Algorithmic Aspects of Computer Networks
CAS CS 559

TR 11:00-12:30, MCS B29

Instructor:	**Prof. John Byers**
Office	MCS 270, Phone: 353-8925
Email	byers@cs.bu.edu
Web	www.cs.bu.edu/fac/byers
Office Hours	Mon: 9:30 - 11:00, Tues: 3 - 4:30

Course Overview: Today's Internet researcher must carry a large toolkit. Expertise in network measurement, network modeling, protocol design and systems engineering are perforce. But while many researchers bring these skills to the table, far fewer have deep insight when it comes to questions of algorithms and algorithmic analysis. This is all the more surprising given the wealth of elegant algorithmic constructs which have successfully been applied to a broad spectrum of problems in computer networking, the Web, and in many other networks arising in science and engineering in recent years. This four-credit elective is designed to strengthen and broaden a student's theoretical background while developing an appreciation for how to integrate algorithmic research results into networked systems. For example, we will consider the case studies of Akamai and Google and their roots in applied algorithms research.

The course focuses on 1) fundamental algorithmic principles as relate to the Internet, 2) how algorithmic methods have been applied to specific networked applications, and 3) the limits of algorithmic practicability, e.g., in deciding if and when heuristics should be employed. Increasingly, research in computer networking closely relate to economic considerations: we will therefore also study the interplay between algorithmics, economics, and computer networks.

Prerequisites: The course is primarily geared toward graduate or advanced undergraduate students who are interested in conducting research in computer networking or in algorithms. Graduate students in other disciplines as well as those students who did excellent work in CS 455/655 are also encouraged to attend. The prerequisites for this course are: CS 455/655 or equivalent, CS 330, and fulfillment of the undergraduate-level CS Background requirements. Most of the papers we read will delve into algorithmic, statistical or information-theoretic techniques, so a solid background in related mathematical foundations is expected. Please see the instructor if you are uncertain about your level of preparation.

Course Expectations and Grading: For class, students will be expected to read approximately two research papers per week and either answer questions about or provide short written critiques of these papers prior to lecture. The course will be subdivided into a set of units, where the lecture format for each unit will start as one or two background lectures, followed by a set of lectures that involve more interactive discussion. For each unit of the course, a group of students chosen in advance will serve as **specialists**, i.e. will be experts on the papers we are discussing by virtue of additional reading or outside study directed by the instructor. These students will be expected to help facilitate the discussion, brainstorm about research directions, and help with the presentation of the material (or with supplemental material). These students will also generate scribe notes reflecting the technical material and the class discussion that will be distributed to the class.

Grading: The capstone of this class is a semester-long research project, to be conducted individually, or in groups of two with the instructor's permission. Groups of two will be expected to accomplish commensurately more to obtain the same grade as an individual working alone. The project culminates in a presentation to the class and a writeup in the style of a networking conference paper due during the last week of class. The project and presentations will constitute 50% of the overall grade. Information about style guidelines for both the paper and presentation will also be given during the course. I will expect students in this class to take the project very seriously and there will be regular interaction with the instructor outside of class to work on the projects — in years past, several of the projects in the class ultimately led to published papers.

Project Milestones:

- *Week 4*: Project format and suggested project topics distributed and discussed.
- *Week 7*: Three page project description and workplan due.
- *Week 8*: Meeting to discuss project description with instructor.
- *Week 11*: Meeting to discuss research progress on project with instructor.
- *Week 12*: Draft of completed sections and outline due (recommended, but optional).
- *Week 14*: Completed writeup due. Oral presentation.

Class participation will constitute 20% of a student's overall grade – this grade will be based both on the student's work as a specialist and contributions to the class discussion throughout the course. For the remaining 30% of the grade, I will periodically pose a few challenging problems as homework questions (10%), and there will be two in-class quizzes testing the main concepts in the papers and class discussions (10% each).

Reading List and Textbooks: The weekly readings will be maintained on the course webpage. A preliminary reading list and topics that the instructor intends to cover is attached. The course will primarily draw from recent research papers in the field, so there is no required text. However, there are recommended texts which cover networking fundamentals, elementary aspects of randomized algorithms, and basics of information theory respectively, and we will periodically draw from these sources in lecture.

The networking text I recommend is: Larry Peterson and Bruce Davie, *Computer Networks: A Systems Approach*, 5th Edition, Morgan Kaufmann, 2007. The randomized algorithms text I recommend is: Michael Mitzenmacher and Eli Upfal, *Probability and Computing*, Cambridge University Press, 2004. A classic reference on information theory is *Information Theory and Reliable Communication*, by Robert G. Gallager, Wiley Publishing, ISBN #0471290483.

If you happen to already own a different high-quality networking or information theory textbook, such as the Kurose-Ross networking text, then that is sufficient for this class.

Academic Conduct: Your work in this class falls under the purview of the College of Arts and Sciences Code of Academic Conduct. Any incidence of cheating or plagiarism in this class will be passed on to the CAS Academic Conduct Committee.

Syllabus: The weekly courseplan is listed below. Lectures will primarily focus on and draw from the technical articles (complete citations are listed in the Reference List), but will also briefly cover background material from the textbooks noted above.

Topic 1: Information Dispersal, Erasure Codes and Network Coding	
Week 1: Secret sharing	Shamir '79
Week 2: Information dispersal and coding basics	Rabin '89, Textbook
Week 3: Forward error correction and network coding	Luby '02, Ahlswede et al '00
Week 4: Coding applications	Byers et al '02, Chachulski et al '07
Topic 2: Concise Representations of Sets with Applications to Information Exchange	
Week 5: Randomization, hashing and entropy	
Week 6: Bloom filters	Bloom, Fan et al '00, Snoeren et al '01
Week 7: Min-wise permutations	Broder et al '98, Byers et al '04
Topic 3: Consistent Hashing and Distributed Hash Tables	
Week 8: Consistent hashing foundations	Karger et al '97, Plaxton et al '97
Week 9: DHTs and Applications	Stoica et al '01, Ratnasamy et al '01
Week 10: Probabilistic counting and sketching	Flajolet et al '85, Cormode et al '05
Week 11: Applications to monitoring	Estan et al '03, Considine et al '04
Topic 4: Algorithms for the Web	
Week 12: PageRank and Kleinberg's algorithm	Page et al '97, Kleinberg '99
Topic 5: The Internet, Economics, and E-Commerce	
Week 13: Internet economics	Fabrikant et al '03, Corbo et al '05
Week 13: Algorithmic game theory	Roughgarden et al, Feigenbaum et al