

Basics of Parallel and Distributed Systems
CIS 4930.008 and CIS 6930.007
Spring 2009

General Information

Class meetings: Tuesdays/Thursdays 3:30-4:45 pm in NES 108

Professor: Adriana (Anda) Iamnitchi (anda@cse.usf.edu)

Office: ENB 334

Phone: 813-974-5357

Office hours: Wednesdays 10am-12pm and other times by appointment

Class Overview

This class introduces the basics of parallel computing and distributed systems. The course will cover knowledge on two plans: theoretical and practical. On the theoretical realm, we will discover the basic principles that are at the foundation of parallel programming and distributed systems. On the practical realm, we will apply the theoretical foundations in developing programming projects in different parallel and distributed environments.

The theoretical side of the parallel computing sections of this course will cover how to think of parallel programs, how to approach the design of parallel programs, performance metrics and limitations. This material will be accompanied by hands-on mini-projects that will introduce students to the practice of parallel programming with POSIX threads, Java threads, MPI, and Google's MapReduce.

On the distributed systems front we will discuss topics such as distributed systems architectures, time and global states, coordination and agreement, distributed transactions and web services. The hands-on experimentation will be based on the PlanetLab testbed and on Amazon Web Services (if time allows).

This course is intended for advanced undergraduate students and graduate students in CS. Graduate students from the College of Engineering might find relevant topics to their own research (if involved in computational efforts).

Prerequisites: for everybody: comfortable with programming in C and/or Java.

Undergraduate students: COP 4600 (Operating Systems)

Graduate students in CS: none

Graduate students outside CS: instructor's permission (email anda@cse.usf.edu)

Required textbook:

Distributed Systems: Concepts and Design, George Coulouris, Jean Dollimore, Tim Kindberg. 4th Edition. Addison-Wesley, 2005. ISBN: 0-321-26354-5 [CDK]

Other textbooks (recommended):

Patterns for Parallel Programming, Timothy G. Mattson, Beverly A. Sanders, Berna L. Massingill. Addison-Wesley Professional; 1 edition (September 25, 2004). ISBN: 0321228111 [MSM]

Parallel Programming: Techniques and Applications Using Networked Workstations and Parallel Computers, by Barry Wilkinson, Michael Allen. Prentice Hall; 2nd edition (March 14, 2004), ISBN: 0131405632 [WA]

Principles of Parallel Programming, by Calvin Lin and Larry Snyder, Addison-Wesley; 1st edition (March 7, 2008). ISBN: 0321487907 [LS]

Tentative Schedule

Week	Course Material	Hands-on Projects
1	Background and Terminology for Parallel Programming [MSM 2]	POSIX threads and Java threads
2	Parallel Abstractions. Parallel Programming Languages [LS 5]	
3	Message Passing Interface [LS 7; WA 2; MSM App B]	MPI
4-5	MapReduce [OSDI 2004 MapReduce paper]	Hadoop/MapReduce
6	Characterization of Distributed Systems. System Models. [CDK1-2]	
7	Distributed Objects and Remote Invocation [CDK5]	
8	Operating Systems Support [CDK6]	
9	Distributed File Systems [CDK8]	PlanetLab
10	Name Services [CDK9]	
11	Time and Global State [CDK11]	
12	Coordination and Agreement [CDK12]	
13	Transactions [CDK13-14]	
14	Web Services [CDK19]	Amazon Web Services
15	Final project presentations	

Grading and Attendance

There will be frequent small assignments with questions on reading material and problems. The 5 programming projects are more involved and will require comfort with programming environments (patience, debugging, reading man pages, etc). There will be a final project that can be a continuation of one of the 5 programming projects. There will be no exams or quizzes.

Since much of the material is collected and compiled from multiple sources, I expect you will find class attendance mandatory.

Grading system:

- Assignments (15%)
- Programming projects (50%)
- Final project (25%)
- In-class contributions and attendance (10%)

For final letter grades, I will use the standard scale of A (100-90), B (89-80), C (79-70), D (69-60), and F (59-0). I will also use pluses and minuses on final grades to indicate either a borderline grade or exceptionally outstanding work (A+). Grades will be posted on Blackboard (<http://my.usf.edu/>).

Academic integrity: Everything you turn in for this class must be your own work. If you are caught cheating, you will receive an FF grade for the class. Material copied from the Web and submitted as your work *is cheating*. Please see the University's Undergraduate Catalog regarding these policies at <http://www.ugs.usf.edu/catalogs/0708/adadap.htm>.

Every part of this syllabus is subject to adjustment as the semester progresses. Please contact me as soon as possible if you have particular interests in material that is relevant to the class topic but not covered in enough detail; I will be happy to accommodate reasonable requests for modifications.