

ECE 580B9 Course Syllabus: Fall 2019

Manycore Systems Design using Machine Learning

Instructor: Dr. Ryan G. Kim

Office Location: Engineering C201G

Email: Ryan.G.Kim@colostate.edu

Office Hours: Fri 9:30 – 10:30 AM or by appointment

Class Lectures: Tuesdays and Thursdays 11:00 AM – 12:15 PM; Scott 231

Course Assistant: Sumant Sakhalkar (Email: Sumant.Sakhalkar@colostate.edu). Office hours: TBD

Required Text: None. Research papers will be assigned throughout the semester.

Useful Materials:

- John L. Hennessy and David A. Patterson, “Computer Architecture: A Quantitative Approach,” ISBN: 978-0128119051, Morgan Kaufmann 2017.
- Hal Daumé III, “A Course in Machine Learning,” [online] <http://cml.info/>

Course Description: Fundamentals of manycore system design and electronic design automation (EDA). Study the design problems created by increased complexity and specialization of modern manycore systems. Exploration of traditional solutions, their deficiencies in this manycore era, and how machine learning can be utilized to tackle these problems.

Prerequisites: ECE452 – Computer Organization and Architecture OR CS470 – Computer Architecture (Note: CS445 – Introduction to Machine Learning is very useful for this course but is NOT a mandatory prerequisite). Python programming skills.

Grading:

The following is the grading breakdown for each major component:

- | | |
|------------------------|-----|
| - Homework: | 40% |
| - Reading Assignments: | 10% |
| - Class Participation: | 10% |
| - Paper Presentation: | 15% |
| - Midterm: | 10% |
| - Comprehensive Final: | 15% |

Submission: All assignments should be submitted before the deadline. However, late submissions will be accepted at a penalty of 10% for each day late up to a maximum penalty of 50%. For example, the maximum score for two days late is 80% while the maximum score for 5+ days late is 50%.

Collaboration: You are expected to work on all homework problems and labs yourself (or within your team), but *reasonable* collaboration is allowed and encouraged. Copying software from another student is not allowed. No collaboration will be allowed on any exam.

Grading Scale:

While lower cutoffs may be used, the following grade cut-offs are guaranteed:

A	A-	B+	B	B-	C+	C	D
90	86	83	80	76	73	70	60

Academic Integrity:

This course will adhere to the CSU Academic Integrity Policy as found in the General Catalog (<http://www.conflictresolution.colostate.edu/academic-integrity>) and the Student Conduct Code (<http://www.conflictresolution.colostate.edu/conduct-code>). At a minimum, violations will result in a grading penalty in this course and a report to the Office of Conflict Resolution and Student Conduct Services.

All submitted work should be your own. Copying of language, structure, images, ideas, or thoughts of another, and representing them as one's own without proper acknowledgement (from web sites, books, papers, other students, solutions from previous offerings of this course, etc.) and failure to cite sources properly is not acceptable. Sources must always be appropriately referenced, whether the source is printed, electronic, or spoken. My policy is that of **zero tolerance**. Minor first infraction in HWs and presentations will lead to a zero score as well as one letter level (e.g. A to B) reduction in the course grade. Project or Major or repeated infractions in HWs and presentations will result in "F" grade for the course as well as reporting to the Dean's Office.

Tentative Course Calendar¹

Week	Date	Lecture Content
1	Aug 26 - Sept 1	Introduction to manycore systems
2	Sept 2 - Sept 8	Fundamentals of manycore systems; Design challenges of Networks-on-Chip, 3D architectures, and Heterogeneous architectures
3	Sept 9 - Sept 15	Overview of manycore system design space exploration (DSE) Part 1: interconnects and layout
4	Sept 16 - Sept 22	Fundamentals of supervised learning and clustering
5	Sept 23 - Sept 29	Local search and genetic algorithms for DSE; Multi-objective optimization; Supervised learning for DSE
6	Sept 30 - Oct 6	Overview of DSE Part 2: hyper-parameter optimization for architecture and software
7	Oct 7 - Oct 13	Fundamentals of Bayesian optimization; Bayesian optimization for DSE
8	Oct 14 - Oct 20	Midterm
9	Oct 21 - 27	Overview of dynamic on-chip resource management
10	Oct 28 - Nov 3	Overview of dynamic on-chip resource management (Cont)
11	Nov 4 - Nov 10	Fundamentals of Markov decision processes, reinforcement learning, and imitation learning
12	Nov 11 - Nov 17	Voltage frequency islands (VFI); Control theoretic approaches and Machine learning for dynamic VFIs; Thermal management
13	Nov 18 - Nov 24	Domain-specific architectures; Future directions
14	Nov 25 - Dec 1	Fall Break
15	Dec 2 - Dec 8	Paper Presentations
16	Dec 9 - Dec 15	Wrapping up and final exam review
17	Dec 18 @ 6:20PM	Final Exam

¹ This schedule may be revised during the term.