

Micah Adler, Friedhelm Meyer auf der Heide,
Dorothea Wagner (editors):

Algorithmic Aspects of Large and Complex Networks

Dagstuhl Seminar Report; 321
17.09.–21.09.2001 (01381)

INTERNATIONALES BEGEGNUNGS- UND
FORSCHUNGSZENTRUM FÜR INFORMATIK

Schloß Dagstuhl

Seminar 01381 Report 321

Algorithmic Aspects of Large and Complex Networks

September 17 – 21, 2001

O V E R V I E W

Large and complex networks play a central role in a variety of today's and future technologies. Communication, information broadcast, and mobile information systems, as well as political and social acting is modeled in terms of such networks. In order to analyze them, the algorithmic aspects of methods to explore them and to derive necessary new information from them have to be understood thoroughly. Scientists from Computer Science, Mathematics, Electrical Engineering and Humanities use these networks, but apply different methods to solve their specific problems.

Algorithmic problems occurring in the design, the analysis, and the application of such networks have been the topic of this seminar. The 33 participants covered a wide area of aspects of this research field, e. g., graph theoretical and combinatorial foundations, graph algorithms, design and analysis of networks via optimization and approximation techniques, reliability and security aspects, and applications in parallel and distributed computing and social networks.

The outstanding environment and organization of Schloß Dagstuhl greatly contributed to the success of the seminar.

Organizers:

Micah Adler (University of Massachusetts, Amherst, MA)

Friedhelm Meyer auf der Heide (Universität Paderborn)

Dorothea Wagner (Universität Konstanz)

Reported by Rolf Wanka

Participants

Rudolf Ahlswede, Universität Bielefeld
Susanne Albers, Universität Freiburg
Andreas Brandstädt, Universität Rostock
Leslie Ann Goldberg, University of Warwick
Sven Grothklags, Universität Paderborn
Horst Hamacher, Universität Kaiserslautern
Klaus Jansen, Universität Kiel
Michael Kaufmann, Universität Tübingen
Georg Kliewer, Universität Paderborn
Spyros Kontogiannis, MPI für Informatik, Saarbrücken
Matthias Krause, Universität Mannheim
Piotr Krysta, MPI für Informatik, Saarbrücken
Ulrich Meyer, MPI für Informatik, Saarbrücken
Friedhelm Meyer auf der Heide, Universität Paderborn
Rolf H. Möhring, TU Berlin
Stefan Näher, Universität Trier
Michael Paterson, University of Warwick
Tobias Polzin, MPI für Informatik, Saarbrücken
André Pönitz, Hochschule Mittweida
Peter Sanders, MPI für Informatik, Saarbrücken
Rüdiger Schultz, Universität-GH Duisburg
Jop F. Sibeyn, Umea University
Christian Sohler, Universität Paderborn
Anand Srivastav, Universität Kiel
Peter Tittmann, Hochschule Mittweida
Siavash Vahdati Daneshmand, Universität Mannheim
Berthold Vöcking, MPI für Informatik, Saarbrücken
Dorothea Wagner, Universität Konstanz
Rolf Wanka, Universität Paderborn
Ingo Wegener, Universität Dortmund
Thomas Willhalm, Universität Konstanz
Philipp Wölfel, Universität Dortmund
Uwe Zimmermann, TU Braunschweig

Contents

RUDOLF AHLWEDE

Remarks about the ZIF Project: General Theory of Information Transfer and Combinatorics

SUSANNE ALBERS

Caching in Large Networks

ANDREAS BRANDSTÄDT

Modular Decomposition, Clique Width and Efficient Graph Algorithms

LESLIE ANN GOLDBERG

Stability of Work-Stealing Algorithms for Load Balancing

SVEN GROTHKLAGS

Airline Fleet Assignment

HORST W. HAMACHER

Dynamic Network Flow Models and their Applications in Emergency Building Evacuation

KLAUS JANSEN

On Preemptive Resource Constrained Scheduling: Polynomial-Time Approximation Schemes

GEORG KLIEWER

Network Design: Modeling and Solving in the Airline Context

MATTHIAS KRAUSE

BDD-based Cryptanalysis of LFSR-based Stream Ciphers

ULRICH MEYER

On the Average-Case Performance of Parallel Shortest-Paths Algorithms

ROLF MÖHRING

Path-based Algorithms for Multicommodity Flow Problems in Large Networks

STEFAN NÄHER

Large Graphs with LEDA – Data Structures for Static Graphs

PETER SANDERS

Graph Algorithms Between Theory and Practice

RÜDIGER SCHULTZ

Gas Pipeline Optimization under Uncertainty

JOP F. SIBEYN

Semi-External Depth First Search

CHRISTIAN SOHLER

Abstract Combinatorial Programming and Efficient Property Testing

ANAND SRIVASTAV

Fast Approximation of Multi-cast Routing with Congestion Minimization

PETER TITTMANN

Computing Network Reliability in Graphs of Restricted Path Width

SIAVASH VAHDATI

Practical Algorithms for the Steiner Problem on Networks

BERTHOLD VÖCKING

Multiple-Choice Algorithms

DOROTHEA WAGNER

Analysis and Visualization of Social Networks

ROLF WANKA

Parallel Bridging Models and Their Impact on Algorithm Design

INGO WEGENER

On OBDD-based Graph Algorithms

THOMAS WILLHALM

On the Single-Source Shortest Path Problem in Embedded Graphs

Abstracts

Remarks about the ZIF Project: General Theory of Information Transfer and Combinatorics

by RUDOLF AHLWEDE

In the period 1.10.2001 – 30.9.2004 we organize a project with this title in the *Zentrum für interdisziplinäre Forschung*, NRW, in Bielefeld.

Some guiding ideas have been explained. In particular, also connections to *SFB-Schwerpunkt “Algorithmische Aspekte großer und komplexer Netzwerke”* have been mentioned. Most of its members are encouraged to participate in the preparatory meetings:

I. General theory of information transfer, 18.2. – 23.2.2002.

II. Information in natural sciences, social sciences, humanities and engineering, 25.2. – 2.3.2002.

Basic concepts of gain of information, transfer of information, and storage of information have been illustrated by our work (with others) on hypothesis testing with communication constraints, identification via noisy channels, and the Erdős/Ko/Rado intersection problem from 1938 with a complete solution.

Caching in Large Networks

by SUSANNE ALBERS (joint work with S. Arora and S. Khanna)

We study caching problems that arise in large networks such as the world-wide web. In the *connection caching* problem, we have to maintain a set of persistent TCP connections. We consider a general setting where the connections may incur varying establishment costs. We develop online algorithms that achieve an optimal competitive ratio and, in particular, present strategies that use different amounts of extra communication among network nodes while maintaining open connections. In the *document caching* problem, we have to maintain local user caches containing documents from the web. We develop polynomial time constant factor approximation algorithms that use a small amount of extra space in cache. Our solutions are based on formulating web caching problems as integer linear programs.

Modular Decomposition, Clique Width and Efficient Graph Algorithms

by ANDREAS BRANDSTÄDT

(joint work with F. F. Dragan, V. Giakoumakis, C. T. Hoàng, D. Kratsch, T. Klemmt, H.-O. Le, V. B. Le, S. Mahfud, R. Mosca, J.-M. Vanherpe, T. Szymczak)

Modular decomposition of undirected graphs is known to be a powerful tool for designing efficient graph algorithms, and the modular decomposition tree can be determined in linear time.

The problems Maximum Weight Stable Set (MWS) and Maximum Weight Clique can be efficiently solved on graphs whose prime nodes of the decomposition tree are “simple” for these problems. We give some new examples of graph classes having simple prime graphs.

The concept of clique width of a graph is closely related to modular decomposition, and it is known that algorithmic problems expressible in terms of Monadic Second Order Logic with quantification only over vertex set predicates can be efficiently solved on graph classes of bounded clique width (assuming that a k -expression of the input graph is given). We give new examples of graph classes of bounded clique width.

Moreover, for some classes of unbounded clique width, efficient algorithms for the MWS problem are given which improve and extend recently published results.

Stability of Work-Stealing Algorithms for Load Balancing

by LESLIE ANN GOLDBERG (joint work with Petra Berenbrink and Tom Friedetzky)

In this work we analyse a very simple dynamic work-stealing algorithm. In the work-generation model, there are n generators which are *arbitrarily* distributed among a set of n processors. The distribution of generators is arbitrary — generators may even move at the beginning of each time step. During each time-step, each generator may generate a unit-time task which it inserts into the queue of its host processor. It generates such a task independently with probability λ . After the new tasks are generated, each processor removes one task from its queue and services it. The natural work-stealing algorithm that we analyse is widely used in practical applications and works as follows. During each time step, each *empty* processor (with no work to do) sends a request to a randomly selected other processor. Any *non-empty* processor having received at least one such request in turn decides (again randomly) in favour of one of the requests. The number of tasks which are transferred from the non-empty processor to the empty one is determined by the so-called *work-stealing function* f . In particular, if a processor that accepts a request has ℓ tasks stored in its queue, then $f(\ell)$ tasks are transferred to the currently empty one. A popular work-stealing function is $f(\ell) = \lfloor \ell/2 \rfloor$, which transfers (roughly) half of the tasks. We analyse the *long-term behaviour* of the system as a function of λ and f . We show that the system is *stable* for any constant generation rate $\lambda < 1$ and for a wide class of functions f . Most intuitively sensible functions are included in this class (for example, every function $f(\ell)$ which is $\omega(1)$ as a function of ℓ is included). We give a quantitative description of the functions f which lead to stable systems. Furthermore, we give *upper bounds* on the average system load (as a function of f and n).

Airline Fleet Assignment

by SVEN GROTHKLAGS

For the fleet assignment problem an airline has to decide on the subfleet type for each flight leg (flight without stopover) in the airline’s schedule while maximizing the overall profit of the assignment. We present a MIP formulation of the problem based on a flow network, the time-space network. To be able to solve real-world instances for both the strategic and tactical planning

phase of an airline in acceptable time some exact and heuristic reduction techniques must be applied. Furthermore, we show that our approach can be extended to incorporate certain soft constraints that are relevant in practice. The MIP formulation is compared to a hill climbing and simulated annealing approach based on an efficient local search for the fleet assignment problem. It achieves near optimal solution qualities ($> 99.9\%$) on a number of real-world problems with running times comparable to those of the simulated annealing heuristics (solution quality $\approx 99.5\%$). The hill climbing approach is an order of magnitude faster with slightly worse solution quality ($\approx 98.5\%$).

Dynamic Network Flow Models and their Applications in Emergency Building Evacuation

by HORST W. HAMACHER (joint work with Stevanus A. Tjandra)

The planning of emergency building evacuation consists in sending the occupants of a given building via predetermined evacuation routes to safety. The routes may vary depending on the time at which individual evacuees reach certain parts of the building. Good evacuation plans should evacuate everyone as fast as possible and with the lowest risk.

The problem can be modeled using several kinds of dynamic network flow models. The underlying (static!) flow network G is determined by the architectural features of the building. Rooms, hallways, stairs, etc., are the nodes, doors and other connecting units correspond to arcs. Both nodes and arcs have capacities, arcs have additionally travel times. By duplicating nodes and arcs over a given discrete time horizon T , one can mirror evacuation plans by flows in the time expanded network G_T , the latter being equivalent to dynamic flows in the original, static graph G . By finding a maximal flow in G_T the maximal number of evacuees for T time periods is found, by choosing specific cost the average evacuation time can be minimized, the number of evacuees in each time period can be maximized.

The disadvantage of this approach is that the network is getting very large since the number of nodes in G_T is of the order nT , where n is the number of nodes in the original static network.

We will therefore present algorithms which work on the original, static network and compute maximal dynamic flows, earliest arrival flows, quickest flows, or minimal cost dynamic flows. While these problems have been tackled in the past by some authors for the case of constant capacities, we consider the case where capacities may vary over time – a case which is highly relevant in the evacuation application due to the impact of hazards like fire, smoke, etc.

On Preemptive Resource Constrained Scheduling: Polynomial-Time Approximation Schemes

by KLAUS JANSEN (joint work with Lorant Porkolab)

We study resource constrained scheduling problems where the objective is to compute feasible preemptive schedules minimizing the makespan and using no more resources than what are available. We present approximation schemes along with some inapproximability results showing how the approximability of the problem changes in terms of the number of resources. The results are based on linear programming formulations (though with exponentially many variables) and some interesting connections between resource constrained scheduling and (multi-dimensional,

multiple-choice, and cardinality constrained) variants of the classical knapsack problem. Finally we present applications of the above results in fractional graph coloring and multiprocessor task scheduling.

Network Design: Modeling and Solving in the Airline Context

by GEORG KLIEWER (joint work with Achim Koberstein)

The airline network design problem appears in the long-term planning phase of an airline. The goal is to decide for a given flight network which arcs can be eliminated or added. The proposed solution must take the passenger flow in the network into account and also the aircraft flow of the airline. For each arc a binary design decision variable is defined. The passenger flow is modeled as a minimum cost multicommodity flow. The model for the aircraft flow considers different aircraft types, balancing conditions, aircraft availability, etc. We work with path-based multi-commodity flow models because of the fact that a passenger travel path is strictly constrained, e.g., the path length cannot be too long. Our experiments include CPLEX-based solution procedures and also a metaheuristic simulated annealing algorithm. We discuss results obtained for the flight network of Lufthansa. A future work direction is to use the network design decisions for the airline alliance flight network planning.

BDD-based Cryptanalysis of LFSR-based Stream Ciphers

by MATTHIAS KRAUSE

Many stream ciphers occurring in practice produce their output bit stream according to a rule $y = C(L(x))$, where $L(x)$ denotes an internal linear bit-stream produced by a fixed number of linear feedback shift registers starting from a secret initial state $x = (x_0, \dots, x_{n-1})$, and C denotes some nonlinear compression function. We present an algorithm for computing the secret key x from a given output bitstream y of length $\approx n$, which uses Free Binary Decision Diagrams (FBDDs), a data structure for minimizing and manipulating Boolean functions. We show that if the decision whether “ $C(z) = y$?” can be performed by polynomial size FBDDs (this is usually the case), the effective key-length of the cipher is bounded by $\frac{1-\alpha}{1+\alpha}n$, where α denotes the information rate (per bit) which y reveals about the internal bit-stream z . This yields the best known upper bounds on the effective key length for several stream ciphers of practical use, for instance a $0.656n$ upper bound on the effective key length of the self-shrinking generator, a $0.6364n$ upper bound on the effective key length of the A5-generator, used in the GSM-standard, a $0.6n$ upper bound on the effective key length of E_0 encryption standard in the one level mode, and a $0.8823n$ upper bound on the effective key length of E_0 in the two level mode, as it is used in the Bluetooth wireless LAN system.

On the Average-Case Performance of Parallel Shortest-Paths Algorithms

by ULRICH MEYER

We study parallel algorithms for the Single-Source Shortest-Path (SSSP) problem on graphs with n nodes, m edges and random independent edge weights uniformly distributed in $[0, 1]$.

The new parallel approaches are derived from our sequential adaptive bucket-splitting SSSP algorithm which was the first to achieve linear $O(n + m)$ average-case time on arbitrary directed graphs. However, direct parallelizations of the bucket-splitting approach are problematic on graphs with unbalanced node degrees: the number of phases is superlinear in the maximum node degree. For example, on graphs modeling the World-Wide Web or telephone calls this means $\Omega(n^{0.48})$ parallel time. We propose a new node selection rule that removes this strong dependence. Combined with an optimized split-free bucket structure we obtain the first $o(n^{1/4})$ average-case time linear work parallel algorithm for SSSP on the graph classes mentioned above.

Path-based Algorithms for Multicommodity Flow Problems in Large Networks

by ROLF MÖHRING

Many practical problems involve flow problems where the standard formulation via air flows is not adequate. Reasons are constraints on the flow carrying paths, global packing constraints, or path-dependent cost functions.

We illustrate this in the context of route guidance models that aim at minimizing total travel time in urban traffic networks with congestion. In this case, an optimal solution may assign very long routes to individual users, which is unacceptable in practice. This is circumvented by enforcing a length restriction on the routes.

The resulting optimization problem is a path-based multicommodity flow problem with convex separable objective, which we solve with a combination of gradient method, linear optimization with column generation and constrained shortest paths as pricing subroutine.

We discuss also extensions to dynamic flows and ideas for speeding up path-based flow algorithms in traffic networks.

Large Graphs with LEDA – Data Structures for Static Graphs

by STEFAN NÄHER

LEDA provides a very general fully dynamic and powerful graph data type. It allows to implement graph algorithms very similar to the typical pseudo-code representation and is used for all kinds of graphs, e. g., in basic graph algorithms, network optimization, graph drawing, or computational geometry. However, this flexibility has its prize. The underlying data structure is based on doubly linked lists and the graph objects (nodes, edges, and faces) are quite big. Therefore, the space requirement is high and memory cannot be used in a cache-optimal (locality preserving) way. This talk presents new (static) graph data structures which are specialized to particular applications, such as network flow algorithms or triangulations in computational geometry. They have been designed to use as little space as possible and to support graph traversals in a much more efficient way by storing adjacent objects closer to each other in memory, and thus avoiding cache misses. The corresponding data types have been realized as C++ class templates that can be parameterized by so-called traits classes describing the particular properties of the graph (e. g., whether incoming edges can be traversed as well). For typical examples (shortest path, maxflow, or Delaunay triangulations) we show that these special graph classes can be used in the same

elegant way as the standard LEDA graph while providing a much better efficiency. For instance, the running time of the maxflow code contained in the LEDA library can be reduced by a factor of two when using one of these special static graph data types.

Graph Algorithms Between Theory and Practice

by PETER SANDERS

The development of asymptotically efficient algorithms for fundamental computational problems is a highly developed discipline. However, there is often a gap to practically used methods that are much simpler, work well for real inputs but are theoretically inferior. The larger the inputs, the more vexing can the apparent contradiction become. The talk gives examples from graph theory and proposes some approaches for bridging the gaps. Ideally, the result is a careful implementation of new algorithms that are both theoretically satisfying and a practical improvement. We start with some results on shortest paths and maximum flows. Then we turn to some ideas on minimum spanning trees and present a concrete application from radar image processing.

Gas Pipeline Optimization under Uncertainty

by RÜDIGER SCHULTZ

Gas pipeline optimization involves the consideration of networks whose complexity is caused not so much by their topology but rather by time dependence and by partial differential equations describing the gas transport in a pipe. Based on the real system of the German gas supplier Ruhrgas AG, we are heading for optimization models that, on the one hand, are computable and, on the other hand, are sufficiently close to reality. With a spatial discretization of the network and a time discretization of the optimization horizon, we propose a mixed-integer linear program (MILP) for minimizing the total costs of gas compression. The model involves linearizations of the pressure drop in the pipes as well as of the gas consumption in the compressors. Preliminary computational experience with a simplified test network is reported.

In the course of the liberalization of energy markets, gas pipeline optimization, in an increasing manner, has to comply with uncertainty of data. Motivated by uncertainty of demand profiles over time, we extend the deterministic MILP for minimizing compression costs into a two-stage stochastic mixed-integer linear program. The latter can be solved by scenario decomposition, a method that has proven successful in solving stochastic integer programs arising in other fields of application, such as unit commitment in power production. Scenario decomposition essentially rests on Lagrangian relaxation of non-anticipativity constraints, on a dual non-smooth concave maximization, and on heuristics to derive promising primal solutions from the results of the dual optimization.

Semi-External Depth First Search

by JOP F. SIBEYN (joint work with Ulrich Meyer)

Depth-first search is a basic and crucial operation on graphs. On undirected graphs it can be used for computing the biconnected components. On directed graphs DFS is the key routine

to computing strongly connected components. For acyclic graphs it can be used to compute a topological sorting. For graphs with n nodes and m edges, sequential DFS can be performed in $O(n + m)$ time. This algorithm accesses the n adjacency lists of the nodes in an *a priori* unpredictable order, which implies that it exploits random memory access in an essential way. In an external-memory context, where the data do not fit in the RAM memory, it performs extremely bad. Surprisingly, there are no substantially better external-memory algorithms, which means that in practice the above mentioned problems cannot be solved for graphs such as the web or the call graphs of telephone companies.

Here we assume that the internal memory can hold $c \cdot n$ data, for some small constant c , but that m is too large. This is called semi-external computing. Even in this context, theory does not provide any practical solution for the DFS problem. Though nothing is proven, it appears that the problem has an intrinsic hardness similar to NP-hardness in the sequential and P-completeness in the parallel domain. It is a common and accepted practice to try to tackle such problems by heuristics.

We have developed such a heuristic. It has the following strong features: minimal usage of internal memory (it can be brought down to $3 \cdot n$ ints), acceptable running speed (in practice the program takes between 4 and 20 times longer than the sequential algorithm if the data would fit in the main memory). Together this implies that it can indeed be applied for very large graphs which are not extremely sparse (for very sparse graphs, semi-external algorithms make no sense). The basic idea of the algorithm is to maintain a tentative forest, which is gradually developing into a DFS forest. This idea is simple and natural, does not easily lead to a good algorithm though. Algorithms of this kind may continue to consider all edges without making (much) progress. In our algorithm these problems are overcome by:

- only replacing an edge in the tentative forest if necessary;
- rearranging the branches of the tentative forest, so that it faster grows deep (as a consequence, from among the many correct DFS forests, our algorithm finds a relatively deep one);
- after considering all edges once, determining as many nodes as possible that have reached their final position in the forest and reducing the set of edges accordingly.

A single simple reduction strategy gives effective reduction for graphs whose DFS forest mainly consists of one very long path. We spent considerable effort on developing three other, less simple strategies to even obtain effective problem reduction for graphs with a less fortunate structure.

The algorithm has been turned into a C program, which in addition to the computation does testing, checking and analysis. It also has build in routines to generate directed and undirected graphs of several classes, to randomize the edges and to filter out double edges and self loops. The central subroutines only take a few hundred lines. All large data structures are maintained as files. Virtual memory is not used. All IO is done in a buffered way, with buffers that are so large

that seek time is negligible. In this way the time for IO has been reduced to a non-dominating fraction (about 20%) of the overall running time.

We are aiming at solving DFS problems efficiently for all graphs, but we particularly also wanted to perform optimal for random graphs. Therefore, the algorithm starts by performing an initial round in which $m_0 = \min\{m, O(n \cdot \log n)\}$ edges are processed (the hidden constant is so small, that this is practical). This is just enough to find a path of length $n - o(n)$. By rearranging the tree, we can guarantee that the nodes on this path are final, which allows to filter out all their ingoing edges and all outgoing edges to other nodes on the path. By tuning the constant, it can be guaranteed that the remaining problem has size $O(n)$. Thus, the whole problem can be solved by processing m_0 edges, then scanning through all edges once, and then solving a problem of size $O(n)$. For $m = \omega(n)$, this implies that DFS can be performed in just slightly more than a scan.

Abstract Combinatorial Programming and Efficient Property Testers

by CHRISTIAN SOHLER (joint work with Artur Czumaj)

The goal of property testing is to distinguish between the case whether a given object has a certain property or is ‘far away’ from any object having this property. In the first part of the talk we prove that k -colorability of graphs can be tested in time independent of the size of the graph (this was first proven by Goldreich, Goldwasser, and Ron). We present a testing algorithm that examines only $\tilde{O}(k^4/\epsilon^4)$ entries in the adjacency matrix of the input graph, where ϵ is a distance parameter independent of the size of the graph.

In the second part of the talk we present a general proof technique that can be used to show that certain properties that are closed under restrictions (if an object has a property then any ‘subobject’ also has the property; e.g., if the object is a graph then colorability is closed under restrictions while connectivity is not). We introduce *abstract combinatorial programming* which can be roughly described as linear programming where you forget about the geometry and allow that a set of constraints defines multiple bases. Then we show that a property that is closed under restrictions can be tested, if there is a gap and feasibility preserving reduction that maps any object to an abstract combinatorial program of small dimension. We illustrate our approach with three examples: testing low degree uni-variate polynomials, radius clustering, and graph coloring.

Fast Approximation of Multi-cast Routing with Congestion Minimization

by ANAND SRIVASTAV (joint work with Andreas Baltz)

A multicast communication network can be modeled as an undirected graph $G = (V, E)$, $|V| = n$, $|E| = m$, where each node represents a computer that is able to receive, copy and send packets of data. A packet is multicast, if there is a simultaneous demand from several nodes of receiving a copy from the packet’s source. We can specify a multicast request by a subset S of nodes called “terminals”. Meeting the request means to establish a connection represented by an S -tree, i.e. a subtree of G containing S . Given G and a set of multicast requests $S_1, \dots, S_k \subseteq V$ it

is natural to ask about S_i -trees minimizing the maximum edge congestion, i. e., the maximum number of times a specific edge is used in the S_i -trees. Solving this minimization problem is NP-hard, since it contains as a special case the well-known NP-hard standard routing problem of finding integral paths with minimum congestion. We present a deterministic algorithm for approximating the minimum multicast congestion problem within $O(\text{OPT} + \log n)$ in $\tilde{O}(k^3 nm)$ time.

Computing Network Reliability in Graphs of Restricted Path Width

by PETER TITTMANN (joint work with André Pönitz)

The *reliability polynomial* $R(G, p)$ of an undirected graph $G = (V, E)$ is the probability that G is connected assuming all edges of G fail independently with probability $1 - p$.

The problem of computing the reliability polynomial of general graphs was proven to be NP-hard by M. O. Ball in 1980.

We propose an algorithm to perform the computation in polynomial time for graphs of restricted path width given a *nice path composition* of the graph.

The algorithm consists of $|V|$ *node activation*, $|E|$ *edge transformation* and $|V|$ *node deactivation* stages, the order of which are determined by the given path decomposition. Each stage transforms a set of up to $\text{Bell}(d)$ *states* into another set of such states. Thereby d depends on the stage but is bounded from above by the path width of the graph. A state is an ordered pair of a partition of certain (stage dependent) set of nodes of cardinality of at most d and a univariate polynomial of at most $|V|$ monoms.

Our algorithm for the computation of the reliability polynomial is member of a wider parametrized class of algorithms that include algorithms for the computation of other graph invariants like the chromatic number, the chromatic polynomial, the matching polynomial, the number of Hamiltonian circles, a shortest Hamiltonian circle and many others.

Practical Algorithms for the Steiner Problem

by SIAVASH VAHDATI DANESHMAND (joint work with Tobias Polzin)

State-of-the-art exact algorithms for \mathcal{NP} -hard problems like the Steiner problem consist of many different components, playing together as an “orchestra”. We begin with giving an overview over our past and current work on such components, namely relaxations and lower bounds, reduction methods and upper bounds. Then we look more closely at two samples. The first one combines the strengths of alternative- and bound-based techniques while extending the scope of reduction methods from single vertices or edges to patterns like trees. The second uses a partitioning scheme through vertex separators as the basis for new reduction methods. Finally, we show how the presented methods interact nicely in the process of exact solution of very large instances.

Multiple-Choice Algorithms

by BERTHOLD VÖCKING

Multiple-choice allocation algorithms have been studied intensively over the last decade. These algorithms have several applications in the areas of load balancing, routing, resource allocation and hashing. The underlying idea is simple and can be explained best in the balls-and-bins model: Instead of assigning balls (jobs, requests, or keys) simply at random to bins (machines, servers, or positions in a hash table), choose first a small set of bins at random, inspect these bins, and place the ball into one of the bins containing the smallest number of balls among them.

The simple idea of first selecting a small set of alternatives at random and then making the final choice after careful inspection of these alternatives leads to great improvements against algorithms that place their decisions simply at random. We illustrate the power of this principle in terms of simple balls-and-bins processes. For example, we present a recent result showing how the usage of asymmetry can improve the load balancing even more.

In addition, we give some examples showing advanced applications of the multiple-choice principle that cannot be modeled directly in terms of the standard balls-and-bins processes.

Analysis and Visualization of Social Networks

by DOROTHEA WAGNER (joint work with Ulrik Brandes)

In social network analysis, relational informations about actors in politics, management or society in general are studied together with additional attributes of the actors or the relations among the actors are studied. A quite new issue is the use of visualizations as basis for the analysis. In this talk, we give an overview of the main aspects of social network visualization.

In the second part, we concentrate on algorithmic questions in this context. We give two examples, where new algorithmic problems arise and report on first attempts to solve these. First, we study network visualizations that communicate centrality of actors. A second study deals with actor status in directed networks.

Parallel Bridging Models and Their Impact on Algorithm Design

by ROLF WANKA (joint work with Friedhelm Meyer auf der Heide)

The aim of this talk is to demonstrate the impact of features of parallel computation models on the design of efficient parallel algorithms. For this purpose, we start with considering Valiant's BSP model and design an optimal multisearch algorithm. For a realistic extension of this model which takes the critical blocksize into account, namely the BSP* model due to Bäumker, Dittrich, and Meyer auf der Heide, this algorithm is far from optimal. We show how the critical blocksize can be taken into account by presenting a modified multisearch algorithm which is optimal in the BSP* model. Similarly, we consider the D-BSP model due to de la Torre and Kruskal which extends BSP by introducing a way to measure locality of communication. Its influence on algorithm design is demonstrated by considering the broadcast problem. Finally, we explain how our Paderborn University BSP (PUB) Library incorporates such BSP extensions.

On OBDD-based Graph Algorithms

by INGO WEGENER

Small graphs can be stored in the main memory, e.g., by adjacency matrices or adjacency lists. For large graphs one has to use the external memory in order to store this information. Giant graphs (like transition graphs in hardware, traffic simulation graphs, or WWW-graphs) are described implicitly. It would take too much time to enumerate all edges. Hence, algorithms on these graphs cannot work with an explicit graph description. Most types of implicit and compact descriptions do not support operations on the graphs. OBDDs seem to be a compact representation of the adjacency matrix (at least for many well-structured graphs) which supports many important operations. However, typical graph algorithms are inefficient if the graphs are given by OBDDs. Therefore, many problems have to be solved to design efficient OBDD-based graph algorithms. As an example how this approach may work, an OBDD-based network flow algorithm is presented.

On the Single-Source Shortest Path Problem in Embedded Graphs

by THOMAS WILLHALM

There are several known speed-up techniques for Dijkstra's Algorithm. We present two of them that use geometric information: Goal-Directed Search and Angle-Restriction. These speed-up techniques are applicable only for embedded graphs. They both reduce the search space in such a way that single-source single-target queries are answered faster in practice, while the result returned is still guaranteed to be correct.

The second part of the talk presents a first attempt to apply these speed-up techniques for non-embedded graphs by generating appropriate embeddings. As a concrete example, the speed-up for different embeddings of the German railway network are examined. Some hints and obstacles of good embeddings are discussed.

E-Mail Addresses

Rudolf Ahlswede	ahlswede@mathematik.uni-bielefeld.de
Susanne Albers	albers@informatik.uni-freiburg.de
Andreas Brandstädt	ab@informatik.uni-rostock.de
Leslie Ann Goldberg	leslie@dcs.warwick.ac.uk
Sven Grothklags	svengroth@upb.de
Horst Hamacher	hamacher@mathematik.uni-kl.de
Klaus Jansen	kj@informatik.uni-kiel.de
Michael Kaufmann	mk@informatik.uni-tuebingen.de
Georg Kliewer	georg.kliewer@upb.de
Spyros Kontogiannis	spyros@mpi-sb.mpg.de
Matthias Krause	krause@informatik.uni-mannheim.de
Piotr Krysta	krysta@mpi-sb.mpg.de
Ulrich Meyer	umeyer@mpi-sb.mpg.de
Friedhelm Meyer auf der Heide	fmadh@upb.de
Rolf H. Möhring	moehring@math.tu-berlin.de
Stefan Näher	naeher@informatik.uni-trier.de
Michael Paterson	msp@dcs.warwick.ac.uk
Tobias Polzin	polzin@mpi-sb.mpg.de
André Pönitz	poenitz@htwm.de
Peter Sanders	sanders@mpi-sb.mpg.de
Rüdiger Schultz	schultz@math.uni-duisburg.de
Jop F. Sibeyn	jopsi@cs.umu.se
Christian Sohler	csohler@upb.de
Anand Srivastav	asr@numerik.uni-kiel.de
Peter Tittmann	peter@htwm.de
Siavash Vahdati Daneshmand	vahdati@informatik.uni-mannheim.de
Berthold Vöcking	voecking@mpi-sb.mpg.de
Dorothea Wagner	Dorothea.Wagner@uni-konstanz.de
Rolf Wanka	wanka@upb.de
Ingo Wegener	wegener@ls2.informatik.uni-dortmund.de
Thomas Willhalm	Thomas.Willhalm@uni-konstanz.de
Philipp Wölfel	woelfel@ls2.cs.uni-dortmund.de
Uwe Zimmermann	u.zimmermann@tu-bs.de

Addresses

Rudolf Ahlswede
Universität Bielefeld
Fakultät Mathematik
Universitätsstr. 25
PF 10 01 31
D-33501 Bielefeld

Andreas Brandstädt
Universität Rostock
FB Informatik
Albert-Einstein-Str. 21
D-18051 Rostock
☎ +49-381-498-3363
📠 +49-381-498-3366
<http://www.informatik.uni-rostock.de/~ab/ab.html>

Sven Grothklags
Universität Paderborn
FB 17 - Mathematik/Informatik
Fürstenallee 11
D-33102 Paderborn
☎ +49-5251-606-705
📠 +49-5251-606-697
<http://www.upb.de/cs/sven/>

Klaus Jansen
Universität Kiel
Institut für Informatik und Prakt. Mathematik
Olshausenstr. 40
D-24098 Kiel
☎ +49-431-880-7501
📠 +49-431-880-7614
<http://www.informatik.uni-kiel.de/inf/Jansen/>

Susanne Albers
Universität Freiburg
Institut für Informatik
Georges-Köhler-Allee
D-79110 Freiburg
☎ +49-761-203-8040
📠 +49-761-203-8042
<http://www.informatik.uni-freiburg.de/~salbers/>

Leslie Ann Goldberg
University of Warwick
Dept. of Computer Science
CV4 7AL Coventry
United Kingdom
☎ +44-2476-52 33 63
📠 +44-2476-52 57 14
<http://www.dcs.warwick.ac.uk/~leslie/>

Horst Hamacher
Universität Kaiserslautern
FB Mathematik
Postfach 3049
D-67653 Kaiserslautern
📠 +49-631-2 90 82
<http://www.mathematik.uni-kl.de/~wwwwi/WWWWI/HOMEPAGES/hintro.html>

Michael Kaufmann
Universität Tübingen
Wilhelm-Schickard-Institut für Informatik
Sand 13
D-72076 Tübingen
☎ +49-7071-2977-404
📠 +49-7071-5061
<http://www-pr.informatik.uni-tuebingen.de/>

Georg Kliewer
Universität Paderborn
FB 17 - Mathematik/Informatik
Fürstenallee 11
D-33102 Paderborn
☎ +49-5251-606704
☎ +49-5251-606-697
<http://www.upb.de/cs/geokl/>

Matthias Krause
Universität Mannheim
Institut für Informatik
Theoretische Informatik D7,27
D-68131 Mannheim
☎ +49-621-181 2670
☎ +49-621-181 3456
<http://th.informatik.uni-mannheim.de/>

Ulrich Meyer
MPI für Informatik
Stuhlsatzenhausweg 85
D-66123 Saarbrücken
☎ +49-681-9325-125
☎ +49-681-9325-199
<http://www.mpi-sb.mpg.de/~umeyer/>

Rolf H. Möhring
TU Berlin
Institut für Mathematik
Skr. MA 6-1
Straße des 17. Juni 136
D-10623 Berlin
☎ +49-30-314-24594 / -25728
☎ +49-30-314-2 51 91
<http://www.math.tu-berlin.de/~moehring/>

Spyros Kontogiannis
MPI für Informatik
AG 1
Stuhlsatzenhausweg 85
D-66123 Saarbrücken
☎ +49-681-9325-117
☎ +49-681-9325-199
<http://www.mpi-sb.mpg.de/~spyros/>

Piotr Krysta
MPI für Informatik
Stuhlsatzenhausweg 85
D-66123 Saarbrücken
☎ +49-681-9325-502
☎ +49-681-9325-199

Friedhelm Meyer auf der Heide
Universität Paderborn
Heinz Nixdorf Institut
FB Mathematik/Informatik
Fürstenallee 11
D-33102 Paderborn
☎ +49-5251-60-6480
☎ +49-5251-60-6482
<http://www.uni-paderborn.de/cs/fmadh.html>

Stefan Näher
Universität Trier
FB IV - Informatik
Universitätsring 15
D-54286 Trier
☎ +49 651 201-3275
☎ +49 651 201-3856
<http://www.informatik.uni-trier.de/~naeher/>

Michael Paterson
University of Warwick
Dept. of Computer Science
CV4 7AL Coventry
United Kingdom
☎ +44-2476-523194
<http://www.dcs.warwick.ac.uk>

Andr Pönitz
Hochschule Mittweida
FB Mathematik/Physik/Informatik
Technikumplatz 17
Postfach 91
D-09642 Mittweida

Rüdiger Schultz
Universität-GH Duisburg
Institut für Mathematik
Lotharstr. 65
D-47048 Duisburg
☎ +49-203-379-1898
☎ +49-203-379-3139

Christian Sohler
Universität Paderborn
FB 17 – Mathematik/Informatik
Fürstenallee 11
D-33102 Paderborn
☎ +49-5251-60 64 82

Peter Tittmann
Hochschule Mittweida
FB Mathematik/Physik/Informatik
Technikumplatz 17
D-09642 Mittweida
☎ +49-3727-58 10 31
☎ +49-3727-58 13 15
<http://www.htwm.de/peter/>

Tobias Polzin
MPI für Informatik
Stuhlsatzenhausweg 85
D-66123 Saarbrücken
☎ +49-681-9325-122
☎ +49-681-9325-999
<http://www.mpi-sb.mpg.de/~polzin>

Peter Sanders
MPI für Informatik
Stuhlsatzenhausweg 85
D-66123 Saarbrücken
☎ +49 681 9325 115
☎ +49 681 9325 199
<http://www.mpi-sb.mpg.de/~sanders/>

Jop F. Sibeyn
Umea University
Dept. of Computing Science
901 87 Umea, Sweden
☎ +46-90-786-6126
<http://www.cs.umu.se/~jopsi/>

Anand Srivastav
Universität Kiel
Mathematisches Seminar
Ludwig-Meyn-Str. 4
D-24098 Kiel
☎ +49-431-880-1725
<http://www.numerik.uni-kiel.de/>

Siavash Vahdati Daneshmand
Universität Mannheim
Fakultät Mathematik und Informatik
Theoretische Informatik D7,27
D-68131 Mannheim
☎ +49-621-181-3466
☎ +49-621-181-3456

Berthold Vöcking
MPI für Informatik
AG1
Stuhlsatzenhausweg 85
D-66123 Saarbrücken
☎ +49 +681 +9325 199
<http://www.mpi-sb.mpg.de/~voecking/>

Dorothea Wagner
Universität Konstanz
FB Informatik & Informationswissenschaft
Universitätsstr. 10
D 188
D-78457 Konstanz
☎ +49-7531-88-2893
☎ +49-7531-88-3577
<http://www.fmi.uni-konstanz.de/~wagner/>

Rolf Wanka
Universität Paderborn
FB 17 – Mathematik / Informatik
D-33095 Paderborn
☎ +49-5251-60-6434
☎ +49-5251-60-6482
<http://www.upb.de/cs/wanka.html>

Ingo Wegener
Universität Dortmund
FB Informatik II
D-44221 Dortmund
☎ +49-231-755-2776
☎ +49-231-755-2047
<http://ls2-www.cs.uni-dortmund.de/~wegener/>

Thomas Willhalm
Universität Konstanz
FB Informatik & Informationswissenschaft
Universitätsstr. 10
D 188
D-78457 Konstanz
☎ +49-7531-88 40 15
☎ +49-7531-88 35 77
<http://www.fmi.uni-konstanz.de/~willhalm>

Philipp Wölfel
Universität Dortmund
FB Informatik II
D-44221 Dortmund
☎ +49-231-755-2120
☎ +49-231-755-2047
<http://www.cs.uni-dortmund.de/~woelfel>

Uwe Zimmermann
TU Braunschweig
Institut für Math. Optimierung
Geb. 4201
Pockelstr. 14
Postfach 3329
D-38106 Braunschweig
☎ +49-531-391-7559
<http://mo.math.nat.tu-bs.de/~uz/>