Advanced Computer Architecture I
ECE 552 / CPS 550

Meeting Time & Location
TuTh 1:25 – 2:40PM, Hudson 208

Professor Benjamin Lee
Office: 210 Hudson Hall
Office Hours: Tu 4-5pm, F 4-5pm
benjamin.c.lee@duke.edu

Teaching Assistant: Sandeep Agrawal
Office: TBD
Office Hours: Tu 12:25-1:25pm, W 4-5pm
sandeep@cs.duke.edu

Teaching Assistant: Marisabel Guevara
Office: TBD
Office Hours: M 4-5pm; Th 2:40-3:40pm
mg@cs.duke.edu

Webpage
http://www.duke.edu/~BCL15/class_ece552fall12.html

Synopsis
This course covers fundamental aspects of advanced computer architecture design and analysis. Topics include processor design, superscalar and out-of-order execution, caches and memory hierarchies, virtual memory, storage systems, simulation techniques, technology trends, and future challenges. Students will complete a collaborative research project. After completing this course, students should be able to

- Understand modern processor architectures.
- Define and evaluate design metrics.
- Define and execute a research project.

Final project and paper required. Appropriate for graduate students, as well as advanced undergraduate students.

Prerequisites
Computer Science 104, or Electrical and Computer Engineering 152, or equivalent.

Text
Hennessy and Patterson. "Computer Architecture: A Quantitative Approach" 5th Edition, 2012.

Grading
Homework (30%), Midterm Exam (15%), Final Exam (25%), Project/Paper (30%)

Academic Policy
University policy as codified by the Duke Undergraduate Honor Code will be strictly enforced. Zero tolerance for cheating and/or plagiarism. If a student is suspect of academic dishonesty (e.g., cheating on an exam, copying a lab report, collaborating inappropriately on an assignment), faculty are required to report the matter to the Office of Student Conduct. A student found responsible for academic dishonesty faces formal disciplinary action, which may include suspension. A student suspended twice for academic dishonesty automatically faces a minimum 5-year separation from Duke University.

Late Policy
Late homework (except with Dean's excuses) will be penalized by 50% if less than a day late and receive zero credit if more than a day late. No late projects will be accepted.

	Topic	Text 5th Ed.	Text 4th Ed.
28 Aug 30 Aug	Introduction Early machines	1, A	1, B
4 Sep 6 Sep	Microcoding CISC to RISC Homework #1 Due		
11 Sep 13 Sep	Paper discussion I Pipelining I	C.1-C.3	A.1-A.3
18 Sep 20 Sep	Pipelining II Instruction-level parallelism I	C.4-C.6 3.1-3.10	A.4 A.5-A.8
25 Sep 27 Sep	Instruction-level parallelism II Instruction-level parallelism III Homework #2 Due	3.1-3.10 3.1-3.10	2 3.1-3.4
2 Oct 4 Oct	Paper discussion II Midterm exam		
9 Oct 11 Oct	Memory I Memory II	B.1-B.3, 2.1-2.3	C.1-C.3
16 Oct 18 Oct	Fall Break Virtual memory Homework #3 Due	B.4-B.5	C.4-C.8
23 Oct 25 Oct	Paper discussion III VLIW	3.2, 3.7	G
30 Oct 1 Nov	Vectors Multi-threading	4.1-4.3 (G) 3.12	F 3.5-3.9
6 Nov 8 Nov	Multiprocessors Homework #4 Due Paper discussion IV	5.1-5.6	4
13 Nov 15 Nov	Advanced topics – technology Advanced topics – specialization		
20 Nov 22 Nov	Advanced topics – datacenters Thanksgiving	6	
27 Nov 29 Nov	Paper discussion IV Homework #5 Due Summary		

Research Readings

Paper Discussion I	• Hill et al. "Classic machines: Technology, implementation, and economics" • Moore. "Cramming more components onto integrated circuits" • Radin. "The 801 minicomputer" • Patterson et al. "The case for the Reduced Instruction Set Computer" • Colwell et al. "Instruction sets and beyond: Computers, complexity, controversy"
Paper Discussion II	• Srinivasan et al. "Optimizing pipelines for power and performance" • Mahlke et al. "A Comparison of Full and Partial Predicated Execution Support for ILP Processors" • Palacharla et al. "Complexity-Effective Superscalar Processors" • Yeh et al. "Two-Level Adaptive Training Branch Prediction"
Paper Discussion III	• Jouppi. "Improving direct-mapped cache performance by the addition of a small fully-associative cache and prefetch buffers" • Kim et al. "An adaptive, non-uniform cache structure for wire-delay dominated on-chip caches" • Fromm et al. "The Energy Efficiency of IRAM Architectures" • Lee et al. "Phase change memory architecture and the quest for scalability"
Paper Discussion IV	• Mudge, "Power: A first-class architectural design constraint" • Lamport. "How to make a multiprocessor computer that correctly executes multiprocess programs" • Lenoski et al. "The Stanford DASH Multiprocessor" • Tullsen et al. "Simultaneous multithreading: Maximizing on-chip parallelism"
Paper Discussion V	• Horowitz et al. "Scaling, power, and the future of CMOS" • Reddi et al. "Web search using mobile cores: Quantifying and mitigating the price of efficiency" • Dally et al. "Efficient Embedded Computing"