

Fall 2006

Georgia Tech, CS 4803-DA

Georgia Tech, CS 8803-DA

Tu/Th 3:05pm - 4:25pm

Classroom: CCB Room 101

Computational Science & Engineering (CSE) Algorithms

<http://www.cc.gatech.edu/~bader/COURSES/GATECH/CS8803-Fall12006/>

Instructor: Prof. David A. Bader, CCB 257, 404-385-0004, bader@cc

Course Assistant: TBD

Office Hours: Tuesday/Thursday 2:00pm-3:00pm, or by appointment

Suggested Textbook:

1. [CLRS]: Cormen, Leiserson, Rivest, and Stein,
Introduction to Algorithms, Second edition, MIT Press, 2001.

Course Description: This course will introduce students to designing high-performance and scalable algorithms for computational science & engineering applications. The course focuses on algorithm design, complexity analysis, experimentation, and optimization, for important science and engineering applications. Students will develop knowledge and skills concerning:

- the design and analysis of real-world algorithms employed in computational science and engineering applications, and
- performance optimization of applications using the best practices of algorithm engineering.

Pre-requisites: design and analysis of algorithms (CS 3510).

Students (from the Sciences, Engineering, and Computing) interested in algorithmic applications in science and engineering are encouraged to take this course.

This course can be taken for satisfying the theory breadth requirement by computer science graduate students (M.S. and non-theory Ph.D. students). This course cannot be taken by ACO students to satisfy their core requirement and theory Ph.D. students in computer science to satisfy the theory breadth requirement.

Grading:

- (25 %) Midterm
- (25 %) Final
- (25 %) Project
- (20 %) Homework
- (5 %) Class participation

CLASS POLICIES

1. Please let me know as soon as possible if you will need to re-schedule an exam, or have any special needs during the semester.
2. Each student must read and abide by the Georgia Tech Academic Honor Code, see www.honor.gatech.edu.
3. Plagiarizing is defined by Websters as “to steal and pass off (the ideas or words of another) as one’s own: use (another’s production) without crediting the source.” If caught plagiarizing, you will be dealt with according to the GT Academic Honor Code.
4. When working on homework, you may work with other students in the class. However; you must turn in separate copies of the homework with the following written on it: your name and the names of everyone with whom you collaborated.
5. Homework is due by 5PM on the given due date. Late homeworks will not be accepted without a legitimate excuse and approval from the instructor.
6. No collaboration is permitted on exams. The midterm and final exams will be in-class, closed-book exams. You will be allowed to take a “cheat sheet” (double-sided 8.5 x 11 sheet of paper) into each exam.
7. Unauthorized use of any previous semester course materials, such as tests, quizzes, homework, projects, and any other coursework, is prohibited in this course. Using these materials will be considered a direct violation of academic policy and will be dealt with according to the GT Academic Honor Code.
8. For any questions involving these or any other Academic Honor Code issues, please consult me, my teaching assistants, or www.honor.gatech.edu.

Coverage of Topics

The course balances the use of theory and practice in algorithm design by presenting the student with case studies from application domains. After an introduction to computational science and engineering applications, algorithm theory, and realistic models of computation, the course consists of two modules. The first focuses on the design and analysis of real-world algorithms, with an application walk-through from algorithmic techniques to construction of the algorithm and finally application-level performance. The second module focuses on attaining high-performance implementations through algorithm engineering. These modules develop the theoretic framework and complements the study with examples from real-world problems using implementations on modern computing systems.

Computational Science & Engineering Introduction

1. Overview of Computational Science and Engineering Applications; characteristics and requirements
2. Computability and Complexity
3. Ideal and Realistic Models of Computation
4. Design of Parallel Algorithms
5. Performance Analysis
6. Multi-scale, multi-discipline applications

Real-World Algorithms: Techniques to Applications

1. Detecting Genomic Repeats: Divide and Conquer, cache-aware data structures, string algorithms
2. Sequence similarity searching: dynamic programming, string algorithms, local alignments, BLAST
3. Parallel Sorting by Regular Sampling: Randomized algorithms, load balancing, partitioning, merging, and sorting
4. Scientific visualization: greedy algorithms, 3d surface construction, marching cubes
5. Graph partitioning: NP-completeness, approximation algorithms, adaptive mesh refinement
6. Graph analysis: search algorithms, semantic networks, betweenness centrality

Algorithm Engineering

1. Modeling of the Memory Hierarchy
2. Cache-friendly, cache-aware, and cache-oblivious algorithms
3. Ideal and realistic models of parallel computation
4. Parallel and distributed algorithms and applications
5. Parallel disk models
6. Problem solving frameworks
7. Tools for Performance analysis
8. Experimental methods and validation

Fall 2006, Tentative Course Schedule

| Week | Date | Lec | Topic |
|------|--------|-----|--|
| 1 | 22 Aug | 1 | Intro. to Computational Science & Engineering |
| | 24 Aug | 2 | Modern Computer Architecture, Models of Computation |
| 2 | 29 Aug | 3 | Review of Computability and Complexity |
| | 31 Aug | 4 | Performance Analysis, High-Performance Computing |
| 3 | 5 Sep | 5 | Multi-scale, multi-discipline applications |
| | 7 Sep | 6 | Scientific visualization |
| 4 | 12 Sep | 7 | Detecting Genomic Repeats |
| | 14 Sep | 8 | Parallel Sorting by Regular Sampling |
| 5 | 19 Sep | 9 | Sequence similarity searching 1 |
| | 21 Sep | 10 | Sequence similarity searching 2 |
| 6 | 26 Sep | 11 | Graph partitioning 1 |
| | 28 Sep | 12 | Graph partitioning 2 |
| 7 | 3 Oct | 13 | Graph analysis 1 |
| | 5 Oct | 14 | Graph analysis 2 |
| 8 | 10 Oct | 15 | Algorithm engineering and experiment design |
| | 12 Oct | 16 | Modeling of the Memory Hierarchy |
| 9 | 17 Oct | - | Fall Break |
| | 19 Oct | 17 | Parallel disk models |
| 10 | 24 Oct | 18 | Cache-friendly and cache-aware algorithms |
| | 26 Oct | 19 | Practical graph algorithms |
| 11 | 31 Oct | 20 | Problem solving environments |
| | 2 Nov | 21 | Tools for performance analysis |
| 12 | 7 Nov | 22 | Realistic models of parallel computation |
| | 9 Nov | 23 | Parallel and distributed algorithms and applications |
| 13 | 14 Nov | 24 | Irregular parallel algorithms |
| | 16 Nov | 25 | Large-scale graph algorithms |
| 14 | 21 Nov | 26 | Multicore and Multithreaded algorithms |
| | 23 Nov | - | Thanksgiving Break |
| 15 | 28 Nov | 27 | Experimental methods and validation |
| | 30 Nov | 28 | Special Topics: Productivity Metrics |
| 16 | 5 Dec | 29 | Special Topics: Petascale algorithms |
| | 7 Dec | 30 | Special Topics: Open Problems in CSE Algorithms |
| 17 | 11 Dec | - | Final Exam (11:30am-2:20pm) |

Additional Readings:

1. Guy Blelloch, Algorithms in the Real World, Lecture Notes
2. Alok Aggarwal and Jeffrey Scott Vitter, The Input/Output Complexity of Sorting and Related Problems, Communications of the ACM, 31:1116-1127, 1988.
3. Sandeep Sen, Siddhartha Chatterjee, Neeraj Dumir, Towards a theory of cache-efficient algorithms, Journal of the ACM, 49(6):828-858, 2002.
4. A. LaMarca and R.E. Ladner, The Influence of Caches on the Performance of Heaps, Journal of Experimental Algorithmics, Vol 1, 1996.
5. A. LaMarca and R.E. Ladner, The Influence of Caches on the Performance of Sorting, Journal of Algorithms, 31:66-104, 1999.
6. R.E. Ladner, J.D. Fix, and A. LaMarca, Cache Performance Analysis of Traversals and Random Accesses, Tenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 1999.
7. John W. Romein, Jaap Heringa, Henri E. Bal, A Million-Fold Speed Improvement in Genomic Repeats Detection, Supercomputing 2003.
8. Joon-Sang Park, Michael Penner, Viktor K Prasanna, Optimizing Graph Algorithms for Improved Cache Performance, International Parallel and Distributed Processing Symposium, Fort Lauderdale, FL, April 2002.
9. Markus Kowarschik, Ulrich Rude, Christian Weiss, and Wolfgang Karl, Cache-Aware Multigrid Methods for Solving Poisson's Equation in Two Dimensions, Computing, 64:381-399, 2000.
10. M. Frigo, C. E. Leiserson, H. Prokop, and S. Ramachandran, Cache-Oblivious Algorithms, IEEE Symposium on Foundations of Computer Science, 1999.
11. Stephen Alstrup, Michael A. Bender, Erik D. Demaine, Martin Farach-Colton, J. Ian Munro, Theis Rauhe, Mikkel Thorup, Efficient Tree Layout in a Multilevel Memory Hierarchy, Extended version of ESA 2002 paper, November 2002.
12. Lars Arge, Michael A. Bender, Erik D. Demaine, Bryan Holland-Minkley, J. Ian Munro, Cache-Oblivious Priority Queue and Graph Algorithm Applications, 34th ACM Symposium on Theory of Computing (STOC), 2002.
13. L. G. Valiant, A Bridging Model for Parallel Computation, Communication of the ACM, 33(8):103-111, 1990.

14. D. E. Culler, R. M. Karp, D. A. Patterson, A. Sahay, K. E. Schauser, E. Santos, R. Subramonian, and T. von Eicken, LogP: Towards a Realistic Model of Parallel Computation, 4th ACM Symp. Principles and Practice of Parallel Programming (PPoPP), pp 1-12, May 1993.
15. V. Ramachandran, A General-Purpose Shared-Memory Model for Parallel Computation, in M. T. Heath and A. Ranade and R. S. Schreiber (eds.), Algorithms for Parallel Processing v 105, pp 1–18, Springer-Verlag, New York, 1999.
16. M. Snir and D.A. Bader, A Framework for Measuring Supercomputer Productivity, The International Journal of High Performance Computing Applications, 2004.
17. D.A. Bader, B.M.E. Moret, and P. Sanders, “High-Performance Algorithm Engineering for Parallel Computation,” Lecture Notes of Computer Science, 2002.
18. William E. Lorensen, Harvey E. Cline. Marching Cubes: A High Resolution 3D Surface Construction Algorithm. Computer Graphics, Volume 21, Number 4, July 1987.
19. M. Erez, J. H. Ahn, A. Garg, W.J. Dally, and E. Darve, Analysis and Performance Results of a Molecular Modeling Application on Merrimac, SC’04, Pittsburgh, PA, November 2004.
20. D.S. Johnson, A Theoretician’s Guide to the Experimental Analysis of Algorithms, in Proceedings of the 5th and 6th DIMACS Implementation Challenges, M. Goldwasser, D. S. Johnson, and C. C. McGeoch, Editors, American Mathematical Society, Providence, 2002.