Matematisches Forschungsinstitut Oberwolfach

Tagungsbericht 32/2000

Efficient Algorithms

August 6, 2000 - August 12, 2000

This workshop was organized by Susanne Albers (Dortmund), Torben Hagerup (Frankfurt), Giuseppe F. Italiano (Rome) and Kurt Mehlhorn (Saarbrücken). There were 36 participants coming from a total of 9 different countries, who contributed talks in many areas related to the design and analysis of algorithms and data structures.

Each day offered plenty of excellent and stimulating talks as well as many opportunities for fruitful discussions and research exchanges. The talks revealed to be interesting also to the outside community at large: on the very first day, even a small snake was seen wandering outside of the lecture hall! Some discussions were saved also for the traditional hike, which took place in a great weather.

The pleasant atmosphere provided by the staff of the Matematisches Forschungsinstitut Oberwolfach contributed greatly to the success and productivity of the meeting.

Abstracts:

Connection caching in the Web

Susanne Albers

In this talk we study the problem of maintaining persistent TCP connections in a network where each network node can maintain a limited number of open connections. We consider a general setting where connections may incur varying establishment costs. We present online algorithms that achieve an optimal competitive ratio and develop tradeoffs between competitiveness and the amount of extra communication used between network nodes.

Optimal Projective List Update Algorithms

Christoph Ambühl

The list update problem is a classical online problem, with an optimal competitive ratio that is still open, somewhere between 1.5 and 1.6. An algorithm with competitive ratio 1.6, the smallest known to date, is COMB, a randomized combination of BIT and TIMESTAMP. This and many other known algorithms, like MTF, are projective in the sense that they can be defined by only looking at any pair of list items at a time. Projectivity simplifies both the description of the algorithm and its analysis, and so far seems to be the only way to define a good online algorithm for lists of arbitrary length. In this talk we characterize all projective list update algorithms and show their competitive ratio is never smaller than 1.6. Therefore, COMB is a best possible projective algorithm, and any better algorithm, if it exists, would need a non-projective approach.

On External-Memory MST, SSSP and Multi-way Planar Graph Separation

Lars Arge

Recently external memory graph algorithms have received considerable attention because massive graphs arise naturally in many applications involving massive data sets. Even though a large number of I/O-efficient graph algorithms have been developed, a number of fundamental problems still remain open. In this talk we discuss an improved algorithm for the problem of computing a minimum spanning tree of a general graph, as well as new algorithms for the single source shortest paths and the multi-way graph separation problems on planar graphs.

Strongly Polynomial Algorithms for the Unsplittable Flow Problem

Yossi Azar

We provide the first strongly polynomial algorithms with the best approximation ratio for all three variants of the unsplittable flow problem (UFP). In this problem we are given a (possibly directed) capacitated graph with n vertices and m edges, and a set of terminal pairs each with its own demand and profit. The objective is to connect a subset of the terminal pairs each by a single flow path as to maximize the total profit of the satisfied terminal pairs subject to the capacity constraints. Classical UFP, in which demands must be lower than edge capacities, is known to have an $O(\sqrt{m})$ approximation algorithm. We provide the same result with a strongly polynomial combinatorial algorithm. The extended UFP case is when some demands might be higher than edge capacities. For that case we both improve the current best approximation ratio and use strongly polynomial algorithms. We also use a lower bound to show that the extended case is provably harder than the classical case. The last variant is the bounded UFP where demands are at most 1/K of the minimum edge capacity. Using strongly polynomial algorithms here as well, we improve the currently best known algorithms. Specifically, for K=2 our results are better than the lower bound for classical UFP thereby separating the two problems.

Joint work with Oded Regev.

Mastering the * operator

Hannah Bast

How often must one iterate a given function f to get from x to y, or, formally, what is $f^*(x,y) = \min\{i : f^{(i)}(x) \leq y\}$? The derivation of closed formulas for this quantity is a frequent subtask in the analysis of algorithms, and is usually done in an ad hoc manner. We present a master method that provides a surprising approximation of $f^*(x,y)$ in the form of a simple integral.

The Strengths and Weaknesses of the Discrepancy Method

Bernard Chazelle

The discrepancy method, which is the linkage between discrepancy theory and peudorandomness, has been the most powerful tool for understanding randomization as a computational resource. It has also been used for proving lower bounds in circuit complexity and communication complexity. I will review some of the milestones in the story of the discrepancy method and I will discuss what it can do and what it (probably) cannot do.

Broadcasting and gossiping in radio networks

Marek Chrobak

We consider two problems of disseminating information in radio networks: broadcasting and gossiping. We present the following results:

1. A deterministic algorithm for broadcasting with running time $O(n \cdot log^2 n)$. The best lower bound for broadcasting is $\Omega(n \cdot \log n)$, so our algorithm is at most one logarithmic factor from optimum.
2. A deterministic algorithm for gossiping with running time $O(n^{3/2} \cdot \log^2 n)$.
3. A randomized algorithm for gossiping with running time $O(n \cdot \log^4 n)$.

Approximability of the Minimum-Cost k-Connected Spanning Subgraph Problem

Artur Czumaj

In the first part of the presentation I will discuss polynomial-time approximation schemes (PTAS) for several basic minimum-cost multi-connectivity problems in geometrical graphs. I describe a randomized approximation scheme for finding a biconnected graph spanning a set of points in a multi-dimensional Euclidean space and having the expected total cost within (1+epsilon) of the optimum. For any constant dimension and epsilon, the scheme runs in time O(n log(n)). It can be turned into Las Vegas one without affecting its asymptotic time complexity, and also efficiently derandomized. Further, I will briefly explain the extension of this approach to a PTAS for finding a minimum-cost k-edge connected graph/multigraph spanning a set of points in a multi-dimensional Euclidean space and to a PTAS for the problem of Euclidean minimum-cost Steiner biconnectivity.

In the second (shorter) part of the talk I will discuss hardness of approximations for the minimum-cost k-vertex- and k-edge-connected spanning subgraph problems. The only inapproximability result known so far states that the k-edge-connectivity problem in unweighted graphs does not have a PTAS unless P=NP, even for k=2. I shall present a simpler proof of this result that holds even for graphs of bounded degree, and provide the first proof that finding a PTAS for the k-vertex-connectivity problem in unweighted graphs is NP-hard even for k=2 and for graphs of bounded degree.

Joint work Andrzej Lingas, Lund University, Sweden.

Dynamic Matrices for Fully Dynamic Transitive Closure

Camil Demetrescu

Fully Dynamic Transitive Closure is the problem of answering reachability queries on a dynamic directed graph subject to both insertions and deletions of edges. This problem has been studied for over two decades. We present a new deterministic algorithm which achieves $O(n^2)$ time per update and $O(1)$ per query in general directed graphs and uses $O(n^2)$ space. This algorithm is based on a novel technique which consists of casting Fully Dynamic Transitive Closure into the problem of reevaluating polynomials over Boolean matrices.

Joint work with Giuseppe F. Italiano.

Making Data Structures Confluently Persistent

Amos Fiat

We address a longstanding open problem of Driscoll, Sarnak, Sleator, and Tarjan. We present a general transformation that takes any data structure and transforms it to a confluently persistent data structure. We model this general problem using the concepts of a version DAG and an instantiation of a version DAG. We introduce the concept of the effective depth of a vertex in the version DAG and use it to derive information theoretic lower bounds on the space expansion of any such transformation for this DAG. We then give a confluently persistent data structure, such that for any version DAG, the time slowdown and space expansion match the information theoretic lower bounds to within a factor of $O(\log(|V|)$.

Joint work with Haim Kaplan.

Routing flow through a strongly connected graph

Thomas Erlebach

It is shown that, for every strongly connected network in which every edge has capacity at least Δ, linear time suffices to send flow from source vertices, each with a given supply, to sink vertices, each with a given demand, provided that the total supply equals the total demand and is bounded by Δ. This problem arises in a maximum-flow algorithm of Goldberg and Rao, the binary blocking flow algorithm.

Joint work with Torben Hagerup.

Opportunistic data structures with applications

Paolo Ferragina

There is an upsurging interest in designing succinct data structures for basic searching problems. The motivation has to be found in the exponential increase of electronic data nowadays available which is even surpassing the significant increase in memory and disk storage capacities of current computers. Space reduction is an attractive issue because it is also intimately related to performance improvements. In designing these implicit data structures the goal is to reduce as much as possible the auxiliary information kept together with the input data without introducing a significant slowdown in the final query performance. Yet input data are represented in their entirety thus taking no advantage of possible repetitiveness into them. The importance of those issues is well known to programmers who typically use various tricks to squeeze data as much as possible and still achieve good query performance. Their approaches, thought, boil down to heuristics whose effectiveness is witnessed only by experimentation.

In this paper, we address the issue of compressing and indexing data by studying it in a theoretical framework. We devise a novel data structure for indexing and searching whose space occupancy is a function of the entropy of the underlying data set. We call the data structure opportunistic since its space occupancy is decreased when the input is compressible at no significant slowdown in the query performance and without any assumption on a particular fixed distribution. It has been a belief that some space overhead should be paid to use full-text indices (i.e., suffix trees or suffix arrays) with respect to the word-based indices (i.e., inverted lists). The results in this paper show that no space overhead is needed at all, and as an application we improve space and query time complexity of the well-known Glimpse tool.

Joint work with Giovanni Manzini.

A bound on the capacity of backoff and acknowledgement-based protocols

Leslie Ann Goldberg

We study contention-resolution protocols for multiple-access channels (such as the Ethernet channel). We show that every backoff protocol is unstable if the arrival rate, lambda, is at least 0.42. Thus, we show that backoff protocols have (provably) smaller capacity than "full-sensing" protocols. Finally, we show that the corresponding results, with the larger arrival bound of 0.531, also hold for every acknowledgement-based protocol. (The talk will not presume that the listener has any prior knowledge of contention-resolution protocols.)

Joint work with Mark Jerrum, Sampath Kannan and Mike Paterson.

Minimal perfect hashing in nearly linear space

Torben Hagerup

We consider the following problem: Given a subset S of size n of a universe $\{0,\ldots,u-1\}$, construct a minimal perfect hash function for S, i.e., a bijection h from S to $\{0,\ldots,n-1\}$. The parameters of interest are the space needed to store h, its evaluation time, and the time required to compute h from S. The number of bits needed for the representation of h, ignoring the other parameters, has been thoroughly studied and is known to be $\log\log u + n\log e \pm O(\log n)$, where "log" denotes the binary logarithm. A construction by Schmidt and Siegel uses $O(\log\log u + n)$ bits and offers constant evaluation time, but the time to find h is not discussed. Building on ideas dating back to 1984, we present a simple randomized scheme that uses $\log\log u + n\log e + o(\log\log u + n)$ bits and has constant evaluation time and $O(n + \log\log u)$ expected construction time.

Joint work with Torsten Tholey, Universität Frankfurt.

A subquadratic algorithm for Fully Dynamic Transitive Closure

Giuseppe F. Italiano

The Fully Dynamic Transitive Closure problem consists of performing an intermixed sequence of the following operations on a directed graph:

`insert(x,y)`:	insert and edge from x to y
`delete(x,y)`:	delete edge from x to y
`query(x,y)`:	is y reachable from x?

We present a randomized Monte-Carlo algorithm which works on directed acyclic graphs and supports each query in $O(n^{0.575})$ time and each update in $O(n^{1.575})$ time.

Joint work with Camil Demetrescu.

Preemptive Parallel Task Scheduling in $O(n) + \mathbf{poly}(m)$ Time

Klaus Jansen

We study the problem of scheduling a set of n independent parallel tasks on m processors, where in addition to the processing time there is a size associated with each task indicating that the task can be processed on any subset of processors of the given size. Based on a linear programming formulation, we propose an algorithm for computing a preemptive schedule with minimum makespan, and show that the running time of the algorithm depends polynomially on m and only linearly on n. Thus for any fixed m, an optimal preemptive schedule can be computed in $O(n)$ time. We also present extensions of this approach to other (more general) scheduling problems with malleable tasks, release times, due dates and maximum lateness minimization.

Joint work with Lorant Porkolab (Imperial College London).

Space-efficient dynamic shortest path algorithms, by counting

Valerie King

I show how counting is the key in recent work to make dynamic shortest path and transitive closure algorithms space efficient. I'll also discuss experiments with implementations.

Covering Rectilinear Polygons with axis-parallel Rectangles

Anil Kumar

We give an $O(\sqrt{\log n})$ factor approximation algorithm for covering a rectilinear polygon with holes using axis-parallel rectangles. This is the first polynomial time approximation algorithm for this problem with a $o(\log n)$ approximation factor.

Non-Clairvoyant Scheduling to Minimize the Average Flow Time

Stefano Leonardi

We consider non-clairvoyant scheduling algorithms to minimize the average flow time of a stream of jobs released over time when the processing time of a job is unknown. A non-clairvoyant algorithm is aware of a job at its release time. The processign time of a job is only known at the completion. This is a classical problem in CPU scheduling in time sharing multitasking operating systems. Our goal is to minimize the average flow time of the jobs, i.e the average time spent from the jobs in the system between release and completion. In this paper we study a randomized version of the Multi-Level-Feedback scheduling algorithm, RMLF, at the basis of the scheduling policy of the Unix Operating system.

In this work we give a first competitive result for non-clairvoyant parallel machine scheduling, we show that RMLF achieves an $O(\log n)^2$ competitive ratio on any number of identical parallel machines against the oblivious adversary. We also show that RMLF achieves an $O(\log n)$ competitive factor on a single machine against the oblivious adversary. This compares with the $O(\log n \log\log n)$ competitive ratio of (Kalyanasundaram and Pruhs, Focs

97) against the adaptive on-line adversary and matches the $\Omega(\log n)$ lower bound (Motwani, Phillips and Torng, Soda 94).

Joint work with Luca Becchetti.

TSP-based Curve Reconstruction

Kurt Mehlhorn

Let γ be an unknown closed curve without self-intersections and let S be a finite set of points from γ. The reconstruction problem is to find the cyclic order in which the points in S are arranged on curve. Giesen has shown that for every semi-regular curve (a curve in which the left and right tangents exist everywhere and form angles strictly less than π) there is a positive ϵ such that if S contains at least one point from every curve segment of length ϵ then the traveling salesman tour of S reconstructs the order of S on γ. We show that under the traveling salesman tour can be constructed in polynomial time under the same conditions.

Joint work with Ernst Althaus.

Single-Source Shortest-Paths on Arbitrary Directed Graphs in Linear Average-Case Time

Uli Meyer

The quest for a linear-time single-source shortest-path (SSSP) algorithm on directed graphs with positive edge weights is an ongoing hot research topic. While Thorup recently found an $O(n+m)$ time RAM algorithm for undirected graphs with n nodes, m edges and integer edge weights in $\{0, ..., 2^w - 1\}$ where w denotes the word length, the currently best time bound for directed sparse graphs on a RAM is $O(n + m \cdot \log\log n)$.

We study the average-case complexity of SSSP. We give a simple algorithm for arbitrary directed graphs with random edge weights uniformly distributed in $[0, 1]$ and show that it needs linear time $O(n+m)$ on average.

Deterministic dictionaries

Rasmus Pagh

This talk surveys recent advances on the classic problem of maintaining a set under insertions and deletions, supporting membership queries. In particular, our interest is (deterministic) worst-case bounds. The techniques used for the advances of the last decade are highlighted, and we speculate how futher improvement might be achieved.

Web-like graphs

Prabhakar Raghavan

In this talk we will review a number of measurements on the world-wide web and derivative graphs such as gnutellanet. The measurements suggest that these graphs do not behave like standard random graph models. We discuss algorithms on these graphs for search and data mining. Finally, we give some recent results on stochastic models for generating these graphs.

Analyzing the cache behaviour of non-uniform distribution sorting algorithms

Rajeev Raman

We analyze the average-case cache performance of distribution sorting algorithms in the case when keys are independently but not uniformly distributed. We use this analysis to tune the performance of the integer sorting algorithm MSB radix sort when it is used to sort independent uniform floating-point numbers (floats). It is well known that integer sorting algorithms can correctly sort floats by sorting the bit-strings representing the floats; however, a uniform distribution on floats does not induce a uniform distribution on the bit-strings.

Like other analyses by Sanders (ICALP 99) and Sen and Chatterjee (SODA 00), ours is based upon an assumption which is clearly not true in the usual implementation of distribution sort. Although one can modify distribution sort to satisfy this assumption, this degrades performance in practice. We touch upon some empirical observations regarding this assumption.

Joint work with Naila Rahman.

Highly Efficient PCPs for Optimization Problems

Ronitt Rubinfeld

Consider the following scenario: A client sends a computational request to a "consulting" company on the internet, by specifying an input and a computational problem to be solved. The company then computes the answer and sends it back to the client. An obvious issue that arises, especially in the case that the company does not have a well established reputation, is: why should the client believe the answer to be correct? We consider the case when it is enough for the client to know that the answer is *close* to correct. We show that for a number of optimization functions, a very short dialogue between the client (the verifier) and the company (the prover) can convince the client of the approximate correctness of the answer. In fact, we give protocols for several optimization problems, in which the running time of the verifier is significantly less than the size of the input. In contrast, results such as $IP = PSPACE$, $MIP = NEXPTIME$ and $NP = PCP(logn, 1)$ require a verifier which runs in $\Omega(n)$ time.

We give protocols for with sublinear time verifiers for approximate lower bounds on the solution quality of a subclass of linear programming problems referred to as fractional packing problems. We also give protocols which allow a prover to convince a polylogarithmic time verifier of the existence of a large cut in a graph, a large matching in a graph or a small bin packing. In addition, our results apply to other optimization problems such as max flow and scheduling.

The protocols are very simple to describe, and the talk is essentially self-contained.

Much of this talk describes joint work with Funda Ergun and S. Ravi Kumar. The rest of this talk describes joint work with Tugkan Batu and Patrick White.

Semi-External Depth First Search

Jop F. Sibeyn

A new algorithm is proposed for semi-external unordered DFS. Different from all previous external memory DFS algorithms it performs no random I/Os but just scans. On random graphs with n nodes and $m \geq n$ edges it requires $\emptyset(m/B \cdot \log n)$ I/Os whp (where B denotes the block size). Also on many other graph classes the performance is amazingly good, though the details of the performance remain to be analyzed. We also present a simple yet efficient bit-based internal dictionary and show how it can be used to reduce the number of rounds in the DFS algorithm of Chiang et al. (SODA '95, pp. 139–149) by a factor of $\Omega(\frac{\log n}{\log(n/M)})$ for internal memory size M. This constitutes a significant improvement for dense graphs if n is just a bit too large for the semi-external case.

Joint work with Ulrich Meyer.

Translating a planar object to maximize point containment: A Monte Carlo approximation algorithm

Michiel Smid

Let S be a set of n points in the plane, and let C be a closed set whose boundary is a closed Jordan curve. (For example, C can be a circle, ellipse, polygon, etc.) We consider the problem of computing a translation of C such that the number of points of S contained in the translate is maximum. If the set C has constant description complexity, then this problem can be solved in $O(n^2)$ time. Even for the case when C is a circle, no subquadratic algorithm is known for this problem. We will give an algorithm that computes, with high probability, a translate of C that approximates the optimal translate. If the set C comes from a rather general class of objects (including circles, ellipses, and convex polygons), then the running time of the algorithm is bounded by $O(n \cdot \log^2 n)$.

Quick k-medians for sparse graphs

Mikkel Thorup

We present a randomized $O(m\ polylog\ n)$ time constant factor approximation algorithm for the k-median problem in a weighted graph with n nodes and m edges: find a set S of k nodes minimizing the sum over all other nodes of the shortest path distance to the nearest node in S. The previous "$O(m \cdot \log n)$" algorithm by Jain and Vazirani [FOCS'99] assumed a complete graph with m quadratic in n. They raised it as an open problem to do better for sparse graphs.

Optimal reconstruction of a sequence from its probes

Eli Upfal

An important combinatorial problem, motivated by DNA sequencing in molecular biology, is the reconstruction of a sequence over a small finite alphabet from the collection of its probes (the sequence spectrum), obtained by sliding a fixed sampling pattern over the sequence. Such construction is required for Sequencing-by-Hybridization (SBH), a novel DNA sequencing technique based on an array (SBH chip) of short nucleotide sequences (probes). Once the sequence spectrum is biochemically obtained, a combinatorial method is used to reconstruct the DNA sequence from its spectrum. Since technology limits the number of probes on the SBH chip, a challenging combinatorial question is the design of a smallest set of probes that can sequence an arbitrary DNA string of a given length.

We present a novel probe design, based on the use of universal bases (bases that bind to any nucleotide that drastically improves the performance of the SBH process and asymptotically approaches the information-theoretic bound up to a constant factor. Furthermore, the sequencing algorithm we propose is substantially simpler than the Eulerian path method used in previous solutions of this problem.

GeneSpectrum, a new biotech start-up, is working on implementation of this method.

How to Draw the Minimum Cuts of a Planar Graph

Dorothea Wagner

In this talk we consider the problem of drawing a graph in such a way that its connectivity structure resp. all its minimum cuts can be easily recognized. We show how to utilize the cactus representation of all minimum cuts of a graph to visualize the minimum cuts of a planar graph in a planar drawing. In a first approach the cactus is transformed into a hierarchical clustering of the graph that contains complete information on all the minimum cuts. We present an algorithm for c-planar orthogonal drawings of hierarchically clustered planar graphs with rectangularly shaped cluster boundaries and the minimum number of bends. This approach is then extended to drawings in which the two vertex subsets of every minimum cut are separated by a simple closed curve.

Joint work with Ulrik Brandes and Sabine Cornelsen.

On the analysis of evolutionary algorithms

Ingo Wegener

Evolutionary algorithms are randomized search heuristics for the optimization of functions f. They have applications if time, money or skills are not available to design a problem-specific algorithm or in black box optimization where f is not known and information on f can only be gained by sampling. Many experiments with evolutionary algorithms have been performed but only few theoretical results are known. Such a theory can be built up only by people with experience in the analysis of algorithms. Here a simple evolutionary algorithm for the optimization of functions $f : \{0,1\}^n \to \mathbb{R}$ is investigated. The algorithm works with population size 1 and mutations only. Most results can be generalized to larger populations. It is shown that such algorithms cannot optimize all unimodal functions efficiently. However, linear functions can be optimized in an expected time of $O(n \log n)$. The effect of varying mutation probabilities is studied. Finally, the first example is presented where it is proven that uniform crossover can improve the expected runtime from superpolynomial to polynomial and also the first example where it is proven that 1-point-crossover can improve the expected runtime from exponential to polynomial.

Some of the results are obtained together with Stefan Droste and Thomas Jansen.

A Log-Approximation Algorithm for Covering a Polygon With a Minimum Number of Convex Polygons

Peter Widmayer

The problem Minimum Convex Cover of covering a given polygon with a minimum number of (possibly overlapping) convex polygons is known to be NP-hard, even for polygons without holes. We propose a polynomial-time approximation algorithm for this problem for polygons with or without holes that achieves an approximation ratio of O(log n), where n is the number of vertices in the input polygon. To obtain this result, we first show that an optimum solution of a restricted version of this problem, where the vertices of the convex polygons may only lie on a certain grid, contains at most three times as many convex polygons as the optimum solution of the unrestricted problem. As a second step, we use dynamic programming to obtain a convex polygon which is maximum with respect to the number of "basic triangles" that are not yet covered by another convex polygon. We obtain a solution that is at most a logarithmic factor off the optimum by iteratively applying our dynamic programming algorithm.

Preemptive scheduling with rejection

Gerhard J. Woeginger

We consider the problem of preemptively scheduling a set of n jobs on m (identical, uniformly related, or unrelated) parallel machines. The scheduler may reject a subset of the jobs and thereby incur job-dependent penalties for each rejected job, and he must construct a schedule for the remaining jobs so as to optimize the preemptive makespan on the m machines plus the sum of the penalties of the jobs rejected.

We provide a complete classification of these scheduling problems with respect to complexity and approximability. Our main results are on the variant with an arbitrary number of unrelated machines. This variant is APX-hard, and we design a 1.58-approximation algorithm for it. All other considered variants are weakly NP-hard, and we provide fully polynomial time approximation schemes for them.

Joint work with Han Hoogeveen and Martin Skutella.

E-mail addresses of the participants:

Name	e-mail address
Susanne Albers	albers@ls2.cs.uni-dortmund.de
Helmut Alt	alt@inf.fu-berlin.de
Christoph Ambühl	ambuehl@inf.ethz.ch
Lars Arge	large@cs.duke.edu
Yossi Azar	azar@math.tau.ac.il
Hannah Bast	hannah@mpi-sb.mpg.de
Bernard Chazelle	chazelle@cs.princeton.edu
Marek Chrobak	marek@cs.ucr.edu
Artur Czumaj	czumaj@cis.njit.edu
Camil Demetrescu	demetres@dis.uniroma1.it
Martin Dietzfelbinger	md@theoinf.tu-ilmenau.de
Thomas Erlebach	erlebach@tik.ee.ethz.ch
Paolo Ferragina	ferragin@di.unipi.it
Amos Fiat	fiat@math.tau.ac.il
Leslie Ann Goldberg	leslie@dcs.warwick.ac.uk
Torben Hagerup	hagerup@informatik.uni-frankfurt.de
Giuseppe F. Italiano	italiano@disp.uniroma2.it
Klaus Jansen	kj@informatik.uni-kiel.de
Valerie King	val@csr.csc.uvic.ca
Anil Kumar	kumar@mpi-sb.mpg.de
Stefano Leonardi	leon@dis.uniroma1.it
Kurt Mehlhorn	mehlhorn@mpi-sb.mpg.de
Uli Meyer	umeyer@mpi-sb.mpg.de
Rasmus Pagh	pagh@brics.dk
Prabhakar Raghavan	pragh@verity.com
Rajeev Raman	raman@dcs.kcl.ac.uk
Ronitt Rubinfeld	ronitt@research.nj.nec.com
Jop Sibeyn	jopsi@mpi-sb.mpg.de
Michiel Smid	michiel@isg.cs.uni-magdeburg.de
Mikkel Thorup	mthorup@research.att.com
Eli Upfal	eli@cs.brown.edu
Dorothea Wagner	Dorothea.Wagner@uni-konstanz.de
Frank Wagner	Frank.H.Wagner@bku.db.de
Ingo Wegener	wegener@ls2.cs.uni-dortmund.de
Peter Widmayer	pw@inf.ethz.ch
Gerhard Woeginger	gwoegi@opt.math.tu-graz.ac.at

Tagungsteilnehmer:

Prof. Dr. Susanne Albers
Institut für Informatik II
Universität Dortmund
44221 Dortmund

Prof. Dr. Helmut Alt
Institut für Informatik (WE 3)
Freie Universität Berlin
Takustr. 9
14195 Berlin

Christoph Ambühl
Institut für Informatik
ETH-Zürich
ETH-Zentrum
CH-8092 Zürich

Prof. Dr. Lars Arge
Department of Computer Science
Duke University
Box 90129
Durham , NC 27708-0129
USA

Prof. Dr. Yossi Azar
Department of Computer Science
Tel Aviv University
Ramat Aviv
69978 Tel Aviv
ISRAEL

Dr. Hannah Bast
Max-Planck-Institut für Informatik
Geb. 46
Im Stadtwald
66123 Saarbrücken

Prof. Dr. Bernard Chazelle
Department of Computer Science
Princeton University
Princeton , NJ 08544-2087
USA

Prof. Dr. Marek Chrobak
Department of Computer Sciences
University of California
Riverside , CA 92521
USA

Prof. Dr. Artur Czumaj
Dept. of Computer and Information Sciences
New Jersey Inst. of Technology
University Heights
Newark , NJ 07102-1982
USA

Camil Demetrescu
Dipartimento di Informatica e Sistemistica
Universita di Roma "La Sapienza"
Via Salaria 113
I-00198 Roma

Prof. Dr. Martin Dietzfelbinger
Fak. für Informatik und Automatisierung
Techn. Universität Ilmenau
Postfach 100565
98684 Ilmenau

Thomas Erlebach
Institut TIK
ETH Zürich
Gloriastr. 35
CH-8092 Zürich

Prof. Dr. Paolo Ferragina
Dipartimento di Informatica
Universita di Pisa
Corso Italia 40
I-56125 Pisa

Prof. Dr. Amos Fiat
Department of Computer Science
Tel Aviv University
Ramat Aviv
69978 Tel Aviv
ISRAEL

Prof. Dr. Leslie Ann Goldberg
Department of Computer Science
University of Warwick
GB-Coventry CV4 7AL

Prof. Dr. Torben Hagerup
Max-Planck-Institut für Informatik
Geb. 46
Im Stadtwald
66123 Saarbrücken

Prof. Dr. Giuseppe F. Italiano
Dipartimento di Informatica,
Sistemi e Prod.
Universita di Roma "Tor Vergata"
via die Tor Vergata 110
I-00133 Roma

Prof. Dr. Klaus Jansen
Institut für Informatik und
Praktische Mathematik, Haus I
Universität Kiel
Olshausenstr. 40
24118 Kiel

Prof. Dr. Valerie King
Dept. of Computer Science
University of Victoria
P.O. Box 3055
Victoria BC
CANADA

Prof. Dr. Anil V.S. Kumar
Max-Planck-Institut
für Informatik
Stuhlsatzenhausen 85
66123 Saarbrücken

Dr. Stefano Leonardi
Dipartimento di Informatica e
Sistemistica
Univ. di Roma "La Sapienza"
Via Salaria 113
I-00198 Roma

Prof. Dr. Kurt Mehlhorn
Max-Planck-Institut für Informatik
Geb. 46
Im Stadtwald
66123 Saarbrücken

Uli Meyer
Max-Planck-Institut für Informatik
Geb. 46
Im Stadtwald
66123 Saarbrücken

Rasmus Pagh
BRICS, Dept. of Computer Sciences
University of Aarhus
Ny Munkegade, Building 540
DK-8000 Aarhus C

Dr. Prabhakar Raghavan
Verity, Inc.
894 Ross Drive
Sunnyvale , CA 94089
USA

Prof. Dr. Rajeev Raman
Department of Computer Science
King's College London
Strand
GB-London WC2R 2LS

Dr. Ronitt Rubinfeld
NEC Research Institute
4 Independence Way
Princeton , NJ 08540
USA

Dr. Jop Sibeyn
Max-Planck-Institut für Informatik
Geb. 46
Im Stadtwald
66123 Saarbrücken

Prof. Dr. Michiel Smid
Fakultät für Informatik
Otto-von-Guericke Universität
Universitätsplatz 2
39106 Magdeburg

Dr. Mikkel Thorup
AT&T Labs-Research
PO Box 971
180 Park Avenue
Florham Park , NJ 07932-0971
USA

Prof. Dr. Eli Upfal
Computer Science Department
Brown University
Box 1910
Providence , RI 02912-1910
USA

Prof. Dr. Dorothea Wagner
Praktische Informatik
Fakultät für Mathematik
Universität Konstanz
Postfach 55 60
78434 Konstanz

Dr. Frank Wagner
Konzernentwicklung
Deutsche Bahn AG
Potsdamer Platz 1
10785 Berlin

Prof. Dr. Ingo Wegener
Institut für Informatik II
Universität Dortmund
44221 Dortmund

Prof. Dr. Peter Widmayer
Institut für Informatik
ETH-Zürich
ETH-Zentrum
CH-8092 Zürich

Prof. Dr. Gerhard Woeginger
Institut für Mathematik
Technische Universität Graz
Steyrergasse 30
A-8010 Graz