

CSCI 473

Introduction to Parallel Systems

Spring Semester, 2011
MWF 11:30AM – 12:20PM
CSCC Rm 318



CSCI 473, an introduction to parallel systems and programming, will be offered for the second time during the Spring 2011 semester. It is a three credit-hour course, and will count as an upper-level elective. In particular, **Dr. Sheel has given permission for CSCI 473 to count for EITHER a “Computational” or “Application” elective (see the typical footnotes on your advisement sheets).** Additionally, it can count as your generic “300-level or above” elective. (Note, although it can count in any of these three categories, it can only count once)

This course will be run concurrently with the same course in the Electrical and Computer Engineering department at Clemson University. We will be making use of Clemson supercomputing resources and may at times participate in a real-time grid-classroom environment with them.

Please see the following tentative syllabus for more information. If you have any questions or concerns, please do not hesitate to contact me.

Dr. William M. Jones
<http://www.parl.clemson.edu/~wjones>

11 SP CSCI 473 01 Jones

CSCI 473 -- Introduction to Parallel Systems Section 01 - Spring Semester, 2011 MWF 11:30AM - 12:20PM, CSCC Rm 318

Catalog Description:

(Prerequisites: A C or better in CSCI 220) This course introduces parallel computer architectures and their programming. It includes an introduction to MPI and OpenMP and a number of engineering problems, including numerical simulations. It also provides an introduction to performance evaluation and modeling as well as scalability analysis.

Coastal Instructor:

Dr. William M. Jones, 843-349-4142, 105 CCSC

Website = <http://www.parl.clemson.edu/~wjones>

Clemson Instructor:

Dr. Walt Ligon, 864-656-1224, 300-D Riggs Hall

Office Hours:

(for Dr. Jones) see schedule at <http://www.parl.clemson.edu/~wjones>

I may also be available for questions outside of scheduled office hours; however depending on my workload, I may have to ask you to come back during regular hours.

Office hours are subject to change without notice.

The best way to contact me is via email. I tend to check my email very regularly, so please do not hesitate to contact me. When you have questions, please be as specific as possible, send complete program sources, page reference numbers, etc.

Required Text:

Parallel Programming in C with MPI and OpenMP

By: Michael J. Quinn,

McGraw Hill ISBN: 0-07-282256-2

Grading Policy:

Grades will be assigned according to the standard 10-point grading scale with possible "+" letter grades.

There will be no curving of grades during the semester. I **MAY** decide after all grades are in at the end of the semester to curve, based on the cumulative difficulty of the material, averages, etc; however you should not assume that any curve will be given.

Grading Scale

| | | | |
|----------|----|---------|----|
| 90 - 100 | A | 70 - 76 | C |
| 87 - 89 | B+ | 67 - 69 | D+ |
| 80 - 86 | B | 60 - 66 | D |
| 77 - 79 | C+ | 0 - 59 | F |

Grading Weights:

Grades will assignments, a mid-term exam, and a final exam. Your final grade will be calculated as follows:

| | |
|---------------------------|-----|
| Mid-Term Exam: | 30% |
| Assignments and Projects: | 40% |
| Final Exam: | 30% |

As you can see, the assignments constitute a large fraction of the course grade.

Attendance:

Attendance is expected, and you are expected to be on time for class. The CCU University Catalog states, with respect to attendance, that “An instructor is permitted to impose a penalty, including assigning the grade of F, for unexcused absences in excess of 25% of the regularly scheduled class meetings.” Attendance will be taken, and absences in excess of 25% of our class meetings will result in a failing grade, no matter what your actual performance in the course happens to be. The Catalog also states that “Absences will be excused for documented cases of:

- a) incapacitating illness,
- b) official representation of the University (excuses for official representation of the University should be obtained from the official supervising the activity),
- c) death of a close relative, and
- d) religious holidays.”

In short: Attendance is required
 No late work accepted.
 Make-up tests by prior appointment only.
 Wait 15 minutes for late instructor.

E-mail Communication:

I will normally respond to e-mails within one day. I do not normally read student e-mails on the weekends, so do not expect an immediate response if an e-mail is sent then; however in some cases I will be able to respond during the weekend. University policy dictates that all e-mail communication regarding class issues be conducted with students via their Coastal e-mail address. E-mail from other addresses will not be answered.

An e-mail must have the correct course number in the subject. If you send me email without a subject, it is very likely that I will throw it away as spa.. If you send me email, use correct spelling, grammar and punctuation. Do not send email to me using instant messaging code. If you do, I will throw away your email.

Team Teaching:

This course will be taught by Dr. Ligon at Clemson University, as well as by Dr. Jones from here at Coastal. We will be making use of various video conferencing technologies as a part of an evaluation of such technologies, ranging from simple to complex. Additionally, we will a cluster computer for running and testing our programs. Important information will be delivered to you from both professors, and as such, class attendance is EXTREMELY important. It may become necessary for us to schedule a recurring out of class meeting time so that we can collectively discuss your questions and concerns, as doing so during class may not always be possible due to the video conferencing.

Academic Misconduct:

Coastal's policy on academic integrity as stated in the Student Code of Conduct:

Coastal Carolina University is an academic community that expects the highest standards of honesty, integrity and personal responsibility. Members of this community are accountable for their actions and reporting the inappropriate action of others and are committed to creating an atmosphere of mutual respect and trust.

Academic misconduct will not be tolerated, and if you are caught committing an academic infraction your action will be reported to the university and my standard sanction is failure for the course. The CCU Student Code of Conduct (URL: <http://www.coastal.edu/judicialaffairs/codeofconduct.pdf>) gives examples of plagiarism and cheating as follows:

a. Examples of plagiarism include but are not limited to the following:

- (i) Words, sentences, ideas, conclusions, examples and/or organization of an assignment are borrowed without proper acknowledgment from a source (for example, a book, article, electronic documents, or another student's paper).
- (ii) A student submits another person's work in place of his/her own.
- (iii) A student allows someone else to revise, correct, or edit an assignment without explicit permission of the instructor.
- (iv) A student submits work without proper acknowledgment from commercial firms, Web sites, fraternity or sorority files, or any other outside sources, whether purchased or not.
- (v) A student allows another person to take all or any part of a course, including quizzes, tests, and final examinations.
- (vi) A student submits any written assignments done with the assistance of another without the explicit permission of the instructor.
- (vii) A student knowingly aids another student who is engaged in plagiarism.

b. Examples of cheating include but are not limited to the following:

- (i) A student uses unauthorized information, materials or assistance of any kind for an assignment, quiz, test, or final examination.
 - (ii) A student knowingly aids another student who is engaged in cheating.
- See the Code of Conduct for more details, as well as other cases of academic misconduct. The simplest rule of thumb here is this: Do your own work, and give properly formatted credit for ideas that aren't your own.

All work on quizzes, tests, exams, design assignments, projects, and labs is to be wholly your own. Possessing, using, providing, or exchanging improperly acquired written, verbal, or electronic information will be considered a violation of the academic honor code. Again, violations will result in a grade of F for the semester.

Linux:

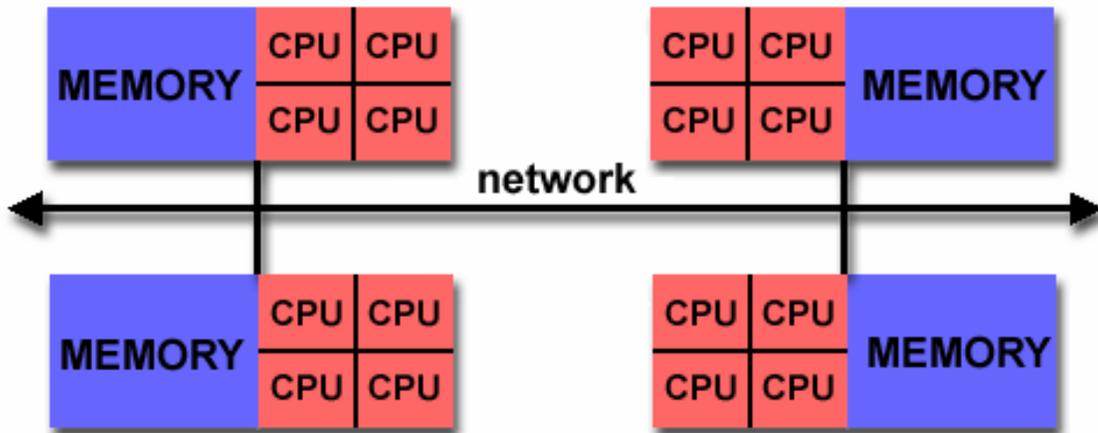
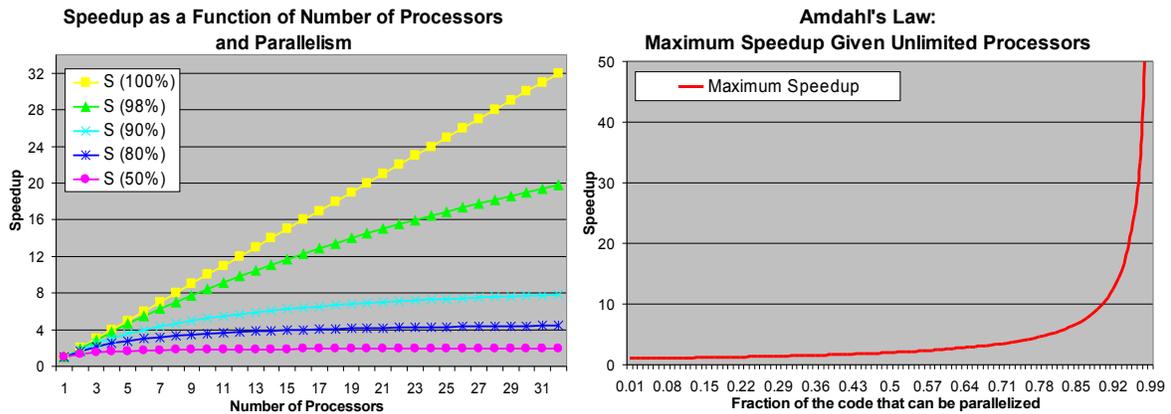
The axiom cluster is a 100 processor linux cluster computer, and we will be using SSH to remotely connect to this machine. I have posted an SSH/ file transfer client on the BB9 server at Clemson. You will be expected to understand how to navigate the commandline in Linux, programming and compiling such programs, and you will have to learn to use the PBS queueing system there.

C Programming:

This course assumes you know how to program in C. If you are not strong in this area, you will likely need to spend considerable time outside of class getting up to speed on this. It might be a good idea to get a good reference / textbook on C programming to help you through this phase as would researching and using tutorials online. Please let me know how I can help.

MPI Programming:

This course will teach you how to use the MPI programming library to implement parallel programs. Our textbook has a good coverage of this, and there are numerous examples online as well. As such, you MUST have a copy of the textbook.



Course Objectives

The student will be able to:

1. define and describe a variety of parallel computer architectures and discuss the performance tradeoffs among these
2. define, describe and discuss both numerical and non-numerical algorithms and to demonstrate an understanding of their parallel formulations
3. apply fundamental principles to synthesize parallel programs to solve problems using both explicit message-passing and shared-memory programming paradigms
4. define and describe challenges faced by the high-performance computing community and explain techniques and technologies appropriate to their solution
5. apply standard analysis techniques to predict the performance and scalability of a parallel program and/or a high-performance computing system

Course Objectives Mapped to Student Learning Objectives

| | | SLO | | | | | | | | | | | |
|----|---|-----|---|---|---|---|---|---|---|---|---|---|-------|
| | | a | b | c | d | e | f | g | h | i | j | k | l (j) |
| CO | 1 | X | | | | | X | | X | | | | |
| | 2 | X | X | | | | | | | | | | |
| | 3 | X | X | X | | | | | | | X | | |
| | 4 | X | X | | | | | | | | | | |
| | 5 | X | X | | | | | | | | X | | |
| | | | | | | | | | | | | | |
| | | | | | | | | | | | | | |

Student Learning Outcomes

- a) An ability to apply fundamental principles of computing and mathematics
- b) An ability to analyze a problem, and identify and define the requirements appropriate to its solution
- c) An ability to design, implement, and evaluate a solution to meet specific requirements subject to a set of constraints
- d) An ability to function effectively on multi-disciplinary teams to accomplish a common goal
- e) An understanding of professional and ethical responsibilities
- f) An ability to communicate effectively, both verbally and in writing
- g) An ability to analyze the local and global impact of computing on individuals, organizations, and society
- h) Recognition of the need for and an ability to engage in life-long learning
- i) An ability to use current techniques, skills, and tools necessary for computing practice Computer Science
- j) [CS ONLY] An ability to apply mathematical foundations, algorithmic principles, and computer science theory in the modeling and design of computer-based systems through the critical analysis of the tradeoffs involved in design choices
- k) [CS ONLY] An ability to apply design and development principles in the construction of complex software systems
- l) [IS ONLY] An understanding of processes that support the development, deployment, and management of information systems within a business-centric application environment

Study Guide 1

Motivation, History, and Architectures (Quinn chapters 1 & 2)

Parallel Computing (notes)

Goal: speeding up computation

Interesting problems communicate

Ways to communicate

- shared memory

- message passing

Message Passing (notes)

send/receive data

issues

- synchronization

 - blocking/non-blocking

 - synchronous

 - asynchronous

- buffering

 - none

 - limited

 - infinite

 - explicit

- naming

 - direct/indirect

 - symbolic

 - symmetric/asymmetric

- data size and type

collective communication

MPI (notes)

basic calls

- Init, Comm_rank, Comm_size, Send, Recv, Finalize

communicators, size ranks

basic data types

tags

semantics

- non-overtaking

- progress

- no fairness

- limited resources

Sendrecv, Sendrecv_replace

IO models

- master task

- independent

- hybrid

- parallel IO

collective operations

- Barrier, Bcast, Reduce, Allreduce, Scatter, Gather, Allgather, All-to-all,

- Reduce_scatter,

- Scatterv, Gatherv, Allgather, Alltoallv, Scan

modes

- normal

- buffered

- synchronous

- ready

non-blocking IO

- Isend, Issend, Ibsend, Irecv, Irecv

- Test, Wait, Testany, Waitany, Testall, Waitall, Testsome, Waitsome

MPI (cont)

derived (user-defined) datatypes

Type_contiguous, Type_vector, Type_hvector, Type_indexed, Type_hindexed,
Type_struct, Type_commit, Type_free, Type_size, Type_extent

groups

What are they and why do we have them?

Group_incl, Group_excl, Group_rank, Group_size

communicators

What are they and why do we need them?

Comm_create, Comm_group, Comm_split

topologies

What are they and how do they help programmers?

Cart_create, Cart_coords, Cart_rank, Cart_shift

Parallel Program Design (Quinn chapter 3)

Foster's methodology

partitioning

communication

agglomeration

mapping

Decomposition

data

for 2D structures:

by row, by column, checkerboard

interleaved

block

functional

Performance Analysis (Quinn chapter 7)

serial and parallel runtime

speedup

Amdahl's Law

efficiency

isoefficiency

Overhead function

Scalability function