

15-495: Parallel Computer Architecture and Programming

Fall 2002

Syllabus

1 Course Details at a Glance

Lectures:	Tuesdays, 12:00-1:20 p.m., BH 237B Wednesdays, 10:30-11:20 a.m., SH 324 Thursdays, 12:00-1:20 p.m., BH 237B
Instructor:	Todd C. Mowry, WeH 8123, 268-3725, tcm@cs.cmu.edu <i>Office Hours:</i> Wednesdays, 3:00-4:00pm, WeH 8123
TA:	Ryan Gallagher, ryan3@andrew.cmu.edu <i>Office Hours:</i> Mondays, 2:00-3:00pm, Wean Computer Cluster
Class Admin:	Jennifer Landefeld, WeH 8124, 268-4740, jennsbl@cs.cmu.edu
Web Page:	www.cs.cmu.edu/afs/cs/academic/class/15495-f02/www/
Newsgroup:	cyrus.academic.cs.cs495
Handouts:	<i>Electronic:</i> /afs/cs.cmu.edu/academic/class/15495-f02/public <i>Hardcopies:</i> In bins outside WeH 8124.

2 Textbook

The following textbook is required for the course:

- David E. Culler and Jaswinder Pal Singh, with Anoop Gupta. *Parallel Computer Architecture: A Hardware/Software Approach*, Morgan Kaufmann, 1998. ISBN: 1558603433.

We will be following this book quite closely, and we will covering much of the material in the book.

3 Course Description

The goal of this course is to provide a deep understanding of the fundamental principles and engineering tradeoffs involved in designing modern parallel computers (aka “multiprocessors”), as well as the programming techniques to effectively utilize these machines. Parallel machines are already ubiquitous from desktops to supercomputers, and the expectation is that they will become even more commonplace in the future. However, very few people exploit the potential processing power of these machines because they do not understand how to write efficient parallel programs. Because one cannot design a good parallel program without understanding how parallel machines are built and vice-versa, this course will cover both parallel hardware and software design, as well as the impact that they have on each other.

Course topics include naming shared data, synchronizing threads, and the latency and bandwidth associated with communication. Case studies on shared-memory, message-passing, data-parallel and dataflow machines will be used to illustrate these techniques and tradeoffs. Programming assignments will be performed on one or more commercial multiprocessors, and there will be a significant course project.

This is a relatively unique course since this material is rarely offered to undergraduates. Because parallel processing has become such an important and mainstream technology, the time has come to integrate this material into the undergraduate systems curriculum. This is the second time that this course has been offered at Carnegie Mellon (it was also offered this past Spring).

4 Prerequisites

15-213 (Intro to Computer Systems) is a strict prerequisite for this course. We will build directly upon the material presented in 15-213, including memory hierarchies, memory management, basic networking, etc.

While 18-347 (Intro to Computer Architecture) would be helpful for understanding the material in this course, it is not a prerequisite.

5 Computer Accounts

To complete your programming assignments and course projects, you will be receiving accounts on machines at the National Center for Supercomputing Applications (NCSA) and the Pittsburgh Supercomputing Center (PSC). Details will be provided later.

Important: please note that the class will be allocated a finite (and not particularly large) amount of time on these machines, so please be careful not to waste time unnecessarily.

6 Course Work

Grades will be based on homeworks, a project, two exams, and class participation.

Homeworks: There will be roughly three parallel programming assignments (which we will call “labs”), which you will work on in groups of two. (If you have difficulty locating a partner, please post a message to the class newsgroup.) Turn in a single writeup per group. In addition, we may have a few written assignments that will focus more on parallel architecture rather than programming.

Project: A major focus of this course is the project. We prefer that you work in groups of two on the project, although groups of up to three may be permitted depending on the scale of project (ask the instructor for permission before forming a group of three). A typical project would involve designing, implementing and evaluating a fairly ambitious parallel program (perhaps on more than one architecture). Some groups may choose to do projects that evaluate a new parallel architecture idea. The project must involve an experimental component—i.e. it is not simply a paper and pencil exercise. We encourage you to try to come up with your own topic for your project (subject to approval by the instructor), although we can make some suggestions if necessary. You will have roughly six weeks to work on the project. You will present your findings in a written report (the collected reports may be published as a technical report at the end of the semester), and also during a poster session during the last day of class. Start thinking about potential project ideas soon!

Exams: There will be two exams, each covering its respective half of the course material. Note that the second exam is not cumulative, and is weighted equally with the first exam. Both exams will be closed book, closed notes.

Class Participation: In general, we would like everyone to do their part to make this an enjoyable interactive experience (one-way communication is no fun). Hence in addition to attending class, we would like you to actively participate by asking questions, joining in our discussions, etc.

6.1 Grading Policy

Your overall grade is determined as follows:

Homework:	25%
Project:	25%
Exams:	40% (20% each)
Class Participation:	10%

Late assignments will not be accepted without prior arrangement.

7 Schedule

Table 1 shows the tentative schedule. The idea is to cover the lecture material in roughly the first $\frac{2}{3}$ of the semester (by meeting three rather than two days a week), so that you will have more time to devote to the class project in the last $\frac{1}{3}$ of the semester, and so that you can take advantage of all of the course lecture material in your projects.

Slack Days: Whenever possible, we will meet on Wednesdays and have a regular lecture. In case it is necessary to cancel a class here or there, I have built some “slack” into the schedule so that we won’t have to reschedule the exam dates. We will keep the schedule on the class web page updated as we go along.

Table 1: 15-495, Fall 2002.

Class	Date	Day	Topic	Reading	Assignments
1	8/27	Tue	Why Study Parallel Architecture?	1.1	
2	8/28	Wed	Evolution of Parallel Architecture	1.2	
3	8/29	Thu	Fundamental Design Issues	1.3-4	
4	9/3	Tue	Parallel Programming: Overview I	2.1-2	L1 Out
5	9/4	Wed	Parallel Programming: Overview II	2.3-4	
6	9/5	Thu	Parallel Programming: Performance I	3.1	
7	9/10	Tue	Parallel Programming: Performance II	3.2	
8	9/11	Wed	Parallel Programming: Performance III	3.3-4	
9	9/12	Thu	Par. Prog: Case Studies & Implications	3.5-6	L2 Out
10	9/17	Tue	Workload-Driven Arch Evaluation I	4.1	L1 Due
11	9/18	Wed	Workload-Driven Arch Evaluation II	4.2-3	
12	9/19	Thu	Shared Memory Multiprocessors I	5.1	
13	9/24	Tue	Shared Memory Multiprocessors II	5.3	
14	9/25	Wed	Shared Memory Multiprocessors III	5.4	
15	9/26	Thu	Directory-Based Cache Coherence I	8.1-5	L2 Due, L3 Out
16	10/1	Tue	Directory-Based Cache Coherence II	8.6-7, 8.9-11	
	10/2	Wed	Relaxed Memory Consistency Models	9.1	
17	10/3	Thu	Snoop-Based Multiprocessor Design I	6.1	
	10/8	Tue	<i>No Class</i>		
	10/9	Wed	<i>No Class</i>		
18	10/10	Thu	Snoop-Based Multiprocessor Design II	6.2	L3 Due
	10/15	Tue	Exam I		
19	10/16	Wed	Snoop-Based Multiprocessor Design III	6.3-4	
20	10/17	Thu	Snoop-Based Multiprocessor Design IV	6.5, 6.7	
21	10/22	Tue	Synchronization	5.5., 7.9, 8.8	Project Proposal
22	10/23	Wed	Earthquake Simulation Case Study		
23	10/24	Thu	Scalable Distributed Memory MPs I	7.1-3	
24	10/29	Tue	Scalable Distributed Memory MPs II	7.4-8	
25	10/30	Wed	Interconnection Network Design	10.1-10	
26	10/31	Thu	Latency Tolerance: Prefetching	11.1, 11.6	
27	11/5	Tue	Latency Tolerance: Multithreading	11.7-9	Project Milestone 1
28	11/14	Thu	Terascale Computing System at PSC		
	11/19	Tue			Project Milestone 2
	11/21	Thu	Exam II		
	12/3	Tue			Project Due
	12/5	Thu	Project Poster Session		