

05361 Abstracts Collection
Algorithmic Aspects of Large and Complex
Networks
— **Dagstuhl Seminar** —

Stefano Leonardi¹, Friedhelm Meyer auf der Heide² and Dorothea Wagner³

¹ Univ. di Roma "La Sapienza", IT

² Univ. Paderborn, DE

`fmadh@uni-paderborn.de`

³ Univ. Karlsruhe, DE

`dwagner@ira.uka.de`

Abstract. From 04.09.05 to 09.09.05, the Dagstuhl Seminar 05361 “Algorithmic Aspects of Large and Complex Networks” was held in the International Conference and Research Center (IBFI), Schloss Dagstuhl. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Algorithms, Large and Complex Networks

05361 Summary – Algorithmic Aspects of Large and Complex Networks

Information systems like the Internet, the World Wide Web, telecommunication networks, sensor networks, or peer-to-peer networks have reached a size and a dynamics that puts them beyond our ability to deploy them and to keep them functioning correctly through traditional algorithmic techniques. Their complexity arises from their growth and dynamics: They change their size and structure over time, large components may be modified deleted or replaced. Challenges for algorithms research include development of models for the dynamics of such systems, the design of decentralized, scalable and adaptive mechanisms for regulation and improvement, and the design, the analysis, and the experimental evaluation of algorithms for analyzing and optimizing such systems.

Today, networks play an important role in many areas of our society. Information acquisition and dissemination as well as the further growing mobility is based on the understanding and employment of such huge networks. Interesting relations between networks in areas like traffic and transport, politics and society make it promising to strengthen the joint effort to tackle them. Dealing with

large networks is a big topic also in areas like physics, biology, and economy, with very different, mainly non-algorithmic methods and different optimization goals.

The seminar brought together 46 researchers from different European countries, Australia and USA. Most participants were from algorithms research, but several also came from other branches of Computer Science as well as from Physics and Biology. As in previous meetings there were not only scientific talks on results respectively ongoing research, but also fruitful and stimulating discussions.

The seminar also offered an opportunity to meet and intensify collaboration for researchers involved in the DFG research cluster "Algorithmic Aspects of Large and Complex Networks" (DFG Schwerpunktprogramm 1126) and from the EU Integrated Project DELIS (Dynamically Evolving Large-scale Information Systems).

Schloss Dagstuhl and its staff provided a very convenient and stimulating environment. The seminar participants appreciated the cordial atmosphere which improved mutual understanding and inspiration. The organizers wish to thank all those who helped make the workshop a fruitful research experience.

On Nash Equilibria for a Network Creation Game

Susanne Albers (Universität Freiburg, D)

We study a network creation game recently proposed by Fabrikant, Luthra, Maneva, Papadimitriou and Shenker. In this game, each player (vertex) can create links (edges) to other players at a cost of α per edge. The player's goal is to minimize the sum consisting of (a) the cost of the links he has created and (b) the sum of the distances to all other players.

Fabrikant et al. conjectured that there exists a constant A such that, for any $\alpha > A$, all non-transient Nash equilibria graphs are trees. In this paper we disprove the tree conjecture. More precisely, we show that for any positive integer n_0 , there exists a graph built by $n \geq n_0$ players which contains cycles and forms a non-transient Nash equilibrium, for any α with $1 < \alpha \leq \sqrt{n/2}$. Our construction makes use of some interesting results on finite affine planes. On the other hand we show that for $\alpha \geq 12n \log n$ every Nash equilibrium forms a tree.

Fabrikant et al. proved an upper bound on the price of anarchy of $O(\sqrt{\alpha})$ where $\alpha \in [2, n^2]$. We improve this bound for every α .

Specifically, we derive a constant upper bound for $\alpha \leq \sqrt{n}$ and for $\alpha \geq 12n \log n$. For the intermediate values we derive an improved bound of $O(1 + (\min\{\frac{\alpha^2}{n}, \frac{n^2}{\alpha}\})^{1/3})$.

Additionally, we develop characterizations of Nash equilibria and extend our results to a weighted network creation game as well as to scenarios with cost sharing.

Joint work of: Eilts, Stefan; Even-Dar, Eyal; Mansour, Yishay; Roditty, Liam

Why Spectral Retrieval Works (and when it does not)

Holger Bast (MPI für Informatik, D)

Spectral retrieval is a popular approach to ranked retrieval on large document collections. The basic scheme is as follows. It is assumed that documents as well as queries are represented as high-dimensional vectors, where the dimensions correspond to the distinct words in the collection. Documents and the query are then projected on a low-dimensional eigenspace computed from the document vectors. Documents are then ranked according to their similarity to the query in that eigenspace.

Previous explanations for why spectral retrieval works are all of the following kind: if there are k "base" documents such that each document can be approximated by a linear combination of these, then reduction to an eigenspace of dimension k works. This kind of explanation leaves open the central question, as to what an appropriate choice for k is.

We argue that what makes spectral retrieval work in practice is its ability to identify pairs of terms with similar co-occurrence patterns. We give a parameterless algorithm that on a number of test collections outperforms all previous algorithms committing to a fixed dimension.

Keywords: Web Search, Ranked Retrieval, LSI, Spectral

Full Paper:

<http://doi.acm.org/10.1145/1076034.1076040>

See also: Proc. SIGIR'05 pp. 11–18

Some new procedures in Pajek

Vladimir Batagelj (University of Ljubljana, SLO)

We shall present some newest additions to Pajek - a program for analysis and visualization of large networks.

1. network multiplication: in general the multiplication of large sparse networks is a 'dangerous' operation since the result can 'explode' - it is not sparse. In some cases we can show that it remains sparse and provides us with useful information. One such application is the computation of basic kinship relations from genealogies.

2. bipartite cores: for analysis of 2-mode (bipartite) networks the notion of core can be extended to (p, q) -cores - all degrees in the first set are at least p , and in the second set at least q . The (p, q) -cores can be efficiently determined, but how to select the right p and q ?

3. 4-rings: to identify dense parts in 2-mode networks we can use as the weight of lines the number (and type, in directed networks) of 4-(semi)cycles that contain the line.

The contribution is a joint work with Andrej Mrvar.

Keywords: Large network, algorithm, network multiplication, kinship, bipartite core, 4-rings, (semi)cycles

Joint work of: Batagelj, Vladimir; Mrvar, Andrej

Fully Distributed Link Analysis for a P2P Search Engine

Geoffrey Canright (Telenor Research - Fornebu, N)

The DELIS EU project has as one of its main goals to develop a distributed, peer-to-peer (P2P) search engine. Motivated by this goal, we have developed a fully distributed form of link analysis, in which the nodes themselves cooperatively carry out a power method calculation which gives the dominant eigenvector used for ranking Web pages.

There are two important forms of global knowledge required by the power method:

(a) length of the current vector (needed for normalization); and (b) a way to avoid the problem of sinks (which arise in directed graphs). We present distributed solutions for both of these problems, and test our solutions for various forms of link analysis. For example, we perform a fully distributed PageRank calculation, so that each page learns its own PageRank score using only local information and local (ie, with neighbors) interactions. Other kinds of link analysis (non-normalized operators, and/or undirected links) can also be handled by our method.

Joint work of: Canright, Geoffrey; Engø-Monsen, Kenth; Jelasity, Mark

Trading off space for passes in graph streaming problems

Camil Demetrescu (Università di Roma I, I)

A natural question in data stream processing is whether we can reduce the space usage, possibly at the price of increasing the number of passes over the data. Unfortunately, this seems to be very hard for many graph problems, unless powerful primitives such as sorting are used. And yet, even using sorting some problems such as shortest paths seem to be difficult.

In this talk we show that, without using sorting, for any space restriction of s bits, single-source shortest paths in directed graphs with small positive integer edge weights can be solved in $O((n \log^{3/2} n)/\sqrt{s})$ passes. For undirected connectivity, we devise an $O((n \log n)/s)$ passes algorithm. Both problems require $\Omega(n/s)$ passes under the restrictions we consider. Our algorithms work in the W-Stream model introduced by Aggarwal et al. in FOCS 2004.

Keywords: Graph algorithms, data streams, space/passes tradeoffs, shortest paths, connectivity

Joint work of: Demetrescu, Camil; Finocchi, Irene; Ribichini, Andrea

Towards a mathematical model of communication for mobile agents

Josep Diaz (TU of Catalonia - Barcelona, E)

We study the distributions in the static and dynamic case when w agents are springled in a grid with N nodes and start moving with a communicating distance d

Joint work of: Diaz, Josep; Perez, Xavier; Serna, Maria; Wormald, Nick

Deterministic boundary recongnition and topology extraction for large sensor networks

Sándor Fekete (TU Braunschweig, D)

We present a new framework for the crucial challenge of self-organization of a large sensor network. The basic scenario can be described as follows: Given a large swarm of immobile sensor nodes that have been scattered in a polygonal region, such as a street network. Nodes have no knowledge of size or shape of the environment or the position of other nodes. Moreover, they have no way of measuring coordinates, geometric distances to other nodes, or their direction. Their only way of interacting with other nodes is to send or to receive messages from any node that is within communication range. The objective is to develop algorithms and protocols that allow self-organization of the swarm into large-scale structures that reflect the structure of the street network, setting the stage for global routing, tracking and guiding algorithms.

Our algorithms work in two stages: boundary recognition and topology extraction. All steps are strictly deterministic, yield fast distributed algorithms, and make no assumption on the distribution of nodes in the environment, other than sufficient density.

Keywords: Distributed algorithms, sensor networks, boundary recognition, topology extraction

Joint work of: Fekete, Sándor; Kröller, Alexander; Pfisterer, Dennis; Fischer, Stefan

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2006/563>

Full Paper: <http://arxiv.org/abs/cs.DC/0510048>

See also: To appear in: SODA 2006, paper available at arxiv.org

A Hybrid Model for Drawing Dynamic and Evolving Graphs

Marco Gaertler (Universität Karlsruhe, D)

Dynamic processes frequently occur in many applications. Visualizations of dynamically evolving data, for example as part of the data analysis, are typically restricted to a cumulative static view or an animation/sequential view. Both methods have their benefits and are often complementary in their use. We present a hybrid model that combines the two techniques. This is accomplished by 2.5D drawings which are calculated in an incremental way. The method has been evaluated on collaboration networks.

Keywords: Visualization dynamic/evolving graphs 2.5D

Joint work of: Gaertler, Marco; Wagner, Dorothea

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2006/568>

Full Paper:

<http://i11www.iti.uka.de/algo/members/gaertler/publications/files/gw-hmdde-05.pdf>

See also: Proceedings of the 13th International Symposium on Graph Drawing (GD'05), Lecture Notes in Computer Science, pages 189 - 200, 2006.

Friends for Free: Self-Organizing Artificial Social Networks for Trust and Cooperation

David Hales (Università di Bologna, I)

By harvesting friendship networks from e-mail contacts or instant message “buddy lists” Peer-to-Peer (P2P) applications can improve performance in low trust environments such as the Internet. However, natural social networks are not always suitable, reliable or available. We propose an algorithm (SLACER) that allows peer nodes to create and manage their own friendship networks. We evaluate performance using a canonical test application, requiring cooperation between peers for socially optimal outcomes. The Artificial Social Networks (ASN) produced are connected, cooperative and robust - possessing many of the desirable properties of human friendship networks such as trust between friends (directly linked peers) and short paths linking everyone via a chain of friends. In addition to new application possibilities, SLACER could supply ASN to P2P applications that currently depend on human social networks thus transforming them into fully autonomous, self-managing systems.

Keywords: Evolution of cooperation, Evolving Networks, Prisoner’s Dilemma, P2P, Tags

Joint work of: Hales, David; Arteconi, Stefano

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2006/565>

Full Paper:

<http://www.davidhales.com>

Approximating the Distortion

Alexander Hall (ETH Zürich, CH)

Kenyon et al. (STOC 04) compute the distortion between one-dimensional finite point sets when the distortion is small; Papadimitriou and Safra (SODA 05) show that the problem is NP-hard to approximate within a factor of 3, albeit in 3 dimensions. We solve an open problem in these two papers by demonstrating that, when the distortion is large, it is hard to approximate within large factors, even for 1-dimensional point sets. We also introduce additive distortion, and show that it can be easily approximated within a factor of two.

Keywords: Distortion, Mapping Point Sets, Inapproximability, Approximation Algorithm

Joint work of: Hall, Alexander; Papadimitriou, Christos

Full Paper:

http://dx.doi.org/10.1007/11538462_10

See also: Proceedings of the 8th Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX 05), LNCS 3624, Springer-Verlag, Berkeley, USA, pages 111-122, August 2005.

Visualisation and Analysis of Large and Complex Scalefree Networks

Seok-Hee Hong (The University of Sydney, AU)

Scale-free networks appear in many application domains such as social networks and biological networks. Roughly speaking, scale-free networks have power law degree distribution and ultra short average path length and high clustering coefficient.

This talk presents new visualisation methods of scale-free networks in three dimensions. To make effective use of the third dimension with minimum occlusion, we use the concept of "layered drawing" and "concentric spheres", constraining nodes to lie on parallel planes or the surface of spheres, based on the degree of nodes.

We implement the algorithms using a variation of the fast force-directed graph layout method, FADE. Experimental results with real world data sets such

as IEEE InfoVis citation network and collaboration network and protein-protein interaction networks show that our method can be useful for visual analysis of large and complex scale-free networks. We further present methods for visualisation of evolution networks and network integration.

Keywords: Visualisation

Maximal cliques in generalized interval graphs

Michael Kaufmann (Universität Tübingen, D)

Driven by an application from bioinformatics (clustering), we consider the problem to find all maximal cliques in a graph of overlapping intervals. Two intervals induce an edge if they symmetrically overlap by at least a certain percentage of their lengths. It turns out that the graph is a special case of the so-called max-tolerance graphs.

Using a geometric characterization as intersection graphs of semi-squares, we are able to give an optimal bound on the number of maximal cliques and an efficient algorithm to construct them.

This is joint work with Katharina Lehmann and Amarendran Subramanian. The results will appear in the proceedings of SODA 2006.

Keywords: Maximal cliques, max-tolerance graphs

Phase Transitions in Satisfiability and Coloring

Lefteris M. Kirousis (University of Patras, GR)

Various techniques for estimating thresholds of decision problems like k -satisfiability or k -coloring will be investigated. Special attention will be given to the Second Moment Method, which was recently introduced as a technique for threshold estimation.

Force-Directed Approaches to Sensor Network Localization

Stephen G. Kobourov (University of Arizona, USA)

In many sensor network applications it is necessary to compute low-error localization of the sensor nodes. Although embedding a GPS unit on each node would solve the problem for many outdoor applications, the cost of this solution for large networks is prohibitively high.

We consider static and mobile network localization approaches that make use of the local neighborhood information, in the form of relative distances and angles to nearby nodes, gathered through simpler and less costly devices (RF,

ultrasound based range sensors, or antenna arrays). Our algorithms do not make any assumptions about the existence of anchor nodes capable of locating themselves, nor about the knowledge of an initial localization to start with. Instead, we rely on a multi-scale force-directed approach, utilizing range and angle data through dead reckoning. We show that our localization algorithms are robust and scale well with network size.

Keywords: Sensor network localization, multi-scale force-directed approach, dead reckoning

Joint work of: Kobourov, Stephen G.; Efrat, Alon; Forrester, David; Iyer, Anand

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2006/569>

Robust Self-Organization of Single Hop Ad Hoc Networks

Mirosław Kutylowski (Institute of Mathematics & Informatics/TU Wrocław, PL)

We concern ad hoc networks of devices communicating through a shared communication channel. Our task is to build a certain structure of the network after it is deployed so that no device has any knowledge on the other nodes that are active.

Even if efficient algorithms for the problems like size approximation of the network, leader election or assigning consecutive numbers to network nodes are known, a significant drawback of this line of research is that it disregards communication failures, or even worse, the failures due to an adversary.

We develop algorithmic tools that provide efficient solutions for the problems mentioned for the model with an adversary having energy resources such as a regular user.

Keywords: Radio network, single hop, size approximation, leader election, initialization, adversary

Joint work of: Kutylowski, Mirosław; Rutkowski, Wojtek; Kabarowski, Jędrzej

Full Paper:

<http://kutyłowski.im.pwr.wroc.pl/bib-html.html#adhoc>

Using Graph Drawing to Search the Web

Giuseppe Liotta (University of Perugia, I)

One of the most challenging issues in mining information from the World Wide Web is the design of systems that can present the data to the end user by clustering them into meaningful semantic categories.

We envision that the analysis of the results of a Web search can significantly take advantage of advanced graph drawing techniques. In this paper we strengthen our point by describing the visual functionalities of WhatsOnWeb. WhatsOnWeb is a meta search clustering engine explicitly designed to make it possible that the user browses the Web by means of drawings of graphs whose nodes represent clusters of coherent data and whose edges describe semantic relationships between pairs of clusters. A prototype of WhatsOnWeb is available at <http://whatsonweb.diei.unipg.it/>.

Keywords: Web Clustering Engines, Graph Visualization Interfaces, Orthogonal Layout

Joint work of: Di Giacomo, Emilio; Didimo, Walter; Grilli, Luca; Liotta, Giuseppe

Full Paper:

<http://whatsonweb.diei.unipg.it/>

Computational Considerations in the Study of Real, Very Large Complex Networks

Madhav Marathe (Virginia Bioinformatics Institute - Blacksburg, USA)

We consider the following questions pertaining to very large, realistic information, social, and infrastructure networks:

- Construction and storage of such networks
- Measurement and analysis of structural properties
- Fast and provable Regeneration of such networks with specified properties
- Scaling laws for Structural and Cultural properties

As specific examples, we will consider methods for constructing and analyzing very large social contact networks that form the basis of epidemiology networks and ad hoc communication networks for next generation wireless systems. Based on our results, we formulate the Cultural and Structural Similarity Hypothesis for Data and Information Transportability.

Latency Constrained Aggregation in Sensor Networks

Alberto Marchetti-Spaccamela (Università di Roma I, I)

Advancements in wireless and sensor technologies have paved the way for the development of tiny and cheap sensor devices equipped with sensors and wireless transceivers. Such devices, named sensor nodes, are able to monitor events, to process the sensed information and to communicate the sensed data to one or more base stations. Being battery powered, sensor networks are highly energy constrained.

Data aggregation is a technique that consists in aggregating redundant or correlated data in order to reduce the overall size of sent data, thus decreasing the network traffic and energy consumption.

In this talk we focus on the impact of data aggregation in terms of energy consumption reduction: we propose a combinatorial optimization problem that models the problem of designing routing algorithms that optimize the energy saving obtained by data aggregation.

Joint work of: Beccchetti, Luca; Marchetti-Spaccamela, Alberto; Skutella, Martin; Stougie, Leen; Vitaletti, Andrea

Cycle Basis of Undirected and Directed Graphs

Kurt Mehlhorn (MPI für Informatik, D)

A cycle basis of an undirected or directed graph is a basis for the set of cycles of the graph. We discuss algorithms for computing minimum weight cycle basis.

Joint work of: Telikapelli, Kavitha; Dimitris, Michail; Kaiza, Paluch; Hariharan, Ramesh

Page Migration in Dynamic Networks

Friedhelm Meyer auf der Heide (Universität Paderborn, D)

Page migration is a basic, simple model for data management in networks: A page of given size D is given, that has to be stored always in exactly one node of a given network. The task is to serve a sequence of requests to data items of size 1 of the page where each request described by the issuing processor. The cost of serving the request is the shortest path length between the issuing processor and the page. The page migration algorithm has to decide, after each served request, if the page is migrated, and, in case of "yes", where it is migrated to. The cost for migration is D times the migration distance. This a very well studied online problem, constant competitive ratio can be achieved.

In this talk, we extend page migration to dynamic networks, where a further input stream dictates network changes. We assume a "speed limit" on the nodes, i.e., each node can move only within a cycle of diameter one in each step. We show that, in this scenario, where the algorithm has to "fight against two adversaries", the competitive ratio is $\Omega(\sqrt{D})$ even for two-node networks. We further show matching upper and lower bounds for deterministic, as well as adaptive and oblivious adversaries. Finally we present results for the cases where one of the input streams is random.

This talk is based on a joint paper with Marcin Bienkowski (invited presentation on MFCS 2005) that surveys results published within the last two years.

Keywords: Page Migration, Online Algorithms, Dynamic Networks

Selfish Routing with Incomplete Information

Burkhard Monien (Universität Paderborn, D)

In his seminal work Harsanyi introduced an elegant approach to study non-cooperative games with incomplete information where the players are uncertain about some parameters. To model such games he introduced the Harsanyi transformation, which converts a game with incomplete information to a strategic game where players may have different types. In the resulting Bayesian game players' uncertainty about each others types is described by a probability distribution over all possible type profiles.

In this work, we introduce a particular selfish routing game with incomplete information that we call Bayesian routing game. Here, n selfish users wish to assign their traffic to one of m links. Users do not know each others traffic. Following Harsanyi's approach, we introduce for each user a set of possible types. This paper presents a comprehensive collection of results for the Bayesian routing game.

- We prove, with help of a potential function, that every Bayesian routing game possesses a pure Bayesian Nash equilibrium. For the model of identical links and independent type distribution we give a polynomial time algorithm to compute a pure Bayesian Nash equilibrium.

- We study structural properties of fully mixed Bayesian Nash equilibria for the model of identical links and show that they maximize individual cost. In general there exists more than one fully mixed Bayesian Nash equilibrium. We characterize the class of fully mixed Bayesian Nash equilibria in the case of independent type distribution.

- We conclude with results on coordination ratio for the model of identical links for three social cost measures, that is, social cost as expected maximum congestion, sum of individual costs and maximum individual cost. For the latter two we are able to give (asymptotic) tight bounds using our results on fully mixed Bayesian Nash equilibria.

To the best of our knowledge this is the first time that mixed Bayesian Nash equilibria have been studied in conjunction with social cost.

Keywords: Selfish routing, bayesian games, existence pure bayesian Nash equilibria, price of anarchy

Joint work of: Gairing, Martin; Monien, Burkhard; Tiemann, Karsten

Full Paper:

<http://doi.acm.org/10.1145/1074000>

See also: Martin Gairing, Burkhard Monien and Karsten Tiemann, *Selfish Routing with Incomplete Information*, Proceedings of SPAA'05, pp. 203–212, 2005.

Hardness and Approximation of Octilinear Steiner Trees

Matthias Müller-Hannemann (TU Darmstadt, D)

Given a point set K of terminals in the plane, the octilinear Steiner tree problem is to find a shortest tree that interconnects all terminals and edges run either in horizontal, vertical, or $\pm 45^\circ$ diagonal direction. This problem is fundamental for the novel octilinear routing paradigm in VLSI design, the so-called X-architecture.

In this talk we present the following results:

We sketch the proof of the NP-completeness of the decision version of the octilinear Steiner tree problem. We also show how to reduce the octilinear Steiner tree problem to the Steiner tree problem in graphs of polynomial size with the following approximation guarantee.

We construct a graph of size $O(\frac{n^2}{\varepsilon^2})$ which contains a $(1 + \varepsilon)$ -approximation of a minimum octilinear Steiner tree for every $\varepsilon > 0$ and $n = |K|$.

We also consider hard and soft obstacles. Hard obstacles have to be completely avoided, whereas for soft obstacles there is a length restriction $L > 0$ which must not be exceeded by any connected subtree lying strictly inside some obstacle. We obtain four different approximation results:

- a) there is a 2-approximation with hard obstacles in $O(n \log^2 n)$.
- b) there is a $(2 + \varepsilon)$ -approximation with soft rectangular obstacles in $O(n^3)$.
- c) there is a $(1.55 + \varepsilon)$ -approximation for hard obstacles .
- d) there is a $(1.55 + \varepsilon)$ -approximation for soft rectangular obstacles.

Joint work of: Schulze, Anna; Müller-Hannemann, Matthias

Keywords: Octilinear Steiner trees, NP-completeness, VLSI design, approximation algorithms, blockages

Iterative Compression for Hard Network Problems

Rolf Niedermeier (Universität Jena, D)

We discuss iterative compression for solving NP-hard graph problems to optimality. This method was introduced by Reed, Smith, and Vetta (2004).

In particular, we explain the method through applications to the problems Vertex Cover and Feedback Vertex Set. We also report on first convincing empirical results when implementing an iterative compression algorithm for Graph Bipartization.

Joint work of: Gramm, Jens; Guo, Jiong; Hueffner, Falk; Wernicke, Sebastian

Highway Hierarchies Hasten Exact Shortest Path Queries

Peter Sanders (Universität Karlsruhe, D)

We present a new speedup technique for route planning that exploits the hierarchy inherent in real world road networks. Our algorithm preprocesses the eight digit number of nodes needed for maps of the USA or Western Europe in a few hours using linear space.

Shortest (i.e. fastest) path queries then take around eight milliseconds to produce exact shortest paths.

This is about 2000 times faster than using Dijkstra’s algorithm.

Full Paper:

<http://www.dominik-schultes.de/hwy/esaHwyHierarchies.pdf>

Peer-to-Peer Networks based on Random Transformations of Connected Regular Undirected Graphs

Christian Schindelhauer (Universität Paderborn, D)

We present k -Flipper, a graph transformation algorithm that transforms regular undirected graphs. Given a path of $k + 2$ edges it interchanges the end vertices of the path. By definition this operation preserves regularity and connectivity. We show that every regular connected graph can be reached by a series of these operations for all $k \geq 1$. We use a randomized version, called Random k -Flipper, in order to create random regular connected undirected graphs that may serve as a backbone for peer-to-peer networks. We prove for degree $d \in \Omega(\log n)$ that a series of $O(dn)$ Random k -Flipper operations with $k \in \Theta(d^2 n^2 \log 1/\epsilon)$ transforms any graph into an expander graph with high probability, i.e. $1 - n^{-\Theta(1)}$.

The Random 1-Flipper is symmetric, i.e. the transformation probability from any labeled d -regular graph G to G' is equal to those from G' to G . From this and the reachability property we conclude that in the limit a series of Random 1-Flipper operations converges against an uniform probability distribution over all connected labeled d -regular graphs. For degree $d \in \omega(1)$ growing with the graph size this implies that iteratively applying Random 1-Flipper transforms any given graph into an expander asymptotically almost surely.

We use these operations as a maintenance operation for a peer-to-peer network based on random regular connected graphs that provides high robustness and recovers from degenerate network structures by continuously applying these random graph transformations. For this, we describe how network operations for joining and leaving the network can be designed and how the concurrency of the graph transformations can be handled.

Keywords: Random Graphs, Peer-to-Peer Networks, Graph Transformations

Joint work of: Schindelhauer, Christian; Mahlmann, Peter

Full Paper:

<http://wwwcs.uni-paderborn.de/fachbereich/AG/agmadh/WWW/schindel/pubs/flipper.pdf>

See also: Peter Mahlmann, Christian Schindelhauer, Peer-to-Peer Networks based on Random Transformations of Connected Regular Undirected Graphs, 17th ACM Symposium on Parallelism in Algorithms and Architectures 2005, 155-164 (SPAA 2005)

Some Remarks on Handling Data Uncertainty in Network Optimization Problems

Rüdiger Schultz (Universität Duisburg-Essen, D)

Multistage stochastic programs are flexible tools to model data uncertainty in optimization problems. When adding risk aversion to conventional risk neutral models non-standard features arise. Additional integer variables and further coupling may occur. In the talk we illustrate these issues at network optimization problems. We present a novel solution algorithm combining relaxation of non-anticipativity with coordinated branching and controlled subproblem solving. First computational results indicate superiority of this approach to applying general-purpose mixed-integer linear programming algorithms and software.

Keywords: Multistage stochastic integer programs, network optimization

Computing earliest arrival flows with multiple sources

Martin Skutella (Universität Dortmund, D)

Earliest arrival flows are motivated by applications related to evacuation. Given a network with capacities and transit times on the arcs, a subset of source nodes with supplies and a sink node, the task is to send the given supplies from the sources to the sink "as quickly as possible". The latter requirement is made more precise by the earliest arrival property which requires that the total amount of flow that has arrived at the sink is maximal for all points in time simultaneously.

It is a classical result from the 1970s that, for the special case of a single source node, earliest arrival flows do exist and can be computed by essentially applying the Successive Shortest Path Algorithm for min-cost flow computations. While it has previously been observed that an earliest arrival flow still exists for multiple sources, the problem of computing one efficiently has been open. We present an exact algorithm for this problem whose running time is strongly polynomial in the input plus output size of the problem.

Keywords: Networks, flows over time, dynamic flows, earliest arrival, evacuation

Joint work of: Baumann, Nadine; Skutella, Martin

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2006/567>

Coresets in dynamic geometric data streams

Christian Sohler (Universität Paderborn, D)

A dynamic geometric data stream consists of a sequence of m insert/delete operations of points from the discrete space $\{1, \dots, \Delta\}^d$. We develop streaming $(1+\epsilon)$ -approximation algorithms for k -median, k -means, MaxCut, maximum weighted matching (MaxWM), maximum travelling salesperson (MaxTSP), maximum spanning tree (MaxST), and average distance over dynamic geometric data streams. Our algorithms maintain a small weighted set of points (a coreset) that approximates with probability $2/3$ the current point set with respect to the considered problem during the m insert/delete operations of the data stream. They use $\text{poly}(\epsilon^{-1}, \log m, \log \Delta)$ space and update time per insert/delete operation for constant k and dimension d . Having a coreset one only needs a fast approximation algorithm for the weighted problem to compute a solution quickly. In fact, even an exponential algorithm is sometimes feasible as its running time may still be polynomial in n .

Keywords: Streaming Algorithms, Computational Geometry, Data Structures

See also: 37th ACM Symposium on Theory of Computing (STOC), pp. 209-217, 2005.

A Cost Mechanism for Fair Pricing of Resource Usage

Paul G. Spirakis (CTI - Patras, GR)

We propose a simple and intuitive cost mechanism which assigns costs for the competitive usage of m resources by n selfish agents. Each agent has an individual demand; demands are drawn according to some probability distribution. The cost paid by an agent for a resource she chooses is the total demand put on the resource divided by the number of agents who chose that same resource. So, resources charge costs in an equitable, fair way, while each resource makes no profit out of the agents.

We call our model the Fair Pricing model. Its fair cost mechanism induces a non-cooperative game among the agents. To evaluate the Nash equilibria of this game, we introduce the Diffuse Price of Anarchy, as an extension of the Price of Anarchy that takes into account the probability distribution on the demands. We prove:

- (1) Pure Nash equilibria may not exist, unless all chosen demands are identical; in contrast, a fully mixed Nash equilibrium exists for all possible choices of the demands. Further on, the fully mixed Nash equilibrium is the unique Nash equilibrium in case there are only two agents.
- (2) In the worst-case choice of demands, the Price of Anarchy is $\Theta(n)$; for the special case of two agents, the Price of Anarchy is less than $2 - \frac{1}{m}$.

- (3) Assume now that demands are drawn from a bounded, independent probability distribution, where all demands are identically distributed and each is at most a (universal for the class) constant times its expectation. Then, the Diffuse Price of Anarchy is at most that same constant, which is just 2 when each demand is distributed symmetrically around its expectation.

Keywords: Cost Sharing, Diffuse Price of Anarchy, Fair Pricing, Resources

Joint work of: Mavronicolas, Marios; Panagopoulou, Panagiota; Spirakis, Paul G.

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2006/564>

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2006/566>

Randomized Partial Covering in Hypergraphs (with application to 3D seed reconstruction in cancer radiation)

Anand Srivastav (Universität Kiel, D)

In this talk we give randomized approximation algorithms for partial set and vertex covering problems in hypergraphs. Previous work of Gandhi, Khuller, Srinivasan (2001) considered the graph case only in case of vertex cover. In particular we give a 3-approximation algorithm for the partial vertex covering problem in a 3-bounded hypergraph. Combining matching and covering conditions the problem of partial set covering with matching constraints is also discussed which is motivated by the 3D seed reconstruction: in brachytherapy of prostate cancer small radioactive implantants are placed in the organ. In this way internal radiation of the tumor is possible. It is an important task to reconstruct the location of the implantants (seeds) by taking 3-Xray fotos from 3 different angles, as the seeds move due to movement of the organ resp. patient. The problem can be directly modelled as a kind of minimum weight perfect matching problem in a 3-uniform hypergraph.

Keywords: Approximation algorithms, randomization, set and vertex cover, matching, hypergraphs

Stability and Similarity of Link Analysis Ranking Algorithms

Panayiotis Tsaparas (University of Helsinki, FIN)

Recently, there has been a surge of research activity in the area of *Link Analysis Ranking*, where hyperlink structures are used to determine the relative *authority* of Web pages.

One of the seminal works in this area is that of Kleinberg, who proposed the HITS algorithm. In this paper, 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, 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.

Joint work of: Donato, Debora; Leonardi, Stefano; Panayiotis, Tsaparas

Full Paper:

<http://www.cs.helsinki.fi/u/tsaparas>

Mechanism Design for Routing Unsplittable Flow

Berthold Vöcking (RWTH Aachen, D)

This talk deals with the design of efficiently computable incentive compatible, or truthful, mechanisms for combinatorial optimization problems with multi-parameter agents. We focus on approximation algorithms for the NP-hard packing problems like routing unsplittable flow problem that can be represented in form of a packing integer problem (PIP). The well known method to solve PIPs in an approximate fashion is LP-based randomized rounding. This method, however, does not satisfy the monotonicity properties required to ensure incentive compatibility.

We show that primal-dual greedy algorithms achieve almost the same approximation ratios for PIPs as randomized rounding. The advantage is that these algorithms are inherently monotone. This way, we can significantly improve the approximation ratios of truthful mechanisms for various fundamental mechanism design problems like single-minded combinatorial auctions (CAs), unsplittable flow routing and multicast routing.

Our approximation algorithms can also be used for the winner determination in CAs with general bidders specifying their bids through an oracle.

Keywords: Approximation algorithms, mechanism design, routing, unsplittable flow

Joint work of: Vöcking, Berthold; Briest, Patrick; Krysta, Piotr

The Spectra of Popular Hypercubic Networks

Rolf Wanka (Universität Erlangen, D)

We investigate the spectra of popular hypercubic networks. We exactly calculate the spectra of the d -dimensional Butterfly network without and with wrap-around edges (including the multiplicities of the eigenvalues).

Furthermore, we characterize the spectra of the Cube-Connected Cycles network and of the Shuffle Exchange network in terms of cycles with self-loops that have weights from $\{-1, 1\}$. In particular, we show that the spectra of the Cube-Connected Cycles and the Shuffle-Exchange graph are almost identical.

Keywords: Graph spectra, eigenvalues, Butterfly, Shuffle-Exchange, Cube-Connected Cycles

Joint work of: Schmidt, Gunnar; Wanka, Rolf