Discrete Algorithms Seminar

List of Presentations

Fall 2008

August 27, 2008

Title: "On the tandem duplication-random loss model of genome rearrangement."

Authors: Kamalika Chaudhuri, Kevin Chen, Radu Mihaescu, and Satish Rao.

Journal or Conference: SODA 2006: 564–570.

Abstract: We initiate the algorithmic study of a new model of genome rearrangement, the tandem duplication-random loss model, in which a genome evolves via successive rounds of tandem duplication of a contiguous segment of genes, followed by the loss of one copy of each of the duplicated genes. This model is well-known in the evolutionary biology literature, where it has been used to explain many of the known rearrangements in vertebrate mitochondrial genomes. Based on the model, we formalize a notion of distance between two genomes and show how to compute it efficiently for two interesting regions of the parameter space. We then consider median problems (i.e. finding the point which minimizes the sum of distances to a given set of points under some distance function) in the context of maximum parsimony phylogenetic reconstruction for these two special cases. Surprisingly, one of them turns out to correspond to the well-known rank aggregation problem, while the other corresponds to the biologically interesting case of whole genome duplication and loss, and we give an $O(loglogn)$ additive approximation algorithm for the latter.

Internet: `http://www.phylo.org/pdf_docs/9_(15)Mihaescu_soda2006.pdf`

Speaker: Andres Varon.

September 03, 2008

Title: "Stochastic Analysis for Online Combinatorial Optimization Problems."

Authors: Naveen Garg, Anupam Gupta, Stefano Leonardi, and Piotr Sankowski.

Journal or Conference: SODA 2008: 942–951.

Abstract: In this paper, we study online algorithms when the input is not chosen adversarially, but consists of draws from some given probability distribution. While this model has been studied for online problems like paging and k-server, it is not known how to beat the F(log n) bound for online Steiner tree if at each time instant, the demand vertex is a uniformly random vertex from the graph. For the online Steiner tree problem, we show that if each demand vertex is an independent draw from some probability distribution $p : V \to [0, 1]$, a variant of the natural greedy algorithm achieves $E_\omega[\mathcal{A}(\omega)]/E_\omega[OPT(\omega)] = O(1)$; moreover, this result can be extended to some other subadditive problems. Both assumptions that the input sequence consists of independent draws from p, and that p is known to the algorithm are both essential; we show (almost) logarithmic lower bounds if either assumption is violated. Moreover, we give preliminary results on extending the Steiner tree results above to the related "expected ratio" measure $E_\omega[\mathcal{A}(\omega)/OPT(\omega)]$. Finally, we use these ideas to give an average-case analysis of the Universal TSP problem.

Internet: A pdf version is Number 7 in the first link. The second link is the ACM portal.

`http://www.cs.cmu.edu/~anupamg/publicat.html`
`http://portal.acm.org/citation.cfm?id=1347185&jmp=cit&coll=GUIDE&dl=GUIDE,`

Speaker: Ou Liu.

September 10, 2008

Title: "Geometric Aspects of Online Packet Buffering: An optimal Randomized Algorithm for Two Buffers."

Authors: Marcin Bienkowski and Aleksander Madry.

Journal or Conference: Latin 2008: 252–263.

Abstract: We study packet buffering, a basic problem occurring in network switches. We construct an optimal 16/13-competitive randomized online algorithm PB for the case of two input buffers of arbitrary sizes. Our proof is based on geometrical transformations which allow to identify the set of sequences incurring extremal competitive ratios. Later we may analyze the performance of PB on these sequences only.

Internet:
`http://www.ii.uni.wroc.pl/~mbi/papers/2008-01-latin/packet-buffering.pdf`

Speaker: Michael Lampis.

September 17, 2008

Title: "Contraction Hierarchies: Faster and Simpler Hierarchical Routing in Road Networks."

Authors: Robert Geisberger, Peter Sanders, Dominik Schultes, and Daniel Delling.

Journal or Conference: WEA 2008: 319–333.

Abstract: We present a route planning technique solely based on the concept of node contraction. The nodes are first ordered by importance. A hierarchy is then generated by iteratively contracting the least important node. Contracting a node v means replacing shortest paths going through v by shortcuts. We obtain a hierarchical query algorithm using bidirectional shortest-path search. The forward search uses only edges leading to more important nodes and the backward search uses only edges coming from more important nodes. For fastest routes in road networks, the graph remains very sparse throughout the contraction process using rather simple heuristics for ordering the nodes. We have five times lower query times than the best previous hierarchical Dijkstra-based speedup techniques and a negative space overhead, i.e., the data structure for distance computation needs less space than the input graph. CHs can be combined with many other route planning techniques, leading to improved performance for many-to-many routing, transit-node routing, goal-directed routing or mobile and dynamic scenarios.

Internet: `http://algo2.iti.uni-karlsruhe.de/documents/gssd-chfsh-08.pdf`

Speaker: Marcin Szczodrak.

September 24, 2008

Title: "Broadcast Scheduling: Algorithms and Complexity."

Authors: Jessica Chang, Thomas Erlebach, Renars Gailis, and Samir Khuller.

Journal or Conference: SODA 2008: 473–482.

Abstract: Broadcast Scheduling is a popular method for disseminating information in response to client requests. There are n pages of information, and clients request pages at different times. However, multiple clients can have their requests satisfied by a single broadcast of the requested page. In this paper we consider several related broadcast scheduling problems. One central problem we study simply asks to minimize the maximum response time (over all requests). Another related problem we consider is the version in which every request has a release time and a deadline, and the goal is to maximize the number of requests that meet their deadlines. While approximation algorithms for both these problems were proposed several years back, it was not known if they were NP-complete. One of our main results is that both these problems are NP-complete. In addition, we use the same unified approach to give a simple NP-completeness proof for minimizing the sum of response times. A very complicated proof was known for this version. Furthermore, we give a proof that FIFO is a 2-competitive online algorithm for minimizing the maximum response time (this result had been claimed earlier with no proof) and that there is no better deterministic online algorithm (this result was claimed earlier as well, but with an incorrect proof).

Internet: `www.cs.umd.edu/~samir/grant/broadcast.ps`

Speaker: Valia Mitsou.

October 15, 2008

Title: "Combinatorial explorations in Su-Doku."

Author: Jean-Marie Chauvet.

Abstract: Su-Doku, a popular combinatorial puzzle, provides an excellent test bench for heuristic explorations. Several interesting questions arise from its deceptively simple set of rules. How many distinct Su-Doku grids are there? How to find a solution to a Su-Doku puzzle? Is there a unique solution to a given Su-Doku puzzle? What is a good estimation of a puzzle's difficulty? What is the minimum puzzle size (the number of "givens")? This paper explores how these questions are related to the well-known alldifferent constraint which emerges in a wide variety of Constraint Satisfaction Problems (CSP) and compares various algorithmic approaches based on different formulations of Su-Doku.

Internet: `http://arxiv.org/PS_cache/arxiv/pdf/0803/0803.4253v1.pdf`

Speaker: Matthew Meyer.

October 22, 2008

Title: "Experimental Evaluation of an Exact Algorithm for the Orthogonal Art Gallery Problem."

Authors: Marcelo C. Couto, Cid C. de Souza, and Pedro J. de Rezende.

Journal or Conference: WEA 2008: 101–113.

Abstract: We consider the Orthogonal Art Gallery problem (oagp) whose goal is to minimize the number of vertex guards required to watch an art gallery whose boundary is an n-vertex orthogonal polygon P. Here, we explore an exact algorithm for oagp, which we proposed in [1], that iteratively computes optimal solutions to Set Cover problems (scps) corresponding to discretizations of P. While it is known [1] that this procedure converges to an exact solution of the original continuous problem, the number of iterations executed is highly dependent on the way we discretize P. Although the best theoretical bound for convergence is $\Theta(n^3)$ iterations, we show that, in practice, it is achieved after only a few of them, even for random polygons of hundreds of vertices. As each iteration involves the solution of an scp, the strategy for discretizing P is of paramount importance. In this paper, we carry out an extensive empirical investigation with five alternative discretization strategies to implement the algorithm. A broad range of polygon classes is tested. As a result, we are able to significantly improve the performance of the algorithm, while maintaining low execution times, to the point that we achieve a fivefold increase in polygon size, compared to the literature.

Internet: `http://www.springerlink.com/content/n80jv60647384711/fulltext.pdf`

Speaker: Andi Toce.

October 29, 2008

Title: "Real-Time Indexing over Fixed Finite Alphabets."

Authors: Amihood Amir and Igor Nor.

Journal or Conference: SODA 2008: 1086–1095.

Abstract: The quest for a real-time indexing algorithm is over three decades old. To date there is no convincing understandable solution to this problem. This paper provides a real-time indexing algorithm over a constant sized alphabet. Assuming the text is arriving at a constant rate, the algorithm spends $O(1)$ time on every text symbol. Whenever a length m pattern is given, the algorithm decides in time $O(m)$ whether there is an occurrence of the pattern in the text thus far.

Internet: `http://portal.acm.org/citation.cfm?doid=1347082.1347201`

Speaker: Shoshana Neuburger.

November 5, 2008

Title: "Algorithms for Secretary Problems on Graphs and Hypergraphs."

Authors: Nitish Korula and Martin Pal.

Journal or Conference: CoRR abs/0807.1139: (2008).

Abstract: We examine several online matching problems, with applications to Internet advertising reservation systems. Consider an edge-weighted bipartite graph G, with partite sets L, R. We develop an 8-competitive algorithm for the following secretary problem: Initially given R, and the size of L, the algorithm receives the vertices of L sequentially, in a random order. When a vertex l in L is seen, all edges incident to l are revealed, together with their weights. The algorithm must immediately either match l to an available vertex of R, or decide that l will remain unmatched. Dimitrov and Plaxton show a 16-competitive algorithm for the transversal matroid secretary problem, which is the special case with weights on vertices, not edges. (Equivalently, one may assume that for each l in L, the weights on all edges incident to l are identical.) We use a similar algorithm, but simplify and improve the analysis to obtain a better competitive ratio for the more general problem. Perhaps of more interest is the fact that our analysis is easily extended to obtain competitive algorithms for similar problems, such as to find disjoint sets of edges in hypergraphs where edges arrive online. We also introduce secretary problems with adversarially chosen groups. Finally, we give a 2e-competitive algorithm for the secretary problem on graphic matroids, where, with edges appearing online, the goal is to find a maximum-weight acyclic subgraph of a given graph.

Internet: `http://arxiv.org/PS_cache/arxiv/pdf/0807/0807.1139v1.pdf`

Speaker: Jeremy Seideman.

November 12, 2008

Title: "An approximation algorithm for conflict-aware broadcast scheduling in wireless ad hoc networks."

Authors: Reza Mahjourian, Feng Chen, Ravi Tiwari, My Thai, Hongqiang Zhai, and Yuguang Fang.

Journal or Conference: Mobicom 2008: 331–340.

Abstract: Broadcast scheduling is a fundamental problem in wireless ad hoc networks. The objective of a broadcast schedule is to deliver a message from a given source to all other nodes in a minimum amount of time. At the same time, in order for the broadcast to proceed as predicted in the schedule, it must not contain parallel transmissions which can be conflicting based on the collision and interference parameters in the wireless network. Most existing work on this problem use a limited network model which accounts only for conflicts occurring inside the transmission ranges of the nodes. The broadcast schedules produced by these algorithms are likely to experience unpredictable delays when deployed in the network. This is because they do not take into consideration other important sources of conflict in parallel transmissions, namely the interference range and the carrier sensing range. In this paper we develop a conflict-aware network model, which uses these parameters to increase the probability of scheduling conflict-free transmissions, and thereby improve the reliability of the broadcast schedule. We present and prove correctness of a constant approximation algorithm for minimum-latency broadcast scheduling under this network model. We also present a greedy heuristic algorithm for the same problem. Experimental results are provided to evaluate the performance of our algorithms. In addition, the algorithms are analyzed to justify their performance trends.

Internet: `http://portal.acm.org/citation.cfm?id=1374663`

Speaker: Rui Zhang.

November 19, 2008

Title: "Incremental String Comparison."

Authors: Gad M. Landau, Eugene W. Myers, and Jeanette P. Schmidt.

Journal or Conference: SIAM J. Comput., 27, 2, 557–582 (1998).

Abstract: The problem of comparing two sequences A and B to determine their longest common subsequence (LCS) or the edit distance between them has been much studied. In this paper we consider the following incremental version of these problems: given an appropriate encoding of a comparison between A and B, can one incrementally compute the answer for A and bB, and the answer for A and Bb with equal efficiency, where b is an additional symbol? Our main result is a theorem exposing a surprising relationship between the dynamic programming solutions for two such "adjacent" problems. Given a threshold k on the number of differences to be permitted in an alignment, the theorem leads directly to an $O(k)$ algorithm for incrementally computing a new solution from an old one, as contrasts the $O(k^2)$ time required to compute a solution from scratch. We further show, with a series of applications, that this algorithm is indeed more powerful than its nonincremental counterpart. We show this by solving the applications with greater asymptotic efficiency than heretofore possible. For example, we obtain $O(nk)$ algorithms for the longest prefix approximate match problem, the approximate overlap problem, and cyclic string comparison.

Internet: `http://cs.haifa.ac.il/LANDAU/gadi/LMS.pdf`

Speaker: Justin Tojeira.

November 26, 2008

Title: "Online Scheduling of Equal-Length Jobs on Parallel Machines."

Authors: Jihuan Ding, Tom Ebenlendr, Jir Sgall, and Guochuan Zhang

Journal or Conference: ESA 2007: 427–438.

Abstract: We study on-line scheduling of equal-length jobs on parallel machines. Our main result is an algorithm with competitive ratio decreasing to $e/(e-1) \approx 1.58$ as the number of machine increases. For $m = 3$, this is the first algorithm better than 2-competitive greedy algorithm. Our algorithm has an additional property called immediate decision: at each time, it is immediately decided for each newly released job if it will be scheduled, and if so, then also the time interval and machine where it is scheduled is fixed and cannot be changed later. We show that for two machines, no deterministic algorithm with immediate decision is better than 1.8-competitive; this lower bound shows that our algorithm is optimal for $m = 2$ in this restricted model. We give some additional lower bounds for algorithms with immediate decision.

Internet: `www.math.cas.cz/~sgall/ps/eqpar.ps`.

Speaker: Antonios Achilleos.

December 3, 2008

Title: "Studying (Non-Planar) Road Networks Through an Algorithmic Lens."

Authors: David Eppstein and Michael T. Goodrich.

Journal or Conference: SIGSPATIAL 2008.

Abstract: This paper studies real-world road networks from an algorithmic perspective, focusing on empirical studies that yield useful properties of road networks that can be exploited in the design of fast algorithms that deal with geographic data. Unlike previous approaches, our study is not based on the assumption that road networks are planar graphs. Indeed, based on the a number of experiments we have performed on the road networks of the 50 United States and District of Columbia, we provide strong empirical evidence that road networks are quite non-planar. Our approach therefore instead is directed at finding algorithmically-motivated properties of road networks as non-planar geometric graphs, focusing on alternative properties of road networks that can still lead to efficient algorithms for such problems as shortest paths and Voronoi diagrams. In particular, we study road networks as multiscale-dispersed graphs, which is a concept we formalize in terms of disk neighborhood systems. This approach allows us to develop fast algorithms for road networks without making any additional assumptions about the distribution of edge weights. In fact, our algorithms can allow for non-metric weights.

Internet: `http://arxiv.org/PS_cache/arxiv/pdf/0808/0808.3694v1.pdf`

Speaker: Seluk Cevher.

December 10, 2008

Title: "Fast Local serach for the Maximum Independent Set Problem."

Authors: Diogo V. Andrade, Mauricio G. C. Resende, and Renato F. Werneck.

Journal or Conference: WEA 2008: 220–234.

Abstract: Given a graph G=(V,E), the independent set problem is that of finding a maximum-cardinality subset S of V such that no two vertices in S are adjacent. We present a fast local search routine for this problem. Our algorithm can determine in linear time whether a maximal solution can be improved by replacing a single vertex with two others. We also show that an incremental version of this method can be useful within more elaborate heuristics. We test our algorithms on instances from the literature as well as on new ones proposed in this paper.

Internet: `http://www.optimization-online.org/DB_FILE/2008/02/1898.pdf`

Speaker: Mohammad Torkaharani.

December 17, 2008

Title I: "Engineering Algorithms for Approximate Weighted Matching."

Authors I: Jens Maue and Peter Sanders.

Journal or Conference I: WEA 2007: 242–255.

Abstract: We present a systematic study of approximation algorithms for the maximum weight matching problem. This includes a new algorithm which provides the simple greedy method with a recent path heuristic. Surprisingly, this quite simple algorithm performs very well, both in terms of running time and solution quality, and, though some other methods have a better theoretical performance, it ranks among the best algorithms.

Internet I: `http://www.ti.inf.ethz.ch/pw/publications/papers07/MaueSanders2007.pdf`

Title: "Computing Minimum-Weight Perfect Matchings."

Authors: William Cook and Andre Rohe.

Journal or Conference: INFORMS Journal on Computing: 138–148.

Abstract: We make several observations on the implementation of Edmonds blossom algorithm for solving minimum-weight perfect matching problems and we present computational results for geometric problem instances ranging in size from 1,000 nodes up to 5,000,000 nodes. A key feature in our implementation is the use of multiple search trees with an individual dual-change e for each tree. As a benchmark of the algorithms performance, solving a 100,000-node geometric instance on a 200 Mhz Pentium-Pro computer takes approximately 3 minutes.

Internet II: `http://www2.isye.gatech.edu/~wcook/papers/match_ijoc.`

Speakers: Mehmet Bicer ans Bilgin Zeki.