CS 258 Parallel Processing
Computer Science Division
Department of Electrical Engineering and Computer Sciences
University of California, Berkeley

Spring 1995 D. E. Culler
Handout 1

Course Information

Instructor:	David E. Culler, Associate Professor
Office:	627 Soda Hall
Phone:	643-7572
E-mail:	culler@cs
Office Hours:	W 11-12, Th 2:30 - 3:30 or by appointment
Prereq:	cs 252 (cs 267 advised)
Meetings:	W, F 9:30 - 11, 405 Soda (note change)

The goal of this course is to provide you with an in-depth understanding of modern parallel computer architecture. This includes a background in the historical evolution of the field, the role of technology developments in shaping this evolution, and the essential interplay of hardware and software. We will seek to understand the fundamental design issues and engineering trade-offs that cut across the diverse spectrum of current and past parallel machines, for these are the principles that will guide future development. Today the field is in the midst of a dramatic and exciting restructuring as previous disparate design styles are converging on a largely common organizational structure. This process is driven by important technological forces, maturing parallel software support, and better understanding of the hardware design trade-offs. This provides the opportunity to attempt to build a systematic understanding of the field and to delve deeply into the design trade-offs, rather than simply offer a descriptive taxonomy.

The course will be based on lectures and readings from a not-yet-alpha version of an upcoming text under development by the instructor, Anoop Gupta of Stanford, and J. P. Singh of Princeton. There will be addition readings from the primary literature. You will get to help shape the book through your feedback, ideas, criticism, confusion, assignments, projects, etc. A preliminary lecture schedule follows. The readings refer to the attached table of contents. (This is a target for both of us.) My intention is that you do the reading prior to class. I will go over the material to bring out the highlights and provide structure for discussions.

My intention is for this to be an interactive, advanced graduate seminar. So enrollment will be limited and prereqs will be enforced. Within the general framework of the content on the schedule, there will be substantial opportunities for class participation. As there are no TAs, you will all get to assist in class administration. There will be an optional group discussion section at a convenient time. We will resolve the grading format and the mix of assignments, projects, and exams on the first meeting.

	Date	Lecture Topic	Reading
W	1/18/95	Why Parallel Architectures	1.1-3
M	1/23/95	Current Parallel Machines	1.4
W	1/25/95	Parallel Programs	2.1-3
M	1/30/95	Programming for Performance	2.4-5
W	2/1/95	Application Case Studies	2.6-7
M	2/6/95	Performance Evaluation	2.8-9
W	2/8/95	Cache Coherency	3.1-3.4
M	2/13/95	Snooping Protocols	3.4
W	2/15/95	Implementation Issues	3.5
M	2/20/95	Design Trade-offs and Case Studies	3.6-3.9
W	2/22/95	Large-scale Machines	4.1-2
M	2/27/95	User-level Networks	4.4-4.6
W	3/1/95	Coprocessors and virtual DMA	4.7-4.8
M	3/6/95	Shared Physical Address	4.9-10
W	3/8/95	Snooping with Directories	5.1-4
M	3/13/95	Case Studies and COMA	5.4-6
W	3/15/95	Software Shared Address Space	5.7-8
M	3/20/95	*Holiday*	
W	3/22/95	Memory Consistency Models	6.1-2
MW	3/27-9/95	*Spring Recess*	
M	4/3/95	Dynamically Scheduled Processors	6.3
W	4/5/95	Prefetching and Mutiple Contexts	6.4-6
M	4/10/95	Network Design Space	7
W	4/12/95	k-ary n-cubes	7
M	4/17/95	Performance under Contention	7
W	4/19/95	Synchronization	8
M	4/24/95	HW / SW Trade-offs	8
W	4/26/95	Dataflow	9
M	5/1/95	SIMD	9
W	5/3/95	Systolic Arrays	9
M	5/8/95	Wrap-up	