

EE/CS 674 – Heterogeneous Computing

Electrical and Computer Engineering Dept.
and Computer Science Dept.
Colorado State University

Prerequisite: Any one of: EE 550, EE 554, CS551, CS570, CS575.
Course Credits: 3

Instructor: Professor H. J. Siegel
Office: Engineering B115
Phone: 491-7982
Email: HJ@ColoState.edu

Description

In a heterogeneous computing environment, a suite of different machines is interconnected to provide a variety of computational capabilities to execute collections of application tasks that have diverse requirements. The execution times of a task will vary from one machine to the next, and tasks will compete for machines in the suite. There are many types of heterogeneous systems, including parallel, distributed, clusters, and grids. They can be found in industrial, laboratory, government, academic, and military settings. Such systems may be used in production, computing center, embedded, or real-time environments. An important research problem for heterogeneous computing is how to assign computation and communication resources to tasks and to schedule the order of their execution to maximize some performance criterion, a process known as mapping. Factors that must be considered include machine and network loading, how well the execution needs of a task match the computational capabilities of a machine, any inter-task communications, operating constraints, and the performance criterion to be optimized.

An overview of the field of heterogeneous computing will be given. Dynamic and static heuristics for mapping tasks to resources in a heterogeneous system will be presented. Tasks that involve priorities, deadlines, and alternate versions of different worths to the user will be considered. Techniques for combining tasks' execution parameters into a single performance metric will be described. Open problems in the field of heterogeneous computing will be discussed.

This course is intended for ECE and CS graduate students who want to learn about the ways in which a collection of heterogeneous machines can be used to execute a single large application task or a set of independent or interrelated application tasks in a way that will optimize some performance criterion.

Course Objectives

This course will enable the student to:

- understand the potential advantages of using heterogeneous computing systems
- analyze some of the factors that must be considered when designing resource management systems for heterogeneous environments
- be familiar with a variety of *dynamic* (on-line, real-time) and *static* (off-line) techniques for assigning resources to tasks and scheduling their execution to optimize some performance criterion
- evaluate the effectiveness of resource allocation schemes for heterogeneous systems
- formulate a performance metric appropriate for the goals of a given computing environment
- be aware of the open research problems in heterogeneous computing that are important areas for future research and development.

Course Materials

This course will be taught using technical papers from the literature.

There will be evening exams.

Course Policies (subject to change)

1. Projects

There will be two research projects, both of which may involve programming and simulation studies. Students will work in pairs on these projects.

2. Exams

- a. There will be three exams, each covering approximately one-third of the course; there will be no comprehensive final (i.e., no final exam that covers the entire course).
- b. Exams will be given in the evening (the third exam may be given in the final exam time slot).

3. Grading

- a. Your final course grade will be based on your projects and examinations. Each of the two projects will be worth 20%, and each of the three exams will be worth 20%.
- b. Your letter course grade will be determined from the total points that you obtain from your projects and tests, and will be based on a combination of a relative and an absolute scale. You determine your own grade by your performance on these items.

Course Outline

Topic	Weeks
1. Class policy, introduction to heterogeneous computing, automatic heterogeneous computing, open problems	1
2. Static mapping of applications composed of communicating tasks	1
3. Static mapping techniques for independent tasks	2
4. Assignment and description of course projects	0.5
5. Dynamic mapping techniques for independent tasks	1
6. Mapping of tasks with priorities, deadlines, and versions	1
7. Scheduling communications in overloaded networks	1
8. Project 1 presentations by students	1
9. On-line use of off-line derived mappings	0.5
10. Classification scheme for resource management techniques	0.5
11. Collective quality of service measures	0.5
12. Advanced current topics	3
13. Project 2 presentations by students	1
14. Exams	1

Additional Information for Spring 2007

Class hours: Tuesday and Thursday, 11 a.m. to 12:15 p.m.

Room: Engineering B3

Instructor Office Hours (beginning January 18, 2007):

Monday: 1:00 p.m. to 2:00 p.m.

Tuesday: 9:00 a.m. to 10:00 a.m.

Wednesday: 10:00 a.m. to 11:00 a.m.

Thursday: 2:00 p.m. to 3:00 p.m.

Friday: 1:00 p.m. to 2:00 p.m.

Please check www.engr.colostate.edu/~hj for changes to the regular schedule of office hours before coming.

You can make an appointment if you have a class conflict with the office hours.

Teaching Assistant:

Name: Luis Briceno

Office: Engineering B5

Phone: 491-3148

E-mail: ldbricen@lamar.colostate.edu

Office hours: Monday, Friday: 11:00 a.m. to 12 noon

Tuesday, Wednesday, Thursday: 1:00 p.m. to 2:00 p.m.

You can make an appointment if you have a class conflict with the office hours.

Additional Course Policies

1. Information

You are responsible for all information given in class verbally and/or in writing. All information about the course (including but not limited to grading policy, exam dates, office hours, and course schedule) may be superseded at any time by information given in class.

2. Learning the Course Material

The Instructor of this course wants you to learn the material. To accomplish this:

- a. please read the relevant assigned material before class;
- b. please attend class;
- c. please pay attention in class;
- d. please take notes in class;
- e. please ask questions in class when you do not understand the material;
- f. please talk to your classmates (outside of the class time) about the course material;

- g. please see the Teaching Assistant or Instructor during their office hours when you need help.

3. **More About Exams**

- a. Each exam will be designed to require only 50 minutes. Thus, each exam is officially only 50 minutes long. However, you will have the option to spend up to two hours on each exam, if you wish.
- b. All exams will be closed book, with no notes and no calculators. Cheating will not be tolerated. For exams, you should use only the information in your own brain. The use of any other information (e.g., reading the paper of the person next to you, accessing information stored in your electronic telephone directory wristwatch) is considered cheating. If one cheats on an exam, one's grade for that exam is zero. Other disciplinary action may also be taken.
- c. If you do not show up for an exam you will receive a zero, unless other arrangements have been made with the Instructor before the exam (see "Disclaimer" below).
- d. Exam dates:
 - #1 – Tuesday, February 13, 2007, 7 – 9 p.m. in Engineering B105
 - #2 – Tuesday, April 3, 2007, 7 – 9 p.m. in Engineering B105
 - #3 – Tuesday, May 10, 2007, 3:40 p.m. – 5:40 p.m. in Engineering B3
- e. The classes on March 27, 2007, and March 29, 2007, will be canceled because of having "out of class" evening exams for exam #1 and exam #2.

4. **More About Grading**

- a. There are *no* homework assignments other than reading the course material and doing the projects. This is subject to change.
- b. There are *no* extra credit assignments.
- c. Please see the Instructor if you feel your grade on any project or test was incorrect.
- d. Graded exams will be returned in class as soon as they are ready.
- e. Please save all of your graded projects and exams until you receive your final grade and you decide you have no questions.

5. **How To Maximize Your Ability To Get An "A" (and How Much You Learn!)**

- a. See "Learning the Course Material" above.
- b. Talk to the Teaching Assistant, Instructor, and/or classmates if you need help understanding the material.
- c. Start the projects as soon as they are assigned and work hard on them to do an excellent job.
- d. To prepare for each exam, review the class notes and reading material.

6. **Disclaimer**

Whenever there are questions relating to the course policies, the Instructor will use his judgment to decide what is fair.