Algorithm Engineering Group

Site report 2005

The research activity of the group of Algorithm Engineering (AE) is concerned with the design, the engineering, the theoretical and experimental performance analysis of combinatorial algorithms for problems arising in modern Computer Systems and Networks, and in applications related to complex resource management problems. Our main research interests deal with the solution of optimization problems and the design of efficient data structures, with special emphasis on those applications involving large data sets. In particular we concentrate on:

1. algorithms that perform efficiently in a dynamically changing environment;
2. models and methodologies for the analysis and design of algorithms for information retrieval;
3. the efficient management of communication and information delivery and recovery in Wireless Networks and on the Internet;
4. the design and analysis of approximation algorithms for NP-hard optimization problems;
5. the design of on-line algorithms that work with incomplete information on the input instance;
6. the design and implementation of tools and platforms for the experimental analysis and visualization of the behavior of algorithms and data structures.

The achievements of the AE group are widely recognized. Giorgio Ausiello is Editor in Chief of Theoretical Computer Science, Series A, Algorithms and Complexity and Member of the Editorial Board of the International Journal of Foundations of Computer Science (IJFCS). He has also been awarded the title of 'Doctor Honoris Causa of Dauphine University'.

Members of the AE group are continuously involved in the Program and Steering Committees of prestigious International Conferences.

Giorgio Ausiello has been in the Selection Committee of the Goedel Prize during years 2003 - 2005. He has also served on the Program Committees of the 32nd International Colloquium on Automata, Languages and Programming (ICALP'05), the 15th International Symposium on Fundamentals of Computation Theory (FCT'05), and the 16th Annual International Symposium on Algorithms and Computation (ISAAC'05).

Alberto Marchetti-Spaccamela is member of the Steering Committee of the European Symposium on Algorithms and has served on the 2005 Selection Committee for the National Olympiads in Informatics.

Stefano Leonardi has been Program Chair of the 13th Annual European Symposium on Algorithms (ESA'05) and has served on the Program Committee of the Workshop on Models and Algorithms for Planning and Scheduling Problems (MAPSP05). He is also guest editor of the Special issue for WAW'04 on the Journal of Internet Mathematics and member of the Editorial Board of the Journal of Interconnection Networks.

Camil Demetrescu has been Program Co-chair (with Roberto Tamassia) of the 7th Workshop on Algorithm Engineering and Experiments (ALENEX'05) and has served on the Program Committee of the 31st International Workshop on Graph-Theoretic Concepts in Computer Science (WG'05). He is member of the Steering Committee of the ACM-SIAM Workshop on Algorithm Engineering and Experiments (ALENEX). He is also guest editor of the Proceedings of ALENEX/ANALCO'05.

A regular Seminar Program, the Interdepartmental Seminar on Algorithms (SIA), is also organized in cooperation with the Department of Computer Science of this University (see `http://www.dis.uniroma1.it/sia/`).

The AE group is currently cooperating with several prestigious research institutions: Max Planck für Informatik (Saarbrücken, Germany), CTI-Patras (Greece), ETH (Zurich, Switzerland), Université de Paris (Dauphine, France), Tel-Aviv University (Israel), AT&T - Research Labs (Florham Park, NJ, USA), ICSI-Berkeley (USA), Brown University (Providence, RI, USA), Carnegie-Mellon University (Pittsburgh, PA, USA), Microsoft Research (Mountain View, CA, USA).

The AE group is presently involved in the following research projects: ALENIA Spazio "Wideband Wireless Local Area Network" (WWLAN); EU-IST "Coevolution and self-organization in dynamical networks (COSIN)"; EU Contract 001907 "Dynamically Evolving Large Scale Information Sytems" (DELIS); MIUR National Project "Algorithms for the Next Generation Internet and Web: Methodologies, Design and Application" (ALGO-NEXT); FIRB National Projects - WEBMINDS and VICOM. A national committee of MIUR has also approved a financial support of 1M-Euro in three years for the institution at the University of Rome "La Sapienza" of an *Excellence Centre* that aims at creating new scientific and technical synergies in the area of transportation and logistics. A major role in this project is played by the Algorithm Engineering group at DIS.

Faculty members. Giorgio Ausiello, Luca Becchetti, Fabrizio d'Amore, Camil Demetrescu, Stefano Leonardi, Alberto Marchetti-Spaccamela, Umberto Nanni.

Post-docs Luigi Laura, Andrea Vitaletti, Luciana Salete Buriol, Guido Schäfer.

PhD students. Luca Allulli, Vincenzo Bonifaci, Debora Donato, Fabiano Sarracco, Andrea Ribichini.

Group visitors Carlos Castillo, Bruno Escoffier, Gabriel Moruz.

Graphs and Networks. In the area of graph and network algorithms major emphasis has been devoted to the analysis and design of algorithms for dynamically evolving graphs and networks.

In [17], we present fully dynamic algorithms for maintaining 3- and 5-spanners of undirected graphs under a sequence of update operations. For unweighted graphs we maintain a 3- or 5-spanner under insertions and deletions of edges; each operation is performed in $O(n)$ amortized time over a sequence of $\Omega(n)$ updates. The maintained 3-spanner (resp., 5-spanner) has $O(n^{3/2})$ edges (resp., $O(n^{4/3})$ edges), which is known to be optimal. On weighted graphs with d different edge cost values, we maintain a 3- or 5-spanner in $O(n)$ amortized time over a sequence of $\Omega(d \cdot n)$ updates. The maintained 3-spanner (resp., 5-spanner) has $O(d \cdot n^{3/2})$ edges (resp., $O(d \cdot n^{4/3})$ edges). The same approach can be extended to graphs with real-valued edge costs in the range $[1, C]$. All our algorithms are deterministic and are substantially faster than recomputing a spanner from scratch after each update.

In [7] we address the problem of dynamically maintaining minimum weight hyperpaths in a directed hypergraph in a decremental setting. For such a problem, we provide a new efficient algorithm that works for a wide class of hyperpath weight measures. This algorithm explicitly updates minimum weight hyperpaths in $O(L \cdot C + \max\{n, C\} \cdot \mathcal{H})$ worst case time under a sequence of L hyperarc weight increments and hyperarc deletions, where C is the maximum weight of minimum hyperpaths in $\mathcal{H}$ and $\mathcal{H}$ is the size of the representation of the hypergraph. Hyperpath weight measures are only required to belong to the class of *strict weakly superior functions*.

In [1], we present an algorithm for directed acyclic graphs that breaks through the $O(n^2)$ barrier on the single-operation complexity of fully dynamic transitive closure, where n is the number of edges in the graph. We can answer queries in $O(n^\epsilon)$ worst-case time and perform updates in $O(n^{\omega(1,\epsilon,1)-\epsilon} + n^{1+\epsilon})$ worst-case time, for any $\epsilon \in [0, 1]$, where $\omega(1, \epsilon, 1)$ is the exponent of the multiplication of an $n \times n^\epsilon$ matrix by an $n^\epsilon \times n$ matrix. The current best bounds on $\omega(1, \epsilon, 1)$ imply an $O(n^{0.575})$ query time and an $O(n^{1.575})$ update time in the worst case. Our subquadratic algorithm is randomized, and has one-sided error. As an application of this result, we show how to solve single-source reachability in $O(n^{1.575})$ time per update and constant time per query.

In [5] we present the first fully dynamic algorithm for maintaining all pairs shortest paths in directed graphs with real-valued edge weights. Given a dynamic directed graph G such that each edge can assume at most S different real values, we show how to support updates in $O(n^{2.5}\sqrt{S \log^3 n})$ amortized time and queries in optimal worst-case time. This algorithm is deterministic: no previous fully dynamic algorithm was known before for this problem. In the special case where edge weights can only be increased, we give a randomized algorithm with one-sided error that supports updates faster in $O(S \cdot n \log^3 n)$ amortized time. We also show how to obtain query/update trade-offs for this problem, by introducing two new families of randomized algorithms. Algorithms in the first family achieve an update bound of $\widetilde{O}(S \cdot k \cdot n^2)$ and a query bound of $\widetilde{O}(n/k)$, and improve over the previous best known update bounds for k in the range $(n/S)^{1/3} \leq k < (n/S)^{1/2}$. Algorithms in the second family achieve an update bound of $\widetilde{O}\left(S \cdot k \cdot n^2\right)$ and a query bound of $\widetilde{O}(n^2/k^2)$, and are competitive with the previous best known update bounds (first family included) for k in the range $(n/S)^{1/6} \leq k < (n/S)^{1/3}$.

In [26, 27] we survey the main techniques developed in the literature for the maintenance of dynamic trees and dynamic graphs.

Algorithms for the Web. Recently major interest in algorithmics is being devoted to the study of algorithms for searching and retrieving information in the Web. What makes such an area challenging is the huge size of the Web graph and the enormous amount of documents stored in its nodes, two aspects that require the design of extremely efficient algorithms. In this domain we have addressed the following topics: information retrieval in the Web and the analysis of the structure of the Web and of its graph theoretic properties.

It is widely known that spectral techniques are very effective for document retrieval. Recently, a lot of effort has been spent by researchers to provide a formal mathematical explanation for this effectiveness. Latent Semantic Indexing, in particular, is a text retrieval algorithm based on the spectral analysis of the occurrences of terms in text documents. Despite of its value in improving the quality of a text search, LSI has the drawback of an elevate response time, which makes it unsuitable for on-line search in large collections of documents (e.g., web search engines). In [19] we present two approaches aimed to combine the effectiveness of latent semantic analysis with the efficiency of text matching retrieval, through the technique of query expansion. We show that both approaches have relatively small computational cost and we provide experimental evidence of their ability to improve document retrieval.

Recently, there has been a surge of research activity in the area of Link Analysis Ranking, where hyperlink structures are used to determine the rel- ative authority of Web pages. One of the seminal works in this area is that of Kleinberg, who proposed the HITS algorithm. In [20], we undertake a theoretical analysis of the properties of the HITS algorithm on a broad class of random graphs. Working within the framework of Borodin et al., we prove that on this class (a) the HITS algorithm is stable with high probability, and (b) the HITS algorithm is similar to the INDEGREE heuristic that assigns to each node weight proportional to the number of incoming links. We demonstrate that our results go through for the case that the expected in-degrees of the graph follow a power-law distribution, a situation observed in the actual Web graph. We also study experimentally the similarity between HITS and INDEGREE, and we investigate the general conditions under which the two algorithms are similar.

Despite of being the sum of the decentralized and uncoordinated efforts by heterogeneous groups and individuals, the World Wide Web exhibits a well defined structure, characterized by several interesting properties. This structure was clearly revealed by Broder et al. who presented the evocative bow-tie structure of the Web, with a CORE comprised of a large strongly connected component, and four sets of vertices distinguishable from their relation with the CORE. The bow-tie structure is a relatively clear abstraction of the macroscopic picture of the Web, but it is very uninformative with respect to the finer details of the Web graph. In [21] we mine the inner structure of the Web graph. We present a series of measurements on the Web, which offer a better understanding of the individual components of the bow-tie. We also document the algorithmic techniques for performing these measurements. We discover that the scale-free properties permeate all the components of the bow-tie which exhibit the same macroscopic properties as the Web graph itself. However, close inspection reveals that their inner structure is quite distinct. We show that the Web graph does not exhibit self similarity within its components, and we propose a possible alternative picture for the Web graph, as it emerges through our experiments. The focus of [22] is on the practical use of data stream algorithms for monitoring statistical and topological properties of large graphs such as the webgraph. By webgraph we mean the directed graph generated from the link structure of webpages: each webpage is a node and each hyperlink is an arc in this graph. We study experimentally the application of the algorithm of Datar and Muthukrishnan for maintaining the indegree rarity distribution and the density of low cardinality bipartite cliques in a graph read in a streamed fashion. The alpha-rarity of a stream is the ratio between the number of elements that appear exactly alpha times and the total number of different items in the stream. We present results and show that the approximated values are very close to the optima even when a low precision is requested. We are interested in computing properties of large graphs, as the webgraph, using data stream algorithms. In [23] we report results on computing the indegree rarity distribution of a graph obtained as a stream of edges. We implement a rarity algorithm proposed in the literature and show experimentally that the results approximate very well the optimal value with very limited use of memory and time. Moreover, considering some structure in the stream, we present results for the algorithm adapted for maintaining the rarity distribution of the number of cliques of size three.

Wikipedia is an online encyclopedia (www.wikipedia.org), available in more than 100 languages. If we consider each article as a node and each hyperlink between articles as a link, we have a wikigraph, the link structure of wikipedia. We can extract one wikigraph for each available language, with size ranging from less then 1000 nodes to more than 500 thousand nodes and more then 5 million links. Associated with each node there are timestamps, indicating the creation and update dates of each page, that allows to study how the graph properties evolve over time. In a first part of a study we presented in [31] we observe that wikigraphs maintain the main characteristics of webgraphs, for which temporal information is usually not available. We then study the temporal evolution of several topological properties of wikigraphs and relate this measures to the number of updates of the documents.

Algorithms for Optimization. In the area of NP-hard optimization problems we have concentrated mainly in two directions: the study of efficient algorithms for the on line version of classical optimization problems and the study of approximability properties and approximation algorithms for network design and network management problems. Besides we have addressed some foundational aspects concerning the completeness of optimization problems in approximation classes, under the differential approximation measure. In [12] we consider two on-line versions of the asymmetric traveling salesman problem with triangle inequality. For the homing version, in which the salesman is required to return in the city where it started from, we give a (3+sqrt(5))/2-competitive algorithm and prove that this is best possible. For the nomadic version, the on-line analogue of the shortest asymmetric hamiltonian path problem, we show that the competitive ratio of any on-line algorithm has to depend on the amount of asymmetry of the space in which the salesman moves. We also give bounds on the competitive ratio of on-line algorithms that are zealous, that is, in which the salesman cannot stay idle when some city can be served.

Vehicle Routing Problems are generalizations of the well known Traveling Salesman Problem. In [18] we focus on the on-line version of these problems, where requests are not known in advance and arrive over time. We introduce a model of lookeahead for this class of problems, the time lookahead, which allows an on-line algorithm to foresee all the requests that will be released during next time units. We present lower and upper bounds on the competitive ratio of known and studied variants of the OLTSP; we compare these results with the ones from the literature. Our results show that the effectiveness of lookahead varies significantly as we consider different problems.

Next generation 3G/4G wireless data networks allow multiple codes (or channels) to be allocated to a single user, where each code can support multiple data rates. Providing fine-grained QoS to users in such networks poses the two dimensional challenge of assigning *both* power (rate) and codes to every user. This gives rise to a new class of parallel scheduling problems. In [2] we abstract general downlink scheduling problems suitable for proposed next generation wireless data systems. Our contribution includes a communication-theoretic model for multirate wireless channels. In addition, while conventional focus has been on throughput maximization, we attempt to optimize the maximum response time of jobs, which is more suitable for streams of user requests. We present provable results on the algorithmic complexity of these scheduling problems. In particular, we are able to provide very simple, online algorithms for approximating the optimal maximum response time. This relies on resource augmented competitive analysis. We also perform an experimental study with realistic data of channel conditions and user requests to show that our algorithms are more accurate than our worst case analysis shows, and that they provide fine-grained QoS to users effectively.

In [3] we introduce the notion of smoothed competitive analysis of online algorithms. Smoothed analysis has been proposed by Spielman and Teng [STOC 2001] to explain the behaviour of algorithms that work well in practice while performing very poorly from a worst case analysis point of view. We apply this notion to analyze the Multi-Level Feedback (MLF) algorithm to minimize the total flow time on a sequence of jobs released over time when the processing time of a job is only known at time of completion. The initial processing times are integers in the range $[1, 2^K]$. We use a partial bit randomization model, where the initial processing times are smoothened by changing the k least significant bits under a quite general class of probability distributions. We show that MLF admits a smoothed competitive ratio of $O((2^k/\sigma)^3 + (2^k/\sigma)^2 2^{K-k})$, where σ denotes the standard deviation of the distribution. In particular, we obtain a competitive ratio of $O(2^{K-k})$ if $\sigma = \Theta(2^k)$. We also prove an $\Omega(2^{K-k})$ lower bound for any deterministic algorithm that is run on processing times smoothened according to the partial bit randomization model. For various other smoothening models, including the additive symmetric smoothening model used by Spielman and Teng, we give a higher lower bound of $\Omega(2^K)$. A direct consequence of our result is also the first average case analysis of MLF. We show a constant expected ratio of the total flow time of MLF to the optimum under several distributions including the uniform distribution.

In the *multicommodity rent-or-buy* (MROB) network design problem we are given a network together with a set of k terminal pairs $(s_1, t_1), \ldots, (s_k, t_k)$. The goal is to provision the network so that a given amount of flow can be shipped between s_i and t_i for all $1 \leq i \leq k$ simultaneously. In order to provision the network one can either *rent* capacity on edges at some cost per unit of flow, or *buy* them at some larger fixed cost. Bought edges have no incremental, flow-dependent cost. The overall objective is to minimize the total provisioning cost. Recently, Gupta et al. presented a 12-approximation for the MROB problem. Their algorithm chooses a subset of the terminal pairs in the graph at random and then buys the edges of an approximate Steiner forest for these pairs. This technique has previously been successfully used by the same authors for the single sink rent-or-buy network design problem. In [9] we give a 5.5-approximation for the MROB problem by refining the algorithm of Gupta et al. and greatly simplifying their analysis. The improvement in our paper is based on a more careful adaptation and simplified analysis of the primal-dual algorithm for the Steiner forest problem due to Agrawal, Klein and Ravi. Our result significantly reduces the gap between the single-sink and multi-sink cases.

In [10] we design an approximately budget-balanced and group-strategyproof cost-sharing mechanism for the Steiner forest game. An instance of this game consists of an undirected graph $G = (V, E)$, non-negative costs c_e for all edges $e \in E$, and a set $R \subseteq V \times V$ of k terminal pairs. Each terminal pair $(s, t) \in R$ is associated with an agent that wishes to establish a connection between nodes s and t in the underlying network. A feasible solution is a forest F that contains an s, t-path for each connection request $(s, t) \in R$. Previously, Jain and Vazirani gave a 2-approximate budget-balanced and group-strategyproof cost-sharing mechanism for the Steiner tree game – a special case of the game considered here. Such a result for Steiner forest games has proved to be elusive so far, in stark contrast to the well known primal-dual $(2-1/k)$-approximate algorithms for the problem. The cost-sharing method presented in this paper is 2-approximate budget-balanced. Our algorithm is an original extension of known primal-dual methods for Steiner forests. An interesting byproduct of the work in this paper is that our Steiner forest algorithm is $(2-1/k)$-approximate despite the fact that the forest computed by our method is usually costlier than those computed by known primal-dual algorithms. In fact the dual solution computed by our algorithm is infeasible but we can still prove that its total value is at most the cost of a minimum-cost Steiner forest for the given instance.

In [25] we consider a game-theoretical variant of the Steiner forest problem in which each player j, out of a set of k players, strives to connect his terminal pair (s_j, t_j) of vertices in an undirected, edge-weighted graph G. In this paper we design an approximately budget balanced and cross-monotonic cost sharing method for this game. we show that a natural adaptation of the primal-dual Steiner forest algorithm of Agrawal, Klein and Ravi [When trees collide: An approximation algorithm for the generalized Steiner problem in networks, *SIAM Journal on Computing*, 24(3):445–456, 1995] yields a 2-budget balanced and cross-monotonic cost sharing method for this game. We also present a negative result, arguing that no cross-monotonic cost sharing method can achieve a budget balance factor of less than 2 for the Steiner tree game. This shows that our result is tight. Our algorithm gives rise to a new linear programming relaxation for the Steiner forest problem which we coin the *lifted-cut relaxation*. We show that this new relaxation is strictly stronger than the standard undirected cut relaxation for the Steiner forest problem. An interesting byproduct of this work is that our algorithm to compute a Steiner forest is $(2-1/k)$-approximate despite the fact that the total dual produced by our algorithm is usually larger than those computed by known primal-dual algorithms. In fact the dual solution computed by our algorithm is infeasible for the undirected cut relaxation but we can still prove that its value is at most the cost of a minimum-cost Steiner forest for the given instance.

A Black Hole is a highly harmful stationary process residing in a node of a network and destroying all mobile agents visiting the node without leaving any trace. The Black Hole Search is the task of locating all black holes in a network, through the exploration of its nodes by a set of mobile agents. In [11] we consider the task of locating a black hole in a (partially) synchronous arbitrary network, assuming an upper bound on the time of any edge traversal by an agent. For a given graph and a given starting node we are interested in finding the fastest possible Black Hole Search by two agents (the minimum number of agents capable to identify a black hole). We prove that this problem is NP-hard in arbitrary graphs, thus solving an open problem stated in [Czyzowicz,Kowalski,Markou,Pelc-2004]. We also give a 7/2-approximation algorithm, thus improving on the 4-approximation scheme observed there. Our approach is to explore the given input graph via some spanning tree. Even if it represents a very natural technique, we prove that this approach cannot achieve an approximation ratio better than 3/2.

In [24] we consider the problem of designing the fastest Black Hole Search, given the map of the network, the starting node and, possibly, a subset of nodes of the network initially known to be safe. We study the version of this problem that assumes that there is at most one black hole in the network and there are two agents, which move in synchronized steps. We prove that this problem is not polynomial-time approximable within $\frac{389}{388}$ (unless P=NP). We give a 6-approximation algorithm, thus improving on the 9.3-approximation algorithm from Czyzowitz et al. 2004. We also prove APX-hardness for a restricted version of the problem, in which only the starting node is initially known to be safe.

In computability and in complexity theory reductions are widely used for mapping sets into sets in order to prove undecidability or hardness results. In the study of the approximate solvability of hard discrete optimization problems, suitable kinds of reductions, called approximation preserving reductions, can also be used to transfer from one problem to another either positive results (solution techniques) or negative results (non approximability results). In [6] various kinds of approximation preserving reductions are surveyed and their properties discussed. The role of completeness under approximation preserving reductions is also analyzed and its relationship with hardness of approximability is explained.

In [8] we study completeness in differential approximability classes. In differential approximation, the quality of an approximation algorithm is the measure of both how far is the solution computed from a worst one and how close is it to an optimal one. The main classes considered are DAPX, the differential counterpart of APX, including the NP optimization problems approximable in polynomial time within constant differential approximation ratio and the DGLO, the differential counterpart of GLO, including problems for which their local optima guarantee constant differential approximation ratio. We define natural approximation preserving reductions and prove completeness results for the class of the NP optimization problems (class NPO), as well as for DAPX and for a natural subclass of DGLO. We also define class 0-APX of the NPO problems that are not differentially approximable within any ratio strictly greater than 0 unless P = NP. This class is very natural for differential approximation, although has no sense for the standard one. Finally, we prove the existence of hard problems for a subclass of DPTAS, the differential counterpart of PTAS,theclassofNPO problems solvable by polynomial time differential approximation schemata.

Game theory. Recently paradigms and mathematical concepts from game theory have been adopted in the analysis of resource management policies in large networks in order to maintain control over the degradation of network performances caused by the selfishness of users. In this area we have addressed the following problems. In [16] we consider the problem of Internet switching, where traffic is generated by selfish users. We study a packetized (TCP-like) traffic model, which is more realistic than the widely used fluid model. We assume that routers have First-In-First-Out (FIFO) buffers of bounded capacity managed by the drop-tail policy. The utility of each user depends on its transmission rate and the congestion level. Since selfish users try to maximize their own utility disregarding the system objectives, we study Nash equilibria that correspond to a steady state of the system. We quantify the degradation in the network performance called the price of anarchy resulting from such selfish behavior. We show that for a single bottleneck buffer, the price of anarchy is proportional to the number of users. Then we propose a simple modification of the Random Early Detection (RED) drop policy, which reduces the price of anarchy to a constant.

In [13] we are interested in the complexity of finding Nash equilibria with one uniformly mixed strategy (that is, equilibria in which at least one of the players plays a uniform probability distribution over some set of pure strategies). We show that, even in imitation bimatrix games, where one player has a positive payoff if he plays the same pure strategy as the opponent, deciding the existence of such an equilibrium is an NP-complete problem. We derive this result from the NP-completeness of graph-theoretical problems strictly related to this class of equilibria.

In [14] we investigate the complexity of finding uniformly mixed Nash equilibria (that is, equilibria in which all played strategies are played with the same probability). We show that, even in very simple win/lose bimatrix games, deciding the existence of uniformly mixed equilibria in which the support of one (or both) of the players is at most or at least a given size is an NP-complete problem. Motivated by these results, we also give NP-completeness results for problems related to finding a regular induced subgraph of a certain size or regularity in a given graph, which can be of independent interest.

Experimentation, visualization and applications. In algorithm engineering is particularly relevant to study and analyze the performance of algorithms not only from the theoretical point of view (through the classical worst case asymptotic analysis) but also by running experiments that allow to assess the practical behaviour of algorithms on real life data. In this area we have continued our research program by addressing both the experimental study of specific graph algorithms and the design and development of an environment for algorithm visualization. In [4] we present the results of an extensive computational study on dynamic algorithms for all pairs shortest path problems. We describe our implementations of the recent dynamic algorithms of King and of Demetrescu and Italiano, and compare them to the dynamic algorithm of Ramalingam and Reps and to static algorithms on random, real-world and hard instances. Our experimental data suggest that some of the dynamic algorithms and their algorithmic techniques can be really of practical value in many situations.

In [15] we describe our own experience in preparing animated presentations of computer science concepts with Leonardo Web, a Java-based animation system that we have previously developed. Our discussion is aimed at highlighting how different visualization tools and techniques turned out to be useful in realizing effective Web-based teaching material. Our experience culminated in the preparation of an on-line repository of animated illustrations for a textbook on algorithms and data structures edited by McGraw-Hill in 2004.

Journals

[1] Camil Demetrescu and Giuseppe F. Italiano. Trade-Offs for Fully Dynamic Reachability on DAGs: Breaking Through the $O(n^2)$ Barrier. *Journal of the Association for Computing Machinery (JACM)*, **52**(2), pp. 147–156, 2005.

[2] Luca Becchetti, S. Diggavi, Stefano Leonardi, Alberto Marchetti-Spaccamela, S. Muthukrishnan, T. Nandagopal and Andrea Vitaletti. Parallel scheduling problems in next generation wireless networks. *Networks*, **45**(1), pp. 9–22, 2005.

[3] Luca Becchetti, Stefano Leonardi, Alberto Marchetti-Spaccamela, Guido Schaefer and T. Vredeveld. Average case and smoothed competitive analysis of the multi-level feedback algorithm. *Accepted for publication in Mathematics of Operations Research*, 2005. A preliminary version appeared in proceedings of the IEEE Symposium on Foundations of Computer Science (FOCS), 2003.

[4] Camil Demetrescu and Giuseppe F. Italiano. Experimental Analysis of Dynamic All Pairs Shortest Path Algorithms. *To appear in ACM Transactions on Algorithms*, 2005. Special issue devoted to the best papers selected from the 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'04).

[5] Camil Demetrescu and Giuseppe F. Italiano. Fully Dynamic All Pairs Shortest Paths with Real Edge Weights. *To appear in Journal of Computer and System Sciences*, 2005. Special issue devoted to the best papers selected from the 42th Annual IEEE Symposium on Foundations of Computer Science (FOCS'01).

[6] Giorgio Ausiello and Vangelis Th. Paschos. Reductions, completeness and the hardness of approximability. *To appear European Journal of Operations Research*, 2005.

[7] Giorgio Ausiello, Paolo Giulio Franciosa and Daniele Frigioni. Partially Dynamic Maintenance of Minimum Weight Hyperpaths. *Journal of Discrete Algorithms*, **3**(1), pp. 27–46, 2005.

[8] Giorgio Ausiello, Cristina Bazgan, Marc Demange and Vangelis Th. Paschos. Completeness in differential approximation classes. *International Journal of Foundations of Computer Science*, **16**(6), 2005.

Conference Proceedings

[9] Luca Becchetti, Jochen Koenemann, Stefano Leonardi and M. Pal. Sharing the cost more efficiently: Improved approximation for multicommodity rent-or-buy. *Proceedings of the SIAM-ACM Symposium on Discrete Algorithms (SODA'05)*, pp. 375–384, 2005. Accepted for publication under minor revision in ACM Transaction on Algorithms.

[10] Jochen Koenemann, Stefano Leonardi and Guido Schaefer. A group-strategyproof mechanism for Steiner forests. *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'05)*, 2005.

[11] Ralf Klasing, Euripides Markou, Tomasz Radzik and Fabiano Sarracco. Hardness and approximation results for Black Hole Search in arbitrary graphs. *Proc. of 12th Colloquium on Structural Information and Communication Complexity (SIROCCO 2005) LNCS 3499*, pp. 200–215, 2005.

[12] Giorgio Ausiello, Vincenzo Bonifaci and Luigi Laura. The On-line Asymmetric Traveling Salesman Problem. *Proceedings of the 9th Workshop on Algorithms and Data Structures 2005 (WADS 2005) - LNCS 3608 - Springer*, 2005.

[13] Vincenzo Bonifaci, Ugo Di Iorio and Luigi Laura. New results on the complexity of uniformly mixed Nash equilibria. *Proceedings of the 1st Workshop on Internet and Network Economics (WINE)*, pp. 1023–1032, 2005.

[14] Vincenzo Bonifaci, Ugo Di Iorio and Luigi Laura. On the complexity of uniformly mixed Nash equilibria and related regular subgraph problems. *Proceedings of the 15th International Symposium on Fundamentals of Computation Theory (FCT 2005)*, 2005.

[15] Vincenzo Bonifaci, Camil Demetrescu, Irene Finocchi, Giuseppe F. Italiano and Luigi Laura. Portraying algorithms with Leonardo Web. *Proceedings of the WISE05 International Workshop on Web-based Learning, New York City, USA*, 2005.

[16] Alex Kesselman, Stefano Leonardi and Vincenzo Bonifaci. Game-theoretic analysis of Internet switching with selfish users. *Proceedings of the 1st Workshop on Internet and Network Economics (WINE)*, pp. 236–245, 2005.

[17] Giorgio Ausiello, Paolo Giulio Franciosa and Giuseppe Francesco Italiano. Small Stretch Spanners on Dynamic Graphs. *Proceedings of European Symposium on Algorithms (ESA 2005) - LNCS 3669 - Springer*, 2005.

[18] Luca Allulli, Giorgio Ausiello and Luigi Laura. On the Power of Lookahead in On-line Vehicle Routing Problems. *COCOON 2005 - LNCS 3595 - Springer*, 2005.

[19] Luigi Laura, Umberto Nanni and Fabiano Sarracco. Fast LSI-based Techniques for Query Expansion in Text Retrieval Systems.. *Proc. of 2nd Workshop on Text-based Information Retrieval (TIR-05)*, pp. 15–28, 2005.

[20] Debora Donato, Stefano Leonardi and Panayiotis Tsaparas. Stability and Similarity of Link Analysis Ranking Algorithms. *Proc. of 32nd International Colloquium on Automata, Languages and Programming (ICALP 2005 - Track A).* , 2005.

[21] Debora Donato, Stefano Leonardi, Stefano Millozzi and Panayiotis Tsaparas. The inner structure of the Webgraph. *Proc. of 8th International Workshop on the Web and Databases (WebDB 2005)*, 2005. A short abstract appears in Proc. of the European Conference on Complex Systems (ECCS05).

[22] Luciana Salete Buriol, Debora Donato, Stefano Leonardi and Tobias Matzner. Using Data Streams Algorithms for Computing Properties of Large Graphs. *Booklet of the Workshop on Massive Geometric Data Sets (Massive2005), Pisa*, pp. 9–14, 2005.

[23] Luciana Salete Buriol, Debora Donato, Stefano Leonardi and Tobias Matzner. Data Stream Computation for Monitoring Statistics of Massive Webgraphs. *Proc. of Conference on Complex Systems (ECCS05)*, 2005.

[24] Ralf Klasing, Euripides Markou, Tomasz Radzik and Fabiano Sarracco. Approximation bounds for Black Hole Search problems. *Proc. of 9th International Conference on Principles of Distributed Systems (OPODIS'2005)*, pp. 198–207, 2005.

[25] Jochen Konemann, Stefano Leonardi, Guido Schaefer and Stefan van Zwam. From Primal-Dual to Cost Shares and Back: A Stronger LP Relaxation for the Steiner Forest Problem. *Automata, Languages and Programming, 32nd International Colloquium, ICALP 2005, Lecture Notes in Computer Science 3580 Springer* , pp. 930–942, 2005.

Articles in books

[26] Camil Demetrescu, Irene Finocchi and Giuseppe F. Italiano. Handbook on Data Structures and Applications. Chapter 36: Dynamic Graphs, Dinesh Mehta and Sartaj Sahni (eds.), CRC Press Series, in Computer and Information Science, 2005.

[27] Camil Demetrescu, Irene Finocchi and Giuseppe F. Italiano. Handbook on Data Structures and Applications. Chapter 35: Dynamic Trees, Dinesh Mehta and Sartaj Sahni (eds.), CRC Press Series, in Computer and Information Science, 2005.

Book

[28] Camil Demetrescu, Roberto Tamassia and Robert Sedgewick. ALENEX/ANALCO 2005: 7th Workshop on Algorithm Engineering and Experiments and 2nd Workshop on Analytic Algorithmics and Combinatorics, Proceedings. SIAM, ISBN 0-89871-596-2, 2005.

[29] Gerth Stolting Brodal and Stefano Leonardi. ESA 2005, 13th Annual European Symposium, Proceedings. Lecture Notes in Computer Science 3669. Springer, 2005.

[30] Ozalp Babaoglu, Mark Jelasity, Alberto Montresor, Christof Fetzer, Aad P. A. van Moorse, Stefano Leonardi and Maarten van Steen. Self-star Properties in Complex Information Systems, Conceptual and Practical Foundations. Lecture Notes in Computer Science 3460, Springer, 2005.

Submitted papers, technical reports and others

[31] Luciana Salete Buriol, Debora Donato, Stefano Leonardi and Stefano Millozzi. Link and Temporal Analysis of Wikigraphs. *Technical Report DELIS European Research Project*, 2005.