CS 524 High-Performance Computing

Instructor's Name: Asim Karim

Office No. & Email: 132, akarim@lums.edu.pk

Office Hours: 2.30 PM – 4.00 TR

TA for the Course: TBA

Year: 2003-04

Quarter: Winter

Category: Graduate

Course Code (Units)

CS 524 High-Performance Computing
(3 Units)

Course Description

With advances in computer architecture, high performance serial and parallel computers have become readily available and affordable. As a result, high-performance and super computing is accessible to a large segment of the industry that was once restricted to military research and large corporations. This course introduces practical performance optimization techniques for serial and parallel computing with an emphasis on algorithms in linear algebra. Popular parallel programming models such as shared- and distributed-memory and parallel libraries such as MPI and OpenMP will be discussed. Performance evaluations will be carried out on modern PCs, LUMS's Linux Cluster, and Sun UltraSparc (suraj) server. Programming assignments will require coding in C/C++ in a Linux/UNIX environment. Students will be required to do readings before and after class from multiple sources to supplement what is discussed in class.

Core/Elective

Elective

Pre-requisites

CS 223 Computer Organization and Assembly Language, CS 213 Data Structures and Algorithms, or permission of instructor (CS 423 Computer Architecture is recommended)
Proficiency in algorithm implementation in C/C++ is required. Most implementations will be done in the Linux/UNIX environment.

Goals

- Mastery of fundamental performance issues for parallel and serial computing;
- Familiarity with shared-memory and message passing architectures and models for parallel program development;
- Hands-on experience in designing and implementing efficient algorithms for high-performance computers;
- Familiarity with the current state-of-art in parallel programming environments, portable software libraries, and program development

TextBooks, Course Materials, aetc.

There is no comprehensive text for this course. Material will be taken from several sources including the following books. Supplementary handouts will be made available to the students.

1. Introduction to Parallel Computing: Design and Analysis of Algorithms, V. Kumar, A. Grama, A. Gupta, G. Karypis, Benjamin/Cummins Publishing Co., 1994.
2. Using MPI – Portable Parallel Programming with the Message-Passing Interface, 2nd Ed., W. Gropp, E. Lusk, A. Skjellum, MIT Press, 1999.
3. Parallel Programming in OpenMP, R. Chandra et al., Academic Press, 2001.
4. Techniques for Optimizing Applications – High Performance Computing, R.P. Garg and I. Sharapov, Sun Microsystems Press, 2002.
5. Introduction to Parallel Processing, M. Sasikumar, D. Shikhare, P.R. Prakash, Prentice-Hall of India, 2000.

CS 524 High-Performance Computing

Year: 2003-04

Quarter: Winter

Lectures, Tutorials & Attendance Policy

There will be 20 sessions of 75 minutes each, one in-class midterm exam and one final exam. There are no points for attendance; however, students who miss lectures will find it very difficult to make up for the content covered, with the possible additional penalty of missing quizzes.

Grading

35%	Assignments
10%	Quizzes
25%	Midterm Exam
30%	Final Exam (Comprehensive)

Additional Details

The course website will be the primary source for announcements and reading materials including lecture slides, handouts, and web links. http://suraj.lums.edu.pk/~cs524w03

Cheating and plagiarism will not be tolerated and will be referred to the disciplinary committee for appropriate action. If an assignment is discussed among students, it is required that each student writes up the solution independently, and without looking at notes from the discussion. Downloading code segments from the internet and presenting them as your own is considered plagiarism.

CS 524 High-Performance Computing

Year: 2003-04

Quarter: Winter

Topics	Sessions	Readings
1. introduction; programming models/techniques; single/multi processor performance issues; cache tuning and characterization	1-2	1:1; 4::1; 5:1; handout;
2. single processor performance enhancements; data locality; loop optimizations; data dependency	3-4	4: 4 and (5), 8 and 9; 5:2 and 4
3. parallel architectures; taxonomy of parallel computers; parallel programming models	5-6	1:2; 5:2; and 3
4. MPP (message passing programming); communication overheads; performance and scalability of parallel algorithms	7-8	1: 3 and 4, 2:1; 4:4
5. MPI (message passing interface)	9-10	2:2 and 3; 5:7; MPI handout
MIDTERM EXAM 25 %		
6. MPI collective communications; parallel algorithms	12-13	MPI handout;; 2:3 and 4; 1:5
7. parallel algorithms (dense matrix and sparse matrix)	14-16	1:5 and 11; 2:4 and 5
8. SMP (shared memory programming), Pthreads; OpenMP	17-19	SSP handout; 5: 5 and 6; 3: all
9. overview of other parallel libraries (PVM, Linda); scientific and engineering applications of HPC; research directions; review	20	Handout

FINAL EXAM 30%