EE/CSCI 451
Parallel and Distributed Computation
Spring 2021

Units: 4.0

Term—Day—Time—Location:
Spring – Tuesday, Thursday 2:00 – 3:50pm
https://usc.zoom.us/j/96191033552?pwd=a3N4RlB5RnduRU1NcVp3emZ5MWk0dz09
Meeting ID: 961 9103 3552
Passcode: J_xv26CAOr

Lab/Lecture: Fri 4:30 – 5:20pm
https://usc.zoom.us/j/95857808208?pwd=S2FQNFpmeU05eTZlUlFndmxyc295Zz09
Meeting ID: 958 5780 8208
Passcode: 3w049e

Instructor: Xuehai Qian
Office: 3740 McClintock, EEB 204
Office Hours: Tuesday/Thursday 4-5pm
https://usc.zoom.us/j/96567209020?pwd=M2ZpSm82TVIxditNSWROd2NJdU9rUT09
Meeting ID: 965 6720 9020
Passcode: 735767
Contact Info: xuehai.qian@usc.edu; 213-740-4459

Teaching Assistant:
Youwei Zhuo (youweizh@usc.edu),
Qinyi Luo (qinyiluo@usc.edu) ,
Office hour: 10-11am, Monday/Wednesday/Friday
https://usc.zoom.us/j/96319784829?pwd=ekRSS2oxc0dYMXV4NTNiMDlJTXkxQT09
Meeting id: 963 1978 4829
Passcode: 8u356g

Course Description

This is a revised version of the course. It will systematically cover the topics from parallel architecture to high-level parallel application development. We will cover the following components:

- Necessary but not exhaustive description of parallel architecture concepts including memory wall, cache coherence, memory consistency, communication and synchronization, and interconnection network, etc. Based on these concepts, some examples of parallel architectures are discussed.
- Abstract machine models: PRAM, including different types of the model and the relationship on computation power. LogP and CTA, why they can model the performance in a more realistic manner. Network model, concepts that measures the performance and cost of a network, network embedding.
- Reasoning about performance and parallelism: explain the various factors affecting the performance of parallel applications, the relationship between communication and computation, load balance, etc. The concept of speedup, efficiency, and isoefficiency. Two important laws, Amdahl's and Gustafson's Law.
- Parallel algorithms based on PRAM and network model. Discuss several parallel algorithms specified in PRAM with different assumptions (e.g., exclusive/concurrent write). A set of communication operators, their semantics and implementation on different network topologies.
- Parallel programming models: shared memory, message passing (MPI)
- Scalable techniques and principles: partition of workload, execution critical path, load balance, locality principle, task dependency, etc.
- GPU architecture and programming. A high-level understanding of GPU architecture, the execution model, and several examples of modern GPUs. Basics of GPU programming, understand factors affecting performance, e.g., divergence, cache management and data placement, communication and computation overlapping.
- Domain-specific language and system: distributed graph processing systems and its programming model, implementation of simple algorithms in the framework. Distributed machine learning with TensorFlow, distributed training.

Prerequisite(s): EE 355 or CSCI 201
Co-Requisite(s): None
Concurrent Enrollment: None
Recommended Preparation: High level programming

Course Notes
Lecture slides will be made available in advance of the lectures.

Technological Proficiency and Hardware/Software Required
Desktop or notebook for accessing the computing resources at USC HPC and remote Cloud.

Suggested Readings and Supplementary Materials
Lecture slides

Description and Assessment of Assignments
There will be approximately five home works and five to seven programming homeworks. The course will also include a parallel programming project.

Grading Breakdown

Homework: 20%
Programming assignments: 20%
Course project: 10%
Midterm 1: 15%
Midterm 2: 15%
Final Exam: 20%

Assignment Submission Policy
Home works to be submitted in USC blackboard. Programming homeworks to be submitted online.

Homeworks must be done independently.

Each student will have two late submissions (24 hours after the deadline) with no penalty, any late submission more than two will receive no credit for that portion.

EE/CSCI 451 Spring 2018
Tentative Course Schedule: A Weekly Breakdown for Lectures

Date	Topics	HW/PA dues	HW/Project/Exam Information
1/19	Introduction		
1/21	Parallel computer architecture fundamentals (1)		
1/26	Parallel computer architecture fundamentals (2)	PA1 due on 1/29 HW1 due on 1/31	HW1: introduction+parallel computer architecture
1/28	Abstract machine models (1)		PA1: C/C++, single-thread programming
2/2	Abstract machine models (2)	PA2 due on 2/5 HW2 due on 2/7	HW2: abstract machine model+performance reasoning
2/4	Reasoning about Performance		PA2: p-thread programming
2/9	Parallel algorithms based on shared memory		
2/11	Parallel algorithms based on network model (1)		
2/16	Parallel algorithms based on network model (2)	HW3 due on 2/18	HW3: parallel algorithms based on shared memory and network
2/18	Parallel programming concepts		
2/23	**Mid-term (1)**		**Scope: before (not including) parallel programming concepts**
2/25	Scalable techniques (1)		
3/2	Scalable techniques (2)	PA3 due on 3/5 HW3 due on 3/7	HW3: scalable techniques
3/4	Shared memory programming (1)		PA3: Mapreduce/Spark programming
3/9	Shared memory programming (2)	PA4 due on 3/14	
3/11	Message passing programming		PA4: OpenMP/MPI programming
3/16	GPU and CUDA programming (1)	HW4 due on 3/19	HW4: shared memory and message passing
3/18	Final project proposal presentation		**Each group present the proposal**
3/25	GPU and CUDA programming (2)	**Proposal due 3/23**	
3/30	GPU and CUDA programming (3)	PA5 due on 3/31	PA5: CUDA programming model and execution model
4/1	**Mid-term (2)**		**Scope: parallel programming concepts to message passing**
4/6	GPU and CUDA programming (4)	PA6 due on 4/10	PA6: CUDA memory model/hierarchy
4/8	GPU and CUDA programming (5)		
4/13	GPU and CUDA programming (6)	PA7 due on 4/17	PA7: concurrency and multi-GPU (tentative)
4/15	Parallel graph algorithms		
4/20	Case study: GPU-based graph processing		
4/27	Machine learning acceleration and system (1)		
4/29	Machine learning acceleration and system (2)		
5/6	**Final exam (2-4pm)**		**Final report due on 5/12**

EE/CSCI 451 Spring 2020
Tentative Course Schedule: A Weekly Breakdown for Lab

Date	Topics	Programming assignment	Due
1/22	Discovery cluster programming environment Discuss single-threaded program basics	Assign PA1: basic C/C++, single-threaded programming	1/28
1/29	P-thread programming	Assign PA2: p-thread programming	2/4
2/5	Homework 1 discussion		
2/12	Homework 2 discussion		
2/19	Mid-term 1 review		
2/26	Mapreduce/Spark cluster programming	Assign PA3: Mapreduce/Spark programming	3/5
3/5	OpenMP/MPI	Assign PA4: OpenMP/MPI programming	3/14
3/19	Homework 3 discussion		
3/26	Homework 4 discussion/Mid-term 2 review		
4/2	CUDA programming discussion (1)	Note: CUDA related programming assignments will be given after class. PA5: problems after class on 3/16 and 3/25 PA6: problems after class on 3/30 and 4/6 PA7: problems after class on 4/8 and 4/13	PA5: 3/31 PA6: 4/10 PA7: 4/17
4/9	CUDA programming discussion (2)		
4/16	Final project presentation (1)		
4/23	Final project presentation (2)		
4/30	Final project presentation (3)		

Note: The above are tentative outlines for Lectures and Lab sessions. Based on the number of students enrolled and the student interest, I expect to make some changes including schedule for the midterm exams, some project materials to be covered in the lecture or in the lab session, and schedule for various topics to be covered in the lecture and lab sessions.

Statement on Academic Conduct and Support Systems

Academic Conduct

Plagiarism – presenting someone else's ideas as your own, either verbatim or recast in your own words – is a serious academic offense with serious consequences. Please familiarize yourself with the discussion of plagiarism in *SCampus* in Part B, Section 11, "Behavior Violating University Standards" https://policy.usc.edu/student/scampus/part-b. Other forms of academic dishonesty are equally unacceptable. See additional information in *SCampus* and university policies on scientific misconduct, http://policy.usc.edu/scientific-misconduct.

Discrimination, sexual assault, intimate partner violence, stalking, and harassment are prohibited by the university. You are encouraged to report all incidents to the *Office of Equity and Diversity*/*Title IX Office* http://equity.usc.edu and/or to the *Department of Public Safety* http://dps.usc.edu. This is important for the health and safety of the whole USC community. Faculty and staff must report any information regarding an incident to the Title IX Coordinator who will provide outreach and information to the affected party. The sexual assault resource center webpage http://sarc.usc.edu fully describes reporting options. Relationship and Sexual Violence Services https://engemannshc.usc.edu/rsvp provides 24/7 confidential support.

Support Systems

A number of USC's schools provide support for students who need help with scholarly writing. Check with your advisor or program staff to find out more. Students whose primary language is not English should check with the *American Language Institute* http://ali.usc.edu, which sponsors courses and workshops specifically for international graduate students. *The Office of Disability Services and Programs* http://dsp.usc.edu provides certification for students with disabilities and helps arrange the relevant accommodations. If an officially declared emergency makes travel to campus infeasible, *USC Emergency Information* http://emergency.usc.edu will provide safety and other updates, including ways in which instruction will be continued by means of Blackboard, teleconferencing, and other technology.