

CS 8803: Memory Models

David Devecsery

Fall, 2018

E-mail: ddevec@gatech.edu

Web: http://classes.cc.gatech.edu/AY2019/cs8803fmm_fall/

Office Hours: By Appointment

Class Hours: M/W 4:30-5:45pm

Office: 3338 KACB

Class Room: Howey (Physics) S104

Course Description

How do we manage our ever increasing use of parallel and shared data? As computer systems have evolved we have gone from an era of single-core monolithic tasks to an era in which parallelism is driving software scalability, operating on several cores in a desktop machine, or possibly thousands of cores in cloud environments. In an era where complex tasks must share data, how do we dictate what happens to concurrent accesses of that data? What operations and semantics are allowed on shared memory? These problems are dictated by the memory model of the system.

Memory models present a unique challenge in computer science. They are necessary to understand and bound how parallel systems operate over shared data; however, they must also balance usability and understand-ability to programmers, with inherent performance trade-offs that stem from the underlying system's construction.

Memory models are by their nature a full software stack problem, with foundations in architecture, compilers, programming languages, and runtime systems. This course explores the full stack of existing memory models. We begin with the architectural foundations behind today's low-level memory models. We work our way up the stack, exploring how memory models effect compiler optimizations and language construction. We also explore proposed language-level and runtime systems proposed to enable higher-level, more understandable runtime systems. Finally, we venture into non-conventional memory models. We look at how distributed systems change how data is shared across wide networks, and look at how lack of specification of timing information has allowed attackers to exploit shared data to discover other user's secrets.

Prerequisites

As memory models are by their nature a full-stack problem, this course will cover a wide range of topics, including software systems, compilers, and computer architecture. Students are not

expected to have complete mastery over all of these topic areas, and background will be provided for more advanced issues in each of these areas. Students should have knowledge in the following:

Required:

- Software Systems (CS 6200) or equivalent: Memory models exist to define how multiple threads can interact through shared memory. Student's should be comfortable writing multi-threaded programs, and should understand the challenges of synchronizing between them.

Useful, but not required:

- Computer Architecture (CS 6290): The properties of the architectural memory system are intrinsic to the memory-model actually visible to programmers. Students do not need to be architects to take this course and more advanced topics will be reviewed, but they should at minimum understand pipeline processors, and caching systems. Knowledge of more advanced topics (out-of-order processing, cache coherence protocols, load/store buffering) is beneficial, but will be reviewed.
- Compilers (CS 6241): Compilers convert from a high-level description of a program, down to lower-level. In order to guarantee a program meets a certain memory model the compiler must follow a set of rules. Knowledge of how compilers transform programs is beneficial to understanding the challenges they face with respect to memory models. Necessary background material will be covered.

Course Objectives

The goal of this course is two fold. First, it will present students with the prerequisite background knowledge to understand and begin to think deeply about the challenges associated with the area of memory models. Second, the course will teach students to begin doing research in the area of memory models. Students are expected to grow in three primary areas required to do research in memory models:

- **Background Knowledge** — Knowledge about the limitations and trade-offs in architecture, compilers, and software systems that form the basis which causes challenges in memory models to exist today.
- **Research Knowledge** — Knowledge about the current state-of-the-art related to memory models, and interesting open and unanswered questions in the area.
- **Research Ability** — Ability to perform a significant research effort in the area of memory models. This includes identifying an open problem, making a hypothesis, then identifying and performing experiments to evaluate that hypothesis.

Assessments and Exams

To evaluate student growth, several assessments will be given over the course of the semester.

Classroom Attendance / Participation

Attendance in class, and involvement in discussions will help demonstrate student's knowledge of research in the area. Students are expected to attend and participate in class each class period. Attendance and participation for the first day of courses will be judged based on the syllabus review quiz found on [canvas](#).

I understand that sickness, travel, or other concerns may cause students to miss class. Therefore students are given **3 excused absences** in which their absence will not negatively impact their grade. If a student must be absent for an extended period of time, this course will adhere to the student absence policy of the institute (<http://www.catalog.gatech.edu/policies/student-absence-regulations/>).

Midterm

An in-class midterm exam will help evaluate student's background and research knowledge in the area. This exam will be administered **in class Monday October 15th**. If students have a conflict, or suspect they may have a conflict with the exam period they are expected to notify the instructor as soon as possible so that an alternative can be arranged.

Paper Reviews

The weekly reading assignments will be posted on the [course website](#) every week by at minimum the week before the reading is due. Readings will typically consist of 2-4 papers per week. Reading reviews must be submitted before the start of class the day the paper to-be-read will be discussed. Reviews should be one-half to one page in length, and should typically focus on following topics:

- The novelty and significance of the research presented in the paper
- The contributions of the paper
- Any particular strengths or weaknesses of the paper
- Anything else related the student thinks interesting and would like to associate with the paper

Paper reviews will not be accepted after the topic is covered in class.

I understand that sickness, travel, or other concerns may cause students to be unable to consistently review papers. Therefore students **2 worst review scores** will be dropped. If a student has difficulty completing their paper reviews for an extended period of time, this course will adhere to the student absence policy of the institute (<http://www.catalog.gatech.edu/policies/student-absence-regulations/>).

Research Project

Students will conduct a significant group research project for roughly 9 weeks during the semester. It will have multiple graded components, including:

- **Proposal** — Initial proposal of the research project, including group mates, hypothesis, and intended evaluation methods
- **Progress Reports** — Periodic progress reports discussing how the research is progressing, and how the project has changed or evolved based on findings.
- **Written Report** — Final written report summarizing the project.
- **Project Presentation** — Group presentation, presenting the project to the course instructor.

More details of the project will be distributed as the project approaches.

Final Exam

This course has no final exam.

Grading Policy

The following graded assignment will contribute to the student's final grade.

- 10% – Attendance / Participation
- 15% – Midterm
- 15% – Paper Reviews
- 10% – Presentations
- 10% – Project Proposal
- 5% – Project Progress Reports
- 25% – Project Written Report
- 10% – Project Presentation

Course Policies

Diversity and Disability Statement

Georgia Tech values diversity and inclusion; we are committed to a climate of mutual respect and full participation. Our goal is to create learning environments that are usable, equitable, inclusive and welcoming. If there are aspects of the instruction or design of this course that result in barriers to your inclusion or accurate assessment or achievement, please notify the instructor as soon as possible. Students with disabilities should contact the Office of Disability Services to discuss options of removing barriers in this course, including accommodations. ODS can be reached at 404.894.2563, dsinfo@gatech.edu, or disabilityservices.gatech.edu.

Access and Accommodations

Your experience in this class is important to me. If you have already established accommodations with the Offices of Disability Services, please communicate your approved accommodations to me at your earliest convenience so we can discuss your needs in this course.

If you have not yet established services through Disability Services, but have a temporary health condition or permanent disability that requires accommodations (conditions include but not limited to; mental health, attention-related, learning, vision, hearing, physical or health impacts), please contact the Office of Disability Services at 404.894.2563 or dsinfo@gatech.edu or disabilityservices.gatech.edu.

Disability Services offers resources and coordinates reasonable accommodations for students with disabilities and/or temporary health conditions. Reasonable accommodations are established through an interactive process between you, your instructor(s) and Disability Services. It is important to the Georgia Institute of Technology to create inclusive and accessible learning environments consistent with federal and state law.

Honor Code

Students are expected to abide by the Georgia Tech Academic Honor Code (honor.gatech.edu). Honest and ethical behavior is expected at all times. All incidents of suspected dishonesty will be reported to and handled by the office of student affairs. You will have to do all assignments individually unless explicitly told otherwise. You may discuss with classmates but you may not copy any solution (or any part of a solution).

Schedule

The most recent version of the course schedule may be found on the [course website](#).