Algorithmic Aspects of Computer Networks
CAS CS 559

MW 9:30-11:00, PSY B51

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: 2 - 3:30, Wed: 11 - 12:30

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 and in Web-related contexts 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.

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*, 3rd Edition, Morgan Kaufmann, 2001. The randomized algorithms

text I recommend is: Michael Mitzenmacher and Eli Upfal, *Probability and Computing*, Cambridge University Press, 2004. An indispensable 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. The one other comparable randomized algorithms text that I recommend is: Rajeev Motwani and Prabhakar Raghavan, *Randomized Algorithms*, Cambridge University Press, 1998.

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	Stoica et al '01, Ratnasamy et al '01
Week 10: Further Applications	Dabek et al '01, Chawathe et al '05
Topic 4: Summarizing Data Streams with Applications to Network Monitoring	
Week 11: Probabilistic counting and sketching	Flajolet et al '85, Cormode et al '05
Week 12: Applications to monitoring	Estan et al '03, Considine et al '04
Topic 5: Indexing the Web	
Week 13: PageRank and Kleinberg's algorithm	Page et al '97, Kleinberg '99
Week 14: Online auctions and markets	TBA

Reading List:

Topic 1: Information Dispersal, Erasure Codes and Network Coding

[D1] A. Shamir, "How to Share a Secret," Communications of the ACM 22(11), pp. 612-613, 1979.

[D2] M. Rabin, "Efficient Dispersal of Information for Security, Load Balancing and Fault Tolerance," Journal of the ACM 38, pp. 335-348, 1989.

[D3] M. Luby, "LT Codes," in Proc. of the 43rd Annual IEEE Symposium on Foundations of Computer Science (FOCS '02).

[D4] J. Byers, M. Luby, and M. Mitzenmacher, "A Digital Fountain Approach to Asynchronous Reliable Multicast," IEEE Journal on Selected Areas in Communications, October 2002.

[D5] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network Information Flow," IEEE Transactions on Information Theory, IT-46, pp. 1204-1216, 2000.

[D6] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Médard, and J. Crowcroft, "XORs in the Air: Practical Wireless Network Coding," in ACM SIGCOMM 2006.

[D7] S. Chachulski, M. Jennings, S. Katti, and D. Katabi, "Trading Structure for Randomness in Wireless Opportunistic Routing," in ACM SIGCOMM 2007.

Topic 2: Concise Representations of Sets with Applications to Information Exchange

[S1] B. Bloom. "Space/time trade-offs in hash coding with allowable errors," Communications of the ACM, 13(7):422-426, 1970. (From the ACM Digital Library).

[S2] L. Fan, P. Cao, J. Almeida and A. Z. Broder, "Summary Cache: A Scalable Wide-area Cache Sharing Protocol," IEEE/ACM Transactions on Networking, 2000.

[S3] A. Snoeren, C. Partridge, L. Sanchez, C. Jones, F. Tchakountio, S. Kent, W. Strayer, "Hash-Based IP Traceback," in Proceedings of ACM SIGCOMM '01.

[S4] A. Z. Broder, M. Charikar, A. M. Frieze, M. Mitzenmacher, "Min-Wise Independent Permutations," Journal of Computer and System Sciences (JCSS), 1998.

[S5] J. Byers, J. Considine, M. Mitzenmacher and S. Rost, "Informed Content Delivery Across Adaptive Overlay Networks," IEEE/ACM Transactions on Networking, October 2004.

[S6] B. Parno, D. Wendlandt, E. Shi, A. Perrig, B. Maggs, and Y.-C. Hu, "Portcullis: Protecting Connection Setup from Denial-of-Capability Attacks," in Proceedings of ACM SIGCOMM '07.

Topic 3: Consistent Hashing and Distributed Hash Tables with Applications to Wide-Area Storage and Retrieval

[H1] D. Karger, E. Lehman, F. T. Leighton, M. Levine, D. Lewin, and R. Panigrahy, "Consistent hashing and random trees: Distributed caching protocols for relieving hot spots on the World Wide Web," in Proceedings of STOC '97. (The basis of Akamai's initial technology).

[H2] G. Plaxton, R. Rajaraman, A. W. Richa, "Accessing Nearby Copies of Replicated Objects in a Distributed Environment," in Proc. of ACM Symposium on Parallel Algorithms and Architectures (SPAA), 1997.

[H3] I. Stoica, R. Morris, D. Karger, F. Kaashoek, and H. Balakrishnan, "Chord: A Scalable Peer-To-Peer Lookup Service for Internet Applications," in Proceedings of ACM SIGCOMM '01.

[H4] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker, "A Scalable Content-Addressable Network," in Proceedings of ACM SIGCOMM '01.

[H5] F. Dabek, M. F. Kaashoek, D. Karger, R. Morris, I. Stoica, "Wide-area cooperative storage with CFS," in Proceedings of the 18th ACM Symposium on Operating Systems Principles (SOSP '01).

[H6] Y. Chawathe, S. Ramabhadran, S. Ratnasamy, A. LaMarca, J. Hellerstein and S. Shenker, "A Case Study in Building Layered DHT Applications," in Proceedings of ACM SIGCOMM 2005

Topic 4: Summarizing Data Streams with Applications to Network Monitoring

[M1] P. Flajolet and G. Martin, "Probabilistic Counting Algorithms for Database Applications," Journal of Computer and System Sciences, 1985.

[M2] G. Cormode and S. Muthukrishnan. "An improved data stream summary: the count-min sketch and its applications," J. Algorithms, 55(1): 58–75, April 2005,

[M3] C. Estan and G. Varghese, "New Directions in Traffic Measurement and Accounting: Focusing on the Elephants, Ignoring the Mice," ACM Transactions on Computer Systems, 21(3):270–313, 2003.

[M4] J. Considine, F. Li, G. Kollios and J. W. Byers. "Approximate Aggregation Techniques for Sensor Databases," in Proc. of the 20th IEEE Int'l Conference on Data Engineering (ICDE '04), Boston, MA, April 2004, pp. 449-460.

Topic 5: Indexing the Web

[W1], L. Page, S. Brin, R. Motwani and T. Winograd, The PageRank Citation Ranking: Bringing Order to the Web," Stanford University Technical report. (The basis of Google's original technology).

[W2], J. Kleinberg, "Authoritative Sources in a Hyperlinked Environment," Journal of the ACM 46, 1999.

[W3] K. Bharat and M. Henzinger, "Improved Algorithms for Topic Distillation in a Hyperlinked Environment," Proceedings of the 21st International ACM SIGIR Conference on Research and Development in Information Retrieval, 1998, pp. 104-111.