyder, L. 2008. *Principles of Parallel Programming*, Addison Wesley.
- Mattson, T. G., Sanders, B. A., and Massingill, B. L. 2004. *Patterns for Parallel Programming*, Addison Wesley.
- Grama, A., Gupta, A., Karypis, G., and Kumar, V. 2003. *Introduction to Parallel Computing*, Second Edition, Addison Wesley.
- Dongarra, J., et al. 2002. *The Sourcebook of Parallel Computing*, Morgan Kaufmann.
- Goedecker, S., and Hoisie, A. 2001. *Performance Optimization of Numerically Intensive Codes*, Society for Industrial Mathematics.