Yuval Rabani, David Shmoys,
Gerhard Woeginger (editors):

Combinatorial Approximation Algorithms

Dagstuhl-Seminar-Report; 187
18.08.-22.08.97 (9734)

ISSN 0940-1121

Tel.: +49 - 6871 - 9050,
Fax: +49 - 6871 - 905 133

Das Internationale Begegnungs- und Forschungszentrum für Informatik (IBFI) ist eine gemeinnützige GmbH. Sie veranstaltet regelmäßig wissenschaftliche Seminare, welche nach Antrag der Tagungsleiter und Begutachtung durch das wissenschaftliche Direktorium mit persönlich eingeladenen Gästen durchgeführt werden.

Bezugsadresse: Geschäftsstelle Schloss Dagstuhl
Universität des Saarlandes
Postfach 15 11 50
D-66041 Saarbrücken, Germany
Tel.: +49 -681 - 302 4396
Fax: +49 -681 - 302 4397
e-mail: office@dag.uni-sb.de
url: http://www.dag.uni-sb.de

Viele Seminarreporte sind auch in elektronischer Form über die Internetseite verfügbar. In einigen Fällen kann man dort von den Abstracts der Vorträge zu den Vollversionen weiterschalten.

Summary

The Dagstuhl seminar on *Combinatorial Approximation Algorithms* brought together 54 researchers with affiliations in Austria (1), France (1), Germany (11), Hungary (1), Iceland (1), Israel (7), Italy (1), Netherlands (2), Russia (1), Sweden (1), Switzerland (1), United Kingdom (3), and USA (23). In 35 talks the participants presented their latest results on approximation algorithms, covering a wide range of topics. The abstracts of most of these talks can be found in this report. Moreover, there is a list of open problems that were stated in the open problem session.

Special events were a hiking tour on Wednesday afternoon and an open problems session held on Thursday evening. In the open problem session, Mark Jerrum, Dorrit Hochbaum, David Shmoys, Vijay Vazirani, and David Williamson presented lists with their favorite open problems. Following the example of Paul Erdős, David Williamson offered money rewards for solutions to his open problems; Sanjeev Arora and Luca Trevisan managed to solve one of his problems (on the Rectilinear Steiner Arborescence problem) by Friday morning.

Due to the outstanding local organization and the pleasant atmosphere, this seminar was a most enjoyable and memorable event.

Yuval Rabani
David Shmoys
Gerhard Woeginger

List of Participants

Karen Aardal, Universiteit Utrecht
Sanjeev Arora, Princeton University
Yossi Azar, Tel Aviv University
Amotz Bar-Noy, Tel Aviv University
Reuven Bar-Yehuda, Technion
Yair Bartal, ICSI Berkeley
Alexander Barvinok, University of Michigan
Martin E. Dyer, University of Leeds
Guy Even, Universität des Saarlandes
Uri Feige, Weizmann Institute
Sandor Fekete, Universitdt Köln
Naveen Garg, MPI Saarbrücken
Magnus Halldorsson, Science Institute Reykjavik
Dorit Hochbaum, University of California
Mark Jerrum, University of Edinburgh
David S. Johnson, AT&T Bell Labs
Viggo Kann, Royal Institute of Technology Stockholm
David Karger, MIT
Howard Karloff, Georgia Institute of Technology
Marek Karpinski, Universität Bonn
Michael Kaufmann, Universität T—bingen
Claire Kenyon, LIP Lyon
Samir Khuller, University of Maryland at College Park
Philip Klein, Brown University
Elias Koutsoupias, University of California
Stefano Leonardi, Universitá di Roma "La Sapienza"
Rolf H. Moehring, TU Berlin
Seffi Naor, Technion
Cindy Phillips, Sandia National Laboratories
Serge Plotkin, Stanford University
Chris Potts, University of Southampton
Hans-Jürgen Prömel, Humboldt Universität Berlin
Yuval Rabani, Technion

Yuri Rabinovich, Ben-Gurion University
Prabhakar Raghavan, IBM Almaden Research Center
Satish B. Rao, NEC Research Institute Princeton
Monika Rauch-Henzinger, DEC Palo Alto
R. Ravi, Carnegie Mellon University
Baruch Schieber, IBM T.J. Watson Research Center
Andreas Schulz, TU Berlin
Uwe Schwiegelshohn, Universität Dortmund
Sergey V. Sevastianov, Russian Academy of Sciences Novosibirsk
David Shmoys, Cornell University
Miklos Simonovits, Hungarian Academy of Sciences Budapest
Angelika Steger, TU München
Cliff Stein, Dartmouth College
Leen Stougie, University of Amsterdam
Luca Trevisan, Universite de Geneve
Vijay Vazirani, Georgia Institute of Technology
Dorothea Wagner, Universität Konstanz
Jeffery Westbrook, AT&T Labs-Research
David Williamson, IBM T.J. Watson Research Center
Gerhard J. Woeginger, TU Graz
Neal Young, Dartmouth College

Abstracts

Minimizing Service and Operation Costs of Periodic Scheduling
AMOTZ BAR-NOY

We study the problem of scheduling activities of several types under the constraint that at most a fixed number of activities can be scheduled in any single period. Any given activity type is associated with a service cost, and an operating cost that increases linearly with the number of periods since the last service of this type. The problem is to find an optimal schedule that minimizes the long-run average cost per period.

Applications of such a model are the scheduling of maintenance service to machines, multi-item replenishment of stock, and minimizing the mean response time in *broadcast disks*. Broadcast disks gained a lot of attention recently, since they are used to model backbone communications in wireless systems, Teletext systems, and web caching in satellite systems.

The first contribution of our work is the definition of a general model that combines into one several important previous models. We prove that an optimal cyclic schedule for the general problem exists and establish the NP-hardness of the problem. Next we show a lower bound on the cost achieved by an optimal solution. Using this bound we analyze several approximation algorithms. Next, we give a non-trivial 9/8-approximation for a variant of the problem that models the broadcast disks application. The algorithm uses some properties of "Fibonacci sequences". Using this sequence we present a 1.57-approximation algorithm for the general problem. Finally, we describe a simple randomized algorithm and a simple deterministic greedy algorithm for the problem, and prove that both achieve approximation factor of 2. To the best of our knowledge this is the first worst case analysis of a widely used greedy heuristic for this problem.

Joint work with Randeep Bhatia, Seffi Naor, and Baruch Schieber. A full version of the paper can be found in: http://www.cs.umd.edu/~randeep/ and in http://www.eng.tau.ac.il/~amotz/publications.html

Approximating the permanent of a non-negative matrix
ALEXANDER BARVINOK

Let $A = (a_{ij})$ be a non-negative $n \times n$ matrix. Let us consider the following algorithm for approximating the permanent of A (Gaussian version of the Godsil-Gutman estimator). Sample n^2 variables u_{ij} at random from the standard Gaussian distribution in R with the density $\psi(x) = (1/\sqrt{2\pi})\exp\{-x^2/2\}$. Compute an $n \times n$ matrix $B = (b_{ij})$ as follows: $b_{ij} = u_{ij}\sqrt{a_{ij}}$. Compute $\alpha = (\det B)^2$. Output α.

It is easy to show that the expectation of α is the permanent of A, and, therefore, by the Chebyshev inequality the probability that $\alpha \geq C per\ A$ is at most $1/C$ for any $C > 1$, so α is unlikely to overestimate the permanent by a large factor (this remains true if we have sampled u_{ij} from any distribution with expectation 0 and variance 1). Furthermore, it turns out that for some constant $c = e^{-\gamma}/2 \approx 0.28$, where γ is the Euler constant, and any $0 < \epsilon < 1$, the probability that $\alpha < (c\epsilon)^n per\ A$ does not exceed $8/(n \ln^2 \epsilon)$. In other words, α is unlikely to underestimate the permanent of A by more than an exponential in n factor. Using Gaussian distribution is essential, this is no longer true for any discrete distribution.

V.D. Milman suggested that for any $1 > \epsilon > 0$ there is a polynomial time algorithm approximating the permanent within a $(1-\epsilon)^n$ factor.

In the course of the talk and shortly thereafter, Martin Dyer and Mark Jerrum suggested to modify the algorithm by using complex Gaussian variables u_{ij}. Thus u_{ij} have zero expectation, the expectation of $|u_{ij}|^2$ is 1, and $\alpha = |\det B|^2$. This, indeed, improves the constant twice: $c = e^{-\gamma} \approx 0.56$.

Fast Approximate Graph Partitioning Algorithms

GUY EVEN

We study graph partitioning problems on graphs with edge capacities and vertex weights. The problems of b-balanced cuts and k-multiway separators are unified with a new problem called minimum capacity ρ-separators. A ρ-separator is a subset of edges whose removal partitions the vertex set into connected components such that the sum of the vertex weights in each component is at most ρ times the weight of the graph.

We present a new and simple $O(\log n)$-approximation algorithm for minimum capacity ρ-separators yielding an $O(\log n)$-approximation algorithm both for b-balanced cuts and k-multiway separators. In particular, this result improves the previous best known approximation factor for k-multiway separators in undirected graphs by a factor of $O(\log k)$. We enhance these results by presenting a version of the algorithm that obtains an $O(\log \text{OPT})$-approximation factor. The algorithm is based on a technique called spreading metrics that enables us to formulate directly the minimum capacity ρ-separator problem as an integer program.

Joint work with Seffi Naor, Satish Rao, and Baruch Schieber.

Improved bounds for acyclic job shop scheduling

URIEL FEIGE

In acyclic job shop scheduling problems there are n jobs and m machines. Each job is composed of a sequence of operations to be performed on different machines. A legal schedule is one in which within each job, operations are carried out in order, and each machine performs at most one operation in any unit of time. If D denotes the length of the longest job, and C denotes the number of time units requested by all jobs on the most loaded machine, then clearly $lb = \max[C, D]$ is a lower bound on the length of the shortest legal schedule. A celebrated result of Leighton, Maggs and Rao shows that if all operations are of unit length, then there always is a legal schedule of length $O(lb)$, independent of n and m. For the case that operations may have different lengths, Shmoys, Stein and Wein showed that there always is a legal schedule of length $\tilde{O}(lb(\log lb)^2)$, where $(\tilde{O})$ notation is used to suppress $\log\log(lb)$ terms. We improve the upper bound to $\tilde{O}(lb \log lb)$. Similar to the proof of Leighton, Maggs and Rao, our proof makes repeated use of the Lovasz local lemma. Our technique for extending the use of the local lemma to the case where operations are of different lengths may be of independent interest. We also show that our new upper bound is essentially best possible, by proving the existence of instances of acyclic job shop scheduling for which the shortest legal schedule is of length $\tilde{\Omega}(lb \log lb)$. This resolves (negatively) a known open problem of whether the linear upper bound of Leighton, Maggs and Rao applies to arbitrary job shop scheduling instances (without the restriction to acyclicity and unit length operations).

Joint work with Christian Scheideler

Geometric Packing and Dispersion Problems

SÁNDOR FEKETE

We present a number of algorithmic aspects of packing and dispersion of two-dimensional objects into a polygonal domain.

In the first part, we describe difficulties of packing in two dimensions, highlighting differences to one-dimensional problems. In particular, we describe what is and what is not known about "Pallet Loading": How many rectangles of size $a * b$ can you fit into a rectangle of size $A * B$? For this problem, it is not known whether there is an algorithm with running time bounded by a polynomial in $\log a$, $\log b$, $\log A$, $\log B$. Moreover, we do not even know whether the problem belongs to the class NP, i. e., whether there always exists an optimal solution that can be described in space polynomial in the input length. We give some basic concepts and a number of specific questions related to this problem.

In the second part, we consider geometric dispersion problems, where the objective is to place k objects in a polygonal region with n edges, such that the objects are "far away" from each other. The problem of deciding whether k objects of size D can be packed into a given region. It was shown by Hochbaum and Maass that the related problem of maximizing k for a given D in the case of L_1 distance allows a polynomial time approxiamtion scheme. We show that the problem of maximizing D for a given k cannot be approximated in polynomial time within a factor better than $\frac{14}{13}$, unless P=NP. For various versions of the problem, we give polynomial time approximation methods, including a method that guarantees a factor of $\frac{3}{2}$ for the scenario where the objects have to be kept away from the boundary of the region.

Parts of the material presened are joint work with Christoph Baur (Köln).

A framework for half integrality and 2-approximations

DORIT S. HOCHBAUM

We define a class of integer programs for which half integral superoptimal solutions are obtained in polynomial time. For some of these problems it is possible to round the half integral solution to a 2-approximate solution. This extends the class of integer programs with at most two variables per constraint that were analyzed in Hochbaum, Meggido, Naor and Tamir, 1993. Interesting problems for which we can get superoptimal half integral solutions include the sparsest cut problem, and minimum weight edge deletion to obtain a bipartite graph. Problems for which we can get 2-approximations include minimum satisfiability, scheduling with precedence constraints, optimization of boolean functions in two variables (generalized 2SAT) and the feasible cut problem. Certain constraint satisfaction problem are included in this framework as well.

The approximation algorithms here work by solving a minimum cut on a certain network associated with the formulation. These algorithms provide an improvement in running time and range of applicability compared to existing 2-approximations, if any. Furthermore, we conclude that problems in the framework are MAX SNP-hard and at least as hard to approximate as vertex cover.

Problems that are amenable to the analysis provided are easily recognized. The analysis itself is entirely technical and only involves manipulation of the constraints, transforming them to a totally unimodular system while losing no more than a factor of 2 in the integrality.

Approximate counting: a short introduction

MARK JERRUM

One of the great triumphs of theoretical computer science in recent years has been the classification of combinatorial optimisation problems according to the accuracy within which their solutions can be approximated in polynomial time. Many degrees of approximability may be exhibited, ranging from approximation within ratio $1 \pm \varepsilon$ for any $\varepsilon > 0$, to non-approximability within ratio n^α for some $\alpha > 0$. Counting problems are different: under weak assumptions, approximation is either (a) possible within ratio $1 \pm \varepsilon$ in time $\mathrm{poly}(n, \varepsilon^{-1})$ or (b) impossible within any polynomial factor in polynomial time. (As usual, n is some measure of input size.) In case (a) we say that a "fully polynomial randomised approximation scheme" or "FPRAS" exists for the problem.

The most successful approach to constructing FPRAS's has been the "Markov chain Monte-Carlo" or "MCMC" method. In this, information about a set of combinatorial structures is obtained by simulating a random walk on the set of structures. Rapid mixing (i.e., convergence to equilibrium) of this random walk may be proved by many methods: canonical paths, geometric arguments or coupling. Non-approximability results are hard to come by: existing ones amount to little more than observations.

On the average case behavior of FFD and BFD under discrete distributions

DAVID S. JOHNSON

We investigate the asymptotic average case behavior of the First Fit Decreasing (FFD) and Best Fit Decreasing (BFD) bin packing algorithms under discrete probability distributions, in particular, ones that allow only a finite number of distinct item sizes. We show that for any such distribution and any of a wide class of algorithms, including FFD and BFD, the expected wasted space in packings generated by the algorithm is either $\Theta(n)$, $\Theta(\sqrt{n})$, or $O(1)$, where n is the number of items. We also provide an algorithm to determine which alternative applies by treating the number of items as a continuous quantity and performing one or more runs of a *fluid packing* algorithm. The algorithm runs in exponential time, which is probably unavoidable since the problem of determining which case holds, given a distribution, is NP-hard.

For the important special case of the discrete uniform distributions $U\{j,k\}$, where items are of size $1/k, 2/k, \ldots, j/k$, for $1 \leq j \leq k$ and bins have unit size, the fluid algorithm runs in time polynomial in k. Here the expected waste is $\Theta(\sqrt{n})$ when $j \in \{k-1, k\}$ and $O(1)$ when $j < \sqrt{k}$ or $j > k - \sqrt{k}$, but is less predictable between these extremes, as we show by using the fluid packing algorithm to determine which alternative applies for each pair (j,k) with $k \leq 2,500$. Both linear waste ($\Theta(n)$) and constant waste ($O(1)$) occur throughout the intermediate region. Even in the case of linear waste, however, the packings remain reasonably efficient. For instance, for all discrete uniform distributions $U\{j,k\}$, the asymptotic ratio of the expected number of bins under FFD or BFD to the expected optimal number is bounded by $169/168 = 1.00595...$, which occurs when $j = 6$ and $k = 13$ (a distribution for which the online algorithm Best Fit surprisingly yields constant waste).

This is joint work with E.G. Coffman, Jr, P.W. Shor, and R.R. Weber.

Collecting the best approximation bounds – is it really a good idea?

VIGGO KANN

A compendium of NP optimization problems, containing the best approximation results known for each problem, is available on the web at

`http://www.nada.kth.se/ viggo/problemlist/`

The compendium is modeled after the famous list of NP complete problems by Garey and Johnson, but is continuously updated with new results. This resource is much used (over 40 accesses every day) and appreciated by researchers in the area from the whole world.

Since the research area of approximation algorithms is very active there are hundreds of new results every year that must be entered into the compendium. In order to minimize the work needed by the authors (Pierluigi Crescenzi and Viggo Kann) for updating the compendium three web forms have been constructed and made available on the home page of the compendium. One form is used for reporting improved results, one for reporting results for new problems, and one for reporting errors in the compendium. We hope that in the future every researcher publishing a new approximation result will report it to us using the web forms.

An Approximation Algorithm for the Bandwidth of Dense Graphs

MAREK KARPINSKI

We study computational complexity of the problem of approximating the Bandwidth of graphs, and matrices. We display an evidence that there are no efficient PTAS for this problem, and further design the first constant ratio approximation algorithm for it on NP-hard dense instances. The best approximation ratio achievable by our method equals 3.

Joint work with J. Wirtgen and A. Zelikovsky

Greedy Strikes Back: Improved Facility Location Algorithms

SAMIR KHULLER

A fundamental facility location problem is to choose the location of facilities, such as industrial plants and warehouses, to minimize the cost of satisfying the demand for some commodity. There are associated costs for locating the facilities, as well as transportation costs for distributing the commodities. This problem is commonly referred to as the *uncapacitated facility location* problem. Hochbaum had shown that a greedy set covering approach yields an approximation bound of $O(\log n)$. Recently, the first constant factor approximation algorithm for this problem was obtained by Shmoys, Tardos and Aardal (STOC 1997).

We show that a simple and natural greedy heuristic combined with the Shmoys, Tardos and Aardal algorithm, can be used to derive a better approximation guarantee. We discuss a few variants of the problem, demonstrating better approximation factors for restricted versions of the problem.

We also show that the problem is Max SNP Hard. However, the inapproximability constants derived from the Max SNP hardness are very close to one. By relaxing the assumption that $P \neq NP$ to $NP \not\subseteq DTIME[n^{O(\log\log n)}]$ we provide a much stronger lower bound on the best possible approximation ratio. This is surprising because we show that the problem can be approximated to within small constants. We expect this reduction approach to be extended to other problems to provide stronger lower bounds on the best possible approximation factors.

Joint work with Sudipto Guha.

A polynomial-time approximation scheme for weighted planar TSP
PHIL KLEIN

Given a planar graph on n nodes with costs (weights) on its edges, define the distance between nodes i and j as the length of the shortest path between i and j. Consider this as an instance of *metric* TSP. For any $\varepsilon > 0$, our algorithm finds a salesman tour of total cost at most $(1+\varepsilon)$ times optimal in time $n^{O(1/\varepsilon^2)}$.

Approximation algorithms for scheduling problems with communication delays
ROLF H. MÖHRING

In the last few years, multi-processor scheduling with interprocessor communication delays has received increasing attention. This is due to the more realistic constraints in modeling parallel processor systems.

We develop approximation algorithms for minimizing the makespan and the weighted sum of completion times for "small" communication delays, arbitrary precedence constraints, and parallel identical processors.

The common underlying idea of our algorithms is to compute first a schedule by LP-relaxation and rounding that regards all constraints except for the processor restrictions. This schedule is then used to construct a provable good feasible schedule for a limited number of processors by a suitable generalization of Graham's list scheduling rule. We thus obtain a simple $\frac{7}{3}$-approximation algorithm for the makespan, and a 6.14232-approximation algorithm for the weighted sum of completion times.

Complementing these results, we obtain new polynomial-time exact algorithms and NP-completeness results on the class of series-parallel precedence constraints.

This is joint work with Markus Schäffter and Andreas Schulz, Berlin.

Relaxed Multicommodity Flow and its Applications in the Design of Approximation Algorithms for Cut-Related Problems
SEFFI NAOR

We introduce the notion of relaxed multi-commodity flow and show its usefulness in designing approximation algorithms. The problems considered are cut-related problems in graphs. A common framework for designing approximation algorithms for such problems is the primal-dual method. In many cases, however, the integrality gap is large, and thus, it is not possible to obtain small (constant) approximation factors via this approach. We have developed a technique for overcoming this obstacle in certain cases by defining A new version of multi-commodity flow which we call *relaxed multi-commodity flow*. The main feature of this new flow function is that it distinguishes between inter-commodity and intra-commodity constraints, and relaxes the inter-commodity constraint. We apply this technique to two problems: the directed multiway cut problem and the subset feedback set problem. In both cases we obtain constant approximation factors which improve upon logarithmic approximation factors previously known. The following polynomial time approximation algorithms are obtained:

1. An approximation algorithm that achieves a factor of 2 for the minimum weight multiway cut problem in directed graphs.
2. An approximation algorithm that achieves a factor of 2 for the minimum weight subset feedback edge set problem.
3. An approximation algorithm that achieves a factor of 8 for the minimum weight subset feedback vertex set problem.

This is joint work with Guy Even, Baruch Schieber, and Leonid Zosin.

Approximating Network Inhibition

CYNTHIA A. PHILLIPS

In the network inhibition problem, we wish to expend a limited budget removing edges or pieces of edges so as to minimize the resulting maximum s-t flow. The problem is strongly $\mathcal{NP}$-hard. Previous approximation algorithms applied only to planar graphs. In this talk, we give a polynomial-time algorithm, based on a linear-programming relaxation of an integer program, that finds an attack with cost B_a and residual network capacity (max flow) C_a such that

$$\frac{B_a}{B} + \epsilon\frac{C_a}{C^*} \leq 1 + \epsilon,$$

where $\epsilon > 0$, B is the budget (the amount of resources to expend damaging the network), and C^* is the minimum (optimal) residual capacity for any attack with budget B. For example, our algorithm returns a $(1, 1 + 1/\epsilon)$-approximation or a $(1 + \epsilon, 1)$-pseudoapproximation, but we do not know which *a priori*. The parameter ϵ biases the nature of the solution, but does not effect the running time. We give a generalization to multiple attack methods/budgets and show that computing the most cost-effective attack is in $\mathcal{P}$.

This is joint work with Carl Burch (CMU), Sven Krumke (U. Würzburg), Madhav Maranthe (Los Alamos National Laboratory), and Eric Sundberg (Rutgers).

Two-Machine No-Wait Flow Shop Scheduling with Missing Operations

CHRIS N. POTTS

This paper considers the no-wait scheduling of n jobs in a two-machine flow shop, where some jobs require processing on the first machine only. The objective is to minimize the maximum completion time, or makespan. In view of its NP-hardness, we propose three heuristic algorithms, each of which is based on solving optimally a related problem that has no missing operations using, for example, the well-known algorithm of Gilmore and Gomory. An analysis of the worst-case performance of these algorithms is given. The main result is an $O(n \log n)$-time heuristic which generates a schedule with makespan no more than 4/3 times that of an optimal schedule. A new proof technique is employed, which has potential application in a variety of bin packing and parallel machine scheduling problems.

Joint work with C.A. Glass and J.N.D. Gupta.

Online Throughput-Competitive Algorithm for Multicast Routing and Admission Control

MONIKA R. HENZINGER

We present the first polylog-competitive online algorithm for the general multicast problem in the throughput model. The ratio of the number of requests accepted by the optimum offline algorithm to the expected number of requests accepted by our algorithm is $O(\log \mathcal{M}(\log n + \log\log \mathcal{M}) \log n)$, where $\mathcal{M}$ is the number of multicast groups and n is the number of nodes in the graph. We show that this is close to optimum by presenting an $\Omega(\log n \log \mathcal{M})$ lower bound on this ratio for any randomized online algorithm against an oblivious adversary, when $\mathcal{M}$ is much larger than the link capacities. We also show that it is impossible to be competitive against an adaptive online adversary.

As in the previous online routing algorithms, our algorithm uses edge-costs when deciding on which is the best path to use. In contrast to the previous competitive algorithms in the throughput model, our cost is not a direct function of the edge load. The new cost definition allows us to decouple the effects of routing and admission decisions of different multicast groups.

Joint work with Ashish Goel and Serge Plotkin.

Coin Tosses and LPs: New Approximations in Scheduling

ANDREAS SCHULZ

We present a new class of randomized approximation algorithms for scheduling problems by directly interpreting solutions to so-called time-indexed LPs as probabilities. The most general model we consider is scheduling unrelated parallel machines with release dates (or even network scheduling) so as to minimize the average weighted completion time. The crucial idea for these multiple machine problems is not to use standard list scheduling but rather to assign jobs randomly to machines (with probabilities taken from an optimal LP solution) and to perform list scheduling on each of them.

For the general model, we give a $(2 + \epsilon)$-approximation algorithm. The best previously known approximation algorithm has a performance guarantee of $(16/3 + \epsilon)$ [Hall, Shmoys, & Wein 1996]. Our algorithm improves as well upon the best previously known approximation algorithms for the special case of identical parallel machine scheduling. Moreover, in this case our analysis also reveals an interesting insight into scheduling by α-points. In addition, for identical parallel machines the algorithm has running time $O(n \log n)$ and can easily be turned into an on-line algorithm with the same performance ratio.

This is joint work with Martin Skutella.

Approximating Weighted Completion Time Schedules for Parallel Jobs

UWE SCHWIEGELSHOHN

A parallel machine with m identical resources is assumed. Further, there are n independent jobs to be scheduled on the machine such that the total weighted completion time is minimized. Each job i has weight the w_i and requires m_i resources for a time period t_i. First, a preemptive algorithm is proposed which achieves a performance guarantee of 2.37 times the optimum. This algorithm is based on the observation that systems can be scheduled in a non preemptive fashion with a performance guarantee of 2 using a simple list schedule, if for all jobs either $m_i \leq \frac{m}{2}$ or $m_i > \frac{m}{2}$ holds. Therefore, the jobs are divided into those two classes and two independent schedules are generated. Then, both schedules are interleaved. The algorithmic analysis is partly based on the results of Kawaguchi and Kyan (SIAMJC 1986). In a second part the preemptive schedule is transformed into a non preemptive schedule by expanding the schedule and thus separating jobs which cause preemption from the other jobs. By careful selection of the parameters this increases the total weighted completion time of the preemptive schedule by at most a factor of 3. Both results together yield a performance guarantee of 7.11 for non preemptive schedules improving on the previously best known value of 8.53.

Compact vector summation: how it works in different fields of discrete mathematics

SERGEY V. SEVASTIANOV

We consider a geometrical problem called *compact vector summation problem.* Suppose that we have an m-dimensional real space $\mathbb{R}m$ with a norm s, and we are given a finite family of "short" vectors $X = \{x_1, \dots, x_n\} \subset \mathbb{R}m$, $\|x_j\|_s \leq 1$, with zero sum: $\sum x_j = 0$. (Such family of vectors is called *s-family.*) We wish to find the permutation $\pi = (\pi_1, \dots, \pi_n)$ for which all partial sums $\sum_{j=1}^{k} x_{\pi_j}$, $k = 1, \dots, n$, are within a ball of a minimum possible radius:

$$\max_k \left\| \sum_{j=1}^{k} x_{\pi_j} \right\|_s \to \min_{\pi}.$$

We show that for any norm s and any s-family of vectors $X \subset \mathbb{R}m$ they can be summed within a ball of radius m, and the desired permutation π can be found in polynomial time ($O(n^2m^2)$).

Next we demonstrate how this result can be applied in different fields of mathematics:

- to prove the Steinitz theorem about the domain of sums of a conditionally convergent series in m-dimensional space;
- to prove Gordon's lemma about a basic family of integral solutions of a homogeneous linear system ($Ax = 0$);
- to find an upper bound on the maximum degree $D(n)$ for which there exists a $D(n)$-regular indecomposable hypergraph on n vertices;
- to find an "almost even" partition of a given annual plan of an enterprise into l parts (4 quarters, 12 month plans, etc.);
- to construct asymptotically optimal solutions to NP-hard shop-scheduling problems (job-shop, e.g.) in polynomial time.

Greed is good, but guessing is better

DAVID B. SHMOYS

We consider approximation algorithms for the *metric uncapacitated facility location problem*: given a set of demand points D and a set of potential facility locations F in a common metric space, where the distance between demand point $j \in D$ and location $i \in F$ is c_{ij}, and, for each $i \in F$, f_i is the cost of opening a facility at location i, find $S \subseteq F$ at which to open facilities such that $\sum_{i \in S} f_i + \sum_{j \in D} \min_{i \in S} c_{ij}$ is minimized. Shmoys, Tardos, and Aardal had previously given a 3.16-approximation algorithm based on rounding an optimal fractional solution to the linear relaxation of the natural integer programming formulation of this problem. We show that their techniques, when combined with randomized rounding as well as information given by the optimal dual solution, leads to a randomized approximation algorithm with expected performance guarantee equal to 2.034.

This is work done jointly with (and primarily by) my Ph.D. student, Fabián A. Chudak.

Random methods and geometric algorithms

MIKLÓS SIMONOVITS

Calculating the volume of an m-dimensional convex polytope is NP-hard: actually it contains several subcases which are intractable. The best one can hope for is to find some approximation, but even this is impossible by deterministic algorithms (Elekes, Bárány–Füredi). Following the break-through of Broder, Jerrum and Sinclair in the

Permanent Problem, Dyer, Frieze and Kannan gave a polynomial approximation algorithm of a convex body given by a weak separation oracle. Several variations, improvements and modification of the original algorithm were found. All these algorithms use multiphase Monte Carlo methods, contain random walks on some (possibly infinite) geometrically defined graphs, use some rounding phase and establish the fact that the Markov chain corresponding to the random walk is rapidly mixing. It turns out that uniform sampling, integration of log-concave functions over convex domains and calculating the volume are (in some sense) equivalent.

Currently the authors, Kannan, Lovász and Simonovits gave a randomized algorithm which, with probability at least $1-\delta$, approximates the volume of a convex body in the n-dimensional space in roughly $O(n^5)$ steps, with relative error smaller than ε. An interesting feature of this algorithm is that it uses the isotropic position of the body. In some sense it also calculates an "almost" Löwner-John ellipsoid in $O^*(n^5)$ steps.

There is a sharp difference between approximating the volume and the diameter (or the width) of a convex body: this later one is superpolynomial even if we use randomized algorithms:

Theorem. Consider a randomized algorithm estimating the diameter of an arbitrary convex body $K \subseteq R^n$, given by a separation oracle. If the algorithm provides a lower bound $L(K)$ and an upper bound $U(K)$ so that

$$L(K) \leq diam(K) < U(K)$$

and

$$U(K)/L(K) = o\left(\sqrt{\frac{n}{\log n}}\right)$$

then it must use superpolynomially many question on the oracle.

This theorem is sharp. The question: "How many steps are needed to approximate the diameter of a convex body" with a given relative errror η (where $2 \leq \eta \leq \sqrt{n}$) is strongly connected to the covering problem of a high dimensional sphere by smaller spherical caps.

Joint work with Ravi Kannan and L. Lovász.

A new randomized approximation algorithm for the Steiner tree problem

ANGELIKA STEGER

In this talk we present an RNC-algorithm for finding a minimum spanning tree in a weighted 3-uniform hypergraph, assuming the edge weights are given in unary, and a fully polynomial time randomized approximation scheme if the edge weights are given in binary. From this result we then derive RNC-approximation algorithms for the Steiner problem in networks with approximation ratio $(1+\varepsilon)\,5/3$ for all $\varepsilon > 0$.

Approximation Algorithms for Unsplittable Flow Problems
CLIFFORD STEIN

In the unsplittable flow problem, we are given a network and set of commodities (source-sink pairs) with associated demands. We seek a single source-sink flow path for each commodity so that the demands are satisfied and the total flow routed across any edge is bounded by its capacity. The problem is an NP-hard variant of max flow and a generalization of single-source disjoint paths with applications to scheduling, load balancing and virtual-circuit routing problems.

We give a simple framework which leads to small constant-factor approximation algorithms for, in the single source case, minimizing congestion, maximizing total flow, and minimizing the number of rounds needed to route all the flow. In the multiple source case, this framework yields an approximation algorithm for maximizing the total flow, and for the disjoint paths problem.

This is joint work with Stavros Kolliopoulos.

Approximating Satisfiable Satisfiability Problems
LUCA TREVISAN

We study the approximability of the Maximum Satisfiability Problem (MAXSAT) and of the boolean k-ary Constraint Satisfaction Problem (MAXkCSP) restricted to satisfiable instances. For both problems we improve on the performance ratios of known algorithms for the unrestricted case.

Our approximation for satisfiable MAX3CSP instances is better than any possible approximation for the unrestricted version of the problem (unless $P = NP$). This result implies that the requirement of perfect completeness weakens the acceptance power of non-adaptive PCP verifiers that read 3 bits.

We also present the first non-trivial results about PCP classes defined in terms of free bits that collapse to P.

Primal-Dual Approximation Algorithms for Feedback Vertex Problems in Planar Graphs
DAVID WILLIAMSON

Given a subset of cycles of a graph, we consider the problem of finding a minimum-weight set of vertices that meets all cycles in the subset. This problem generalizes a number of problems, including the minimum-weight feedback vertex set problem in both directed and undirected graphs, the subset feedback vertex set problem, and the graph bipartization problem, in which one must remove a minimum-weight set of vertices so that the remaining graph is bipartite. We give a $\frac{9}{4}$-approximation algorithm for the general problem in planar graphs, given that the subset of cycles obeys certain properties. This results in $\frac{9}{4}$-approximation algorithms for the aforementioned feedback and bipartization problems in planar graphs. Our algorithms use the primal-dual method for approximation algorithms as given in the survey of Goemans and Williamson (Chapter 4 in Hochbaum, ed., "Approximation Algorithms for NP-hard Problems," PWS, 1997).

One for the Price of Two: A Unified Approach for Approximating Covering Problems
REUVEN BAR-YEHUDA

We present a simple and unified approach for developing and analyzing approximation algorithms for covering problems. We illustrate this on approximation algorithms for the following problems: Vertex Cover, Set Cover, Feedback Vertex Set, Generalized Steiner Tree and related problems.

The main idea can be phrased as follows: iteratively, pay two dollars (at most) to reduce the total optimum by one dollar (at least), so the rate of payment is no more than twice the rate of the optimum reduction. This implies a total payment (i.e., approximation cost) $\leq$ twice the optimum cost.

Our main contribution is based on a formal definition for covering problems, which includes all the above fundamental problems and others. We further extend the Bafna, Berman and Fujito Local-Ratio theorem. This extension eventually yields a short generic r-approximation algorithm which can generate most known approximation algorithms for most covering problems.

Another extension of the Local-Ratio theorem to randomized algorithms gives a simple proof of Pitt's randomized approximation for Vertex Cover. Using this approach, we develop a modified greedy algorithm, which for Vertex Cover, gives an expected performance ratio ≤ 2.

Open Problems

A collection of open problems

DAVID WILLIAMSON

During the open problems session at Dagstuhl, I listed the following 10 problems as interesting and (in my opinion) eminently solvable, except possibly for problems 8 and 9. For each I will give a cash reward for solving the problem, to be awarded if a paper containing the result is accepted to a refereed journal by 1/1/1 (or another arrangement at the judge's discretion).

1. ($30) Find an approximation algorithm for the Rectilinear Steiner Arborescence problem with a constant performance guarantee less than 2. See Rao, Sadayappan, Hwang, and Shor, "The Rectilinear Steiner Arborescence Problem," *Algorithmica*, 7:277-288, 1992.

   Claimed by Sanjeev Arora and Luca Trevisan, August 22, 1997.

2. ($40) Find a 2-approximation algorithm for the Survivable Network Design Problem with multiple edge copies allowed. See Section 4.2 of Goemans and Williamson, "A General Approximation Technique for Constrained Forest Problems," *SIAM Journal on Computing*, 24:296-317, April 1995.

3. ($40) Find a constant approximation algorithm for the scheduling problem $Q|prec|C_{\max}$ (i.e. scheduling related machines with precedence constraints so as to minimize makespan). See Chudak and Shmoys, "Approximation algorithms for precedence-constrained scheduling problems on parallel machines that run at different speeds," SODA '97, pp. 581-590.

4. ($30) Give a 2-approximation algorithm for the subset feedback vertex set problem in undirected graphs. See Even, Naor, and Zosin, "An 8-approximation algorithm for the subset feedback vertex set problem," FOCS '96, 310-319.

5. ($40) Give a new proof of the Linial-London-Rabinovich $O(\log n)$-distortion embedding for arbitrary metrics by using the dual of Garg's semidefinite program. See Linial, London, and Rabinovich, "The geometry of graphs and some of its algorithmic applications," *Combinatorica*, 15:215-246, 1995; and Garg, "A deterministic $O(\log k)$-approximation algorithm for the sparsest cut problem" (request from Naveen Garg, `naveen@mpi-sb.mpg.de`).

6. ($40) Find an approximation algorithm for the Steiner tree problem with constant performance guarantee better than 2 that bounds against the bidirected linear programming formulation of Steiner tree. Alternatively, give an example that shows that a factor of 2 is the best achievable using this formulation. See p. 321-322 of Chopra, Gorres, and Rao, "Solving the Steiner Tree problem on

a graph using branch and cut," *ORSA Journal of Computing* 4:320-335, 1992, for a definition of the bidirected formulation and empirical evidence that it performs well in practice.

7. ($40) Show that every weighted connected graph can be $O(\log n)$-probabilistically-approximated by tree metrics. See Bartal, "Probabilistic Approximation of Metric Spaces and its Algorithmic Applications," FOCS '96, 184-193.
8. Find an improved approximation algorithm for the sparsest cut problem. See Shmoys, "Cut problems and their application to divide-and-conquer," in Hochbaum, ed., "Approximation algorithms for NP-hard problems," PWS, 1997. $50 for a performance guarantee of $o(\log n)$, $100 for a constant.
9. Analyze the max cut semidefinite program given the addition of triangle constraint inequalities. See Goemans and Williamson, "Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming," *JACM*, 42:1115-1145, 1995; and Karloff, "How good is the Goemans-Williamson MAX CUT algorithm?," STOC '96, 427-434. $100 for a tight analysis, otherwise a sliding scale for improvement from .878 up to 16/17.
10. ($50) Give a combinatorial .878-approximation algorithm for MAX CUT (see above).

Some Open Problems

MARK JERRUM

Permanent of a 0,1-matrix. A counting problem is specified by a function $f : \Sigma^* \to \mathbb{N}$ mapping problem instances encoded in some alphabet Σ to numbers. A *randomised approximation scheme* for f is a randomised algorithm that takes as input an instance $x \in \Sigma^*$ and a tolerance $\varepsilon > 0$, and outputs a number Z (a random variable) satisfying

$$\Pr\Big((1-\varepsilon)f(x) \le Z \le (1+\varepsilon)f(x)\Big) \ge \frac{3}{4}.$$

The randomised approximation scheme is *fully polynomial* if its running time is $\mathrm{poly}(n, \varepsilon^{-1})$. The phrase "Fully Polynomial Randomised Approximation Scheme" is often abbreviated to FPRAS.

The permanent of an $n \times n$ matrix $A = (a_{ij})$ with (say) integer entries is

$$\mathrm{per}A = \sum_{\pi} a_{1,\pi(1)} a_{2,\pi(2)} \cdots a_{n,\pi(n)},$$

where the sum ranges over all permutations π of $\{1, \ldots, n\}$. Formally, $\mathrm{per}A$ is similar to $\det A$ except that all terms have positive sign. Is there an FPRAS for the permanent of a 0,1-matrix? The crucial point here is that the matrix is non-negative.

If negative entries are allowed, it is hard to determine even the correct sign for per$A$. On the other hand, there is a technique for encoding the permanent of a matrix with arbitrary non-negative integer entries as a permanent of a (not too much bigger) 0,1-matrix. FPRASs are known for some special cases, for example when the matrix A is "dense" [JS96], but the existence of an FPRAS for general 0,1-matrices is open.

In the absence of an FPRAS, Barvinok asks whether per$A$ can be approximated in a weaker sense. More precisely, is there, for every $\varepsilon > 0$, a polynomial-time randomised algorithm that outputs a number Z satisfying

$$\Pr\Big((1-\varepsilon)^n f(x) \le Z \le (1+\varepsilon)^n f(x)\Big) \ge \frac{3}{4}?$$

See [B97] for an important step in this direction.

Random walk on the bases of a matroid. There is a very natural random walk on the bases of a matroid M on ground set S: if the current base is B, select $e, f \in S$ uniformly at random and let $B' \leftarrow B \oplus \{e, f\}$, where $\oplus$ denotes symmetric difference; if B' is a base then move to B', otherwise stay at B. Is this random walk rapidly mixing? Precisely, is the t-step distribution within variation distance ε of the stationary distribution in $t = \mathrm{poly}(n, \log \varepsilon^{-1})$ steps, where $n = |S|$? Feder and Mihail [FM92] show rapid mixing for "balanced" matroids, a class of matroids that strictly includes regular matroids. An extension of this result to an interesting class of non-regular matroids would lead to new FPRASs: e.g., for forests of given size in a graph (bases of a truncation of a graphic matroid), or connected spanning subgraphs of given size in a graph (a truncation of the dual of a graphic matroid). Note that the latter counting problem corresponds to "all-terminal reliability" estimation for a network with independent edge failures. Rapid mixing is not known even for binary matroids.

Random walk on the independent sets of a matroid. There is an equally natural random walk on the independent sets of a matroid M on ground set S: if the current independent set is I, select $e \in S$ uniformly at random and let $I' \leftarrow I \oplus \{e\}$; if I' is an independent set then move to I', otherwise stay at I. Even less is known about this random walk than the one on bases: rapid mixing is not known even for graphic matroids. (Note that regular matroids include graphic matroids.)

Words of given length in a context free grammar. Is there an FPRAS for the following problem: given a context-free grammar G (with terminal symbols Σ) and a positive integer m (in unary notation), how many words of length m are there in the language $L(G)$ generated by G? Equivalently, and perhaps more naturally, given G and m, select a word u.a.r. from $L(G) \cap \Sigma^m$. The difficulty lies in ambiguity: if G is an unambiguous grammar then the required number could be computed exactly by dynamic programming. Gore et al. [GJKSM97] provide a "quasi-polynomial" approximation algorithm—i.e., one with running time $\exp(\mathrm{polylog}(|G|+m))$—based on iterated Karp-Luby sampling from unions of sets. The existence on an FPRAS is open. The non-uniform version of the problem, in which G is not considered to be part of the input, appears to makes sense, and is also open, to the best of my knowledge.

Interesting non-approximability results. There can be no FPRAS for the number of independent sets of size at least k in a graph G unless RP = NP, since an FPRAS must, a fortiori, distinguish between no independent sets of size k, and some. Somewhat less trivially, it is known [Sin93] that there can be no FPRAS for counting independent sets of all sizes, unless RP = NP. Roughly, this is because the graph G can be "powered" to a graph G' in such a way that almost all independent sets in G' correspond to maximum independent sets in G. It would be a significant advance to have an example of a counting problem that is non-approximable for some more interesting reason. The general 0,1-permanent forms an excellent test case for those who conjecture that it does not admit an FPRAS.

[B97] Alexander Barvinok, *A simple polynomial time algorithm to approximate the permanent within a simply exponential factor*, MSRI preprint #1997-031.

[FM92] Tomás Feder and Milena Mihail, Balanced matroids, *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*, ACM Press, 1992, 26–38.

[GJKSM97] Vivek Gore, Mark Jerrum, Sampath Kannan, Z. Sweedyk and Steve Mahaney, A quasi-polynomial-time algorithm for sampling words from a context-free language, *Information and Computation 134* (1997), 59–74.

[JS96] Mark Jerrum and Alistair Sinclair, The Markov chain Monte Carlo method: an approach to approximate counting and integration. In *Approximation Algorithms for NP-hard Problems* (Dorit Hochbaum, ed.), PWS, 1996, 482–520.

[Sin93] A. J. Sinclair, *Randomised algorithms for counting and generating combinatorial structures*, Advances in Theoretical Computer Science, Birkhäuser, Boston, 1993.

A List of Open Problems

MAREK KARPINSKI

1. Verify the approximation hardness status of the MINIMUM BISECTION problem (MAXIMUM BISECTION is known to be MAX-SNP hard). The problem admits a PTAS on dense instances (Arora, Karger, Karpinski, 1995), and the standard reductions from MAXIMUM BISECTION do not seem to yield lower approximation bounds for this problem.

2. Design sublinear approximation algorithms working in polynomial time for the MINIMUM BISECTION.

3. Verify approximation hardness of the MINIMUM BANDWIDTH problem. The problem does not admit an efficient PTAS (Karpinski, Wirtgen, 1997). The problem admits a 3-approximation algorithms on dense instances (Karpinski, Wirtgen, Zelikovsky, 1997). How about the existence of a PTAS for this problem on dense instances?

4. Design a sublinear approximation algorithm working in polynomial time for the MINIMUM BANDWIDTH.

5. Design a sublogarithmic approximation algorithm working in polynomial time for the MINIMUM SET COVER problem on dense instances. MINIMUM SET COVER on dense instances is known to be approximable within ratio $c\ln(n)$ for any $c > 0$ (Karpinski, Zelikovsky, 1997).

[AKK95] Arora, S., Karger, D., and Karpinski, M.(1995), "Polynomial Time Approximation Schemes for Dense Instances of NP-Hard Problems", Proc. 27th ACM STOC (1995), pp. 284-293.

[K97] Karpinski, M.(1997), "Polynomial Time Approximation Schemes for Some Dense Instances of NP-Hard Optimization Problems", Proc. 1st Symp. on Randomization and Approximation Techniques in Computer Science, RANDOM'97, Lecture Notes in Computer Science Vol.1269, Springer Verlag, 1997, pp. 1-14.

[KW97] Karpinski, M., and Wirtgen, J.(1997), "On Approximability Hardness of the Bandwidth Problem", Research Report, Dept. Computer Science, Univ. Bonn, 1997.

[KWZ97] Karpinski, M., Wirtgen, J., and Zelikovsky, A.(1997), "An Approximation Algorithm for the Bandwidth Problem on Dense Graphs", ECCC Technical Report TR 97-017.

[KZ97] Karpinski, M., and Zelikovsky, A.(1997),"Approximating Dense Cases of Covering Problems", ECCC Technical Report TR 97-004.)

An Open Problem on Graph Spanners

PHIL KLEIN

This spanner question would be useful for generalizing the approximation scheme for weighted planar TSP: Given a weighted planar graph G, a subset S of nodes, and a parameter ε, does there exist a subgraph H such that

- the weight of H is at most a polynomial in $1/\varepsilon$ times the weight of a Steiner tree in G spanning the nodes in S, and
- for every pair of nodes u, v belonging to S, the u-to-v distance in H is at most $1+\varepsilon$ times the u-to-v distance in G.

Some Open Problems

ANDREAS SCHULZ

Acyclic subdigraph and linear ordering problem. Trivial 2-approximation algorithms are known. Find an approximation algorithm with worst case performance better than 2 or prove that this is impossible.

Total completion time with release dates. There is one machine and n jobs J_j, $j = 1, \ldots, n$, with processing times p_j and release dates r_j. The goal is to find a feasible schedule that minimizes the sum of all job completion times. An $e/(e-1)$-approximation algorithm is known. Does their exist a PTAS? An FPTAS is not possible unless P=NP, since the problem is strongly NP-hard.

Total completion time with precedence constraints. There is one machine and n precedence constrained jobs J_j, $j = 1, \ldots, n$, with processing times p_j and weights w_j. The goal is to find a feasible schedule that minimizes the weighted sum of job completion times. Several 2-approximation algorithms are known for this problem. Does their exist an approximation algorithm with performance guarantee strictly less than 2? Does there exist a PTAS for the special case where $w_j = 1$ for all jobs J_j? Or can we show that this special case is not easier than the general case with arbitrary nonnegative weights?

An Open Problem on Non Preemptive Weighted Completion Time Schedule for Parallel Jobs

UWE SCHWIEGELSHOHN

A parallel machine with m identical resources is assumed. Further, there are n independent jobs to be scheduled on the machine in a non preemptive fashion such that the total weighted completion time is minimized. Each job i has weight the w_i and requires m_i resources for a time period t_i. Find a lower bound (or a PTAS) for the problem.

Known Results: 3 different approximation algorithm are known:

- Chakrabarti et al. [CPS96] presented a method which gives an approximation factor of 12 in the deterministic case and an approximation factor of 8.67 if randomization is allowed.
- The use of the SMART algorithm [TSWY94a] gives an approximation factor of 8.53 (deterministic) [SLWTY98].
- The transformation of a preemptive schedule into a non preemptive one results in an approximation factor of 7.11 (see the abstracts of the meeting).

There are no results known on lower bounds. However, in the case $m_i = 1$ for all jobs Sahni [Sah76] gave a PTAS.

Award: \$ 50.- (the rules defined by David Williamson apply.)

[CPS96] S. Chakrabarti, C. Phillips, A.S. Schulz, D.B. Shmoys, C. Stein, and J. Wein. Improved approximation algorithms for minsum criteria. In *Proceedings of the 1996 International Colloquium on Automata, Languages and Programming.* Springer Verlag Lecture Notes in Computer Science, 1996.

[Sah76] S. Sahni. Algorithms for scheduling independent tasks. *Journal of the ACM*, 23:116–127, 1976.

[SLWTY98] U. Schwiegelshohn, W. Ludwig, J.L. Wolf, J.J. Turek, and P. Yu. Smart SMART bounds for weighted response time scheduling. *SIAM Journal on Computing.* Accepted for publication.

[TSWY94a] J.J. Turek, U. Schwiegelshohn, J.L. Wolf, and P. Yu. Scheduling parallel tasks to minimize average response time. In *Proceedings of the 5^{th} SIAM Symposium on Discrete Algorithms*, pages 112–121, 1994.

A Favorite Open Problem

DAVID B. SHMOYS

In this scheduling problem, there are n jobs and m machines, where each job j has a processing time p_j and a subset S_j of machines that are capable of processing it. Each job j must be assigned to exactly one machine $i \in S_j$. The load on a machine is total processing time of jobs assigned to it. We wish to find an assignment that minimizes the maximum machine load. Lenstra, Shmoys, and Tardos give a 2-approximation algorithm, and show, for any $\rho < 3/2$, that there does not exist a ρ-approximation algorithm for this problem. Give an approximation algorithm with constant performance guarantee better than 2.

Dagstuhl-Seminar 9734 — List of Participants

Karen **Aardal**
Universiteit Utrecht
Department of Computer Science
Padualaan 14
NL-3508 TB Utrecht
The Netherlands
aardal@cs.ruu.nl
http://www.cs.ruu.nl/staff/aardal.html
tel: +31-30-253-4114

Sanjeev **Arora**
Princeton University
Department of Computer Science
35 Olden Street
Princeton NJ 08544
USA
arora@cs.princeton.edu

Yossi **Azar**
Tel Aviv University
Department of Computer Science
Ramat Aviv
Tel Aviv 69978
Israel
azar@math.tau.ac.il
tel: +972-3-640-63 54

Amotz **Bar-Noy**
Tel Aviv University
Dept. of Electrical Engineering Systems
Ramat Aviv
Tel Aviv 69978
Israel
amotz@eng.tau.ac.il
www.eng.tav.ac.il/~amotz
tel: +972-3-640-7766

Reuven **Bar-Yehuda**
Computer Science Department
Technion
Haifa 32000
Israel
reuven@cs.technion.ac.il
tel: +972-4-829-4332

Yair **Bartal**
ICSI
Suite 600
1947 Center Street
Berkeley CA 94704-1198
USA
yairb@icsi.berkeley.edu
http://www.icsi.berkeley.edu/~yairb
tel: +1-510-642-42 74 ext. 149

Alexander **Barvinok**
University of Michigan
Dept. of Mathematics
Ann Arbor MI 48109-1009
USA
barvinok@math.lsa.umich.edu
http://www.math.lsa.umich.edu/~barvinok
tel: +1-313-763-0294

Martin E. **Dyer**
University of Leeds
School of Computer Science
Leeds LS2 9JT
Great Britain
dyer@scs.leeds.ac.uk
tel: +44-532-335442

Guy **Even**
Universität des Saarlandes
Fachbereich 14 - Informatik
Lehrstuhl Prof. Paul
Postfach 15 11 50
66041 Saarbrücken
Germany
guy@cs.uni-sb.de
tel: +49-681-302-55 58

Uri **Feige**
Weizmann Institute
Departement of Appl. Mathematics
P.O. Box 26
76100 Rehovot
Israel
feige@wisdom.weizmann.ac.il
tel: +972-8-934-3364

Sandor **Fekete**
Universität Köln
Center for Parallel Computing
Weyertal 80
D-50923 Köln
Germany
sandor@zpr.uni-koeln.de
http://www.apr.uni-koeln.de/~sandor
tel: +49-221-470-6024

Naveen **Garg**
Max-Planck-Institut für Informatik
Im Stadtwald
66123 Saarbrücken
Germany
naveen@mpi-sb.mpg.de
http://www.mpi-sb.mpg.de/~naveen
tel: +49-681-9325-122

Magnus **Halldorsson**
Science Institute
Loftskeytastod - Sudurgotu
Dunhaga 4
IS-107 Reykjavik
Iceland
mmh@rhi.hi.is

Dorit **Hochbaum**
University of California
Department of IE & OR
Etcheverry Hall
Berkeley CA 94720-1777
USA
dorit@hochbaum.berkeley.edu
http://isor.berkeley.edu/~hochbaum/
tel: +1-510-642-4998

Mark **Jerrum**
University of Edinburgh
Department of Computer Science
King's Building
Mayfield Road
Edinburgh EH9 3JZ
Great Britain
mrj@dcs.ed.ac.uk
tel: +44-31-650-5134

David S. **Johnson**
AT&T Bell Labs
Room 2D-150
600 Mountain Avenue
Murray Hill NJ 07974-0636
USA
dsj@research.att.com
tel: +1-908-582-4742

Viggo Kann
Royal Institute of Technology
NADA KTH
Dept. of Numerical Analysis & Comp. Sc.
S-10044 Stockholm
Sweden
viggo@nada.kth.se
http://www.nada.kth.se/~viggo

David **Karger**
MIT
Lab. of Computer Science / NE43-321
545 Technology Square
Cambridge MA 02139
USA
karger@theory.lcs.mit.edu

Howard **Karloff**
Georgia Institute of Technology
College of Computing
GVU Center
Atlanta GA 30332-0280
USA
howard@cc.gatech.edu
tel: +1-404-894-9677

Marek **Karpinski**
Universität Bonn
Institut für Informatik
Römerstraße 164
53117 Bonn
Germany
marek@cs.uni-bonn.de
tel: +49-228-734-224 (Sekr.: 224)

Michael **Kaufmann**
Universität Tübingen
Wilhelm- Schickard- Institut für Informatik
Sand 13
D-72076 Tübingen
Germany
mk@informatik.uni-tuebingen.de
tel: +49-7071-2977-404

Claire **Kenyon**
ENS
LIP
46 Allée d'Italie
F-69364 Lyon
France
kenyon@ens-lyon.fr
tel: +33-4-72.72.85.04

Samir **Khuller**
Univ. of Maryland at College Park
Department of Computer Science
3217 AV Williams Bldg.
College Park MD 20742
USA
samir@cs.umd.edu

Philip **Klein**
Brown University
Department of Computer Science
P.O. Box 1910
Providence RI 02912
USA
klein@cs.brown.edu
tel: +1-401-863-76 44

Stefano **Leonardi**
Università di Rome "La Sapienza"
Dipartimento di Informatica e Sistemistica
Via Salaria 113
I-00198 Roma
Italy
leon@dis.uniroma1.it
tel: +39-6-49 91-83 41

Rolf H. **Moehring**
TU Berlin
FB Mathematik / MA 6-1
Str. d. 17. Juni 136
10623 Berlin
Germany
moehring@math.tu-berlin.de
tel: +49-30-314-24594

Seffi **Naor**
Technion
Computer Sciene Department
Haifa 32000
Israel
naor@cs.technion.ac.il
tel: +972-2-8294328

Cindy **Phillips**
Sandia National Laboratories
P.O. Box 5800
Albuquerque NM 87185
USA
caphill@cs.sandia.gov

Serge **Plotkin**
Stanford University
Department of Computer Science
Gates 472
Stanford CA 94305
USA
plotkin@cs.stanford.edu
http://www.stanford.edu/~plotkin

Chris **Potts**
University of Southampton
Faculty of Mathematics
Southampton SO17 1BJ
Great Britain
cnp@maths.soton.ac.uk
tel: +44-1-703-59 36 51

Hans-Jürgen **Prömel**
Humboldt Universität
Institut für Informatik
Unter den Linden 6
D-10099 Berlin
Germany
proemel@informatik.hu-berlin.de
tel: +49-30-20181-2 95

Yuval **Rabani**
Computer Science Department
Technion
Haifa 32000
Israel
rabani@csa.cs.technion.ac.il
http://www/cs/technion.ac.il/~rabani/
tel: +972-4-829-2056

Yuri **Rabinovich**
Ben-Gurion University
Dept. of Mathematics & Computer Science
Beer Sheva 85105
Israel
yuri@indigo.cs.bgu.ac.il
tel: +972-7-6461-623

Prabhakar **Raghavan**
IBM Almaden Research Center
650 Harry Road
San José CA 95120-6099
USA
pragh@almaden.ibm.com

Satish B. **Rao**
NEC Research Institute
4 Independence Way
Princeton NJ 08540
USA
satish@research.nj.nec.com

Monika **Rauch-Henzinger**
DEC SRC
130 Lytton Ave.
Palo Alto CA 94301
USA
monika@pa.dec.com
http://www.research.digital.com/src/personal/monika-henzinger/home.html
tel: +1-415-853-2129

R. **Ravi**
Carnegie Mellon University
GSIA
5000 Forbes Avenue
Pittsburgh PA 15213-389
USA
ravi@andrew.cmu.edu
http://www.gsia.cmu.edu/andrew/ravi/home.htlm
tel: +1-412-268-3694

Baruch **Schieber**
IBM T. J. Watson Research Center
Mathematical Sciences Dept.
P.O. Box 218
Yorktown Heights NY 10598
USA
sbar@watson.ibm.com
tel: +1-914-945-11 69

Andreas **Schulz**
TU Berlin
FB Mathematik MA 6-1
Straße des 17. Juni 136
D-10623 Berlin
Germany
schulz@math.tu-berlin.de
http://www.math.tu-berlin.de/~schulz
tel: +49-30-314-2 57 91

Uwe **Schwiegelshohn**
Universität Dortmund
Lehrstuhl für Datenverarbeitungssysteme
44221 Dortmund
uwe@carla.e-technik.uni-dortmund.de
http://www-ds.e-technik.uni-dortmund.de/
tel: +49-231-755-26 34

Sergey V. **Sevastianov**
Siberian Branch of the
Russian Academy of Sciences
Institute of Mathematics
Universitetskii pr. 4
Novosibirsk 630090
Russia
seva@math.nsc.ru
tel: +7-3832-35-6328

David **Shmoys**
Cornell University
School of Operations Research and
Industrial Engineering
232 Rhodes Hall
Ithaca NY 14853
USA
shmoys@cs.cornell.edu
tel: +1-607-255-9146

Miklos **Simonovits**
Hungarian Academy of Sciences
Mathematical Research Institute
H-1364 Budapest
Hungary
miki@math-inst.hu

Angelika **Steger**
TU München
Institut für Informatik
D-80290 München
steger@informatik.tu-muenchen.de
http://www.informatik.tu-muenchen.de/~steger
tel: +49-89-289-2 84 73

Cliff **Stein**
Dartmouth College
Dept. of Computer Science
6211 Sudikoff Laboratory
Hanover NH 03755-3551
USA
cliff@cs.dartmouth.edu
http://www.cs.dartmouth.edu/~cliff
tel: +1-603-646-2760

Leen **Stougie**
University of Amsterdam
Faculty of Economics and Econometrics
Dept. of Operation Research
Roeterstraat 11
NL-1018 WB Amsterdam
The Netherlands
leen@sara.nl
tel: +31-20-5254319

Luca **Trevisan**
Universite de Geneve
CUI
24 Rue General-Dufour
CH-1211 Geneve 4
Switzerland
trevisan@cui.unige.ch

Vijay **Vazirani**
Georgia Institute of Technology
College of Computing
801 Atlantic Drive
Atlanta GA 30332-0280
USA
vazirani@cc.gatech.edu
tel: +1-404-894-5850

Dorothea **Wagner**
Universität Konstanz
FB Mathematik u. Informatik
Postfach 55 60 D188
78457 Konstanz
Dorothea.Wagner@uni-konstanz.de
tel: +49-7531-88-2893

Jeffery **Westbrook**
AT&T Labs-Research
180 Park Avenue
Florham Park NJ 07932
USA
jeffw@research.att.com
tel: +1-973-360-8703

David **Williamson**
IBM T. J. Watson Research Center
Room 33-219
P.O. Box 218
Yorktown Heights NY 10598
USA
dpw@watson.ibm.com
http://www.research.igm.com/people/w/
williamson/
tel: +1-914-945-1743

Gerhard J. **Woeginger**
TU Graz
Institut für Mathematik
Steyrergasse 30
A-8010 Graz
Austria
gwoegi@opt.math.tu-graz.ac.at
tel: +43-316-873-53 57

Neal **Young**
Dartmouth College
Dept. of Computer Science
6188 Bradley Hall
Hanover NH 03755-3551
USA
ney@cs.dartmouth.edu

Bisher erschienene bzw. geplante Dagstuhl-Seminar-Reports:

T. Dietterich, W. Maass, H.U. Simon, Richard Sutton (editors):
Theory and Practice of Machine Learning, Dagstuhl-Seminar-Report; 163; 06.-10.01.97 (9702)

K. Honda, M. Odersky, B. Pierce, G. Smolka, P. Wadler (editors):
High-Level Concurrent Languages, Dagstuhl-Seminar-Report; 164; 20.01.-22.01.97 (97041)

P. Gritzmann, M. Nivat (editors):
Discrete Tomography: Algorithms and Complexity, Dagstuhl-Seminar-Report; 165; 20.01.-24.01.97 (97042)

B. Becker, R. Bryant, M. Fujita, Ch. Meinel (editors):
Computer Aided Design and Test Decision Diagrams - Concepts and Applications, Dagstuhl-Seminar-Report; 166; 27.01.-31.01.97 (9705)

J. Carletta, N. Dahlbäck, N. Reithinger, M.A. Walker (editors):
Standards for Dialogue Coding in Natural Language Processing, Dagstuhl-Seminar-Report; 167; 03.02.-07.02.97 (9706)

R. Klein, R. Seidel, S. Teller (editors):
Computational Geometry, Dagstuhl-Seminar-Report; 168; 10.02.-14.02.97 (9707)

M. Cole, S. Gorlatsch, Ch. Lengauer, D. Skillicorn (editors):
Theory and Practice of Higher-Order Parallel Programming, Dagstuhl-Seminar-Report; 169; 17.02.-21.02.97 (9708)

W. Bibel, U. Furbach, R. Hasegawa, M. Stickel (editors):
Deduction, Dagstuhl-Seminar-Report; 170; 24.02.-28.02.97 (9709)

S. Jähnichen, J. Loeckx, M. Wirsing (editors):
Logic for System Engineering, Dagstuhl-Seminar-Report; 171; 03.03.-07.03.97 (9710)

D. Barrington, N. Nisan, R. Reischuk, I. Wegener (editors):
Complexity of Boolean Functions, Dagstuhl-Seminar-Report; 172; 10.03.-14.03.97 (9711)

N. Barghouti, K. Dittrich, D. Maier, W. Schäfer (editors):
Software Engineering and Database Technology, Dagstuhl-Seminar-Report; 173; 17.03.-21.03.97 (9712)

H.-D. Ehrich, Y. Feng, D. Kung (editors):
Object-Oriented Software Development, Dagstuhl-Seminar-Report; 174; 07.04.-11.04.97 (9715)

N. Fuhr, K. v. Rijsbergen, A. F. Smeaton (editors):
Evaluation of Multimedia Information Retrieval, Dagstuhl-Seminar-Report; 175; 14.04.-18.04.97 (9716)

Ker-I Ko, A. Nerode, K. Weihrauch (editors):
Computability and Complexity in Analysis, Dagstuhl-Seminar-Report; 176; 21.-25.04.97 (9717)

N. Gilbert, U. Mueller, K. Troitzsch (editors):
Social Science Microsimulation: Tools for Modeling, Parameter Optimization, and Sensitivity Analysis, Dagstuhl-Seminar-Report; 177; 05.05.-09.05.97 (9719)

E. Börger, P. Joannou, D. Parnas (editors):
Practical Methods for Code Documentation and Code Inspection, Dagstuhl-Seminar-Report; 178; 12.05.-16.05.97 (9720)

E. Becker, W. Decker, D. Eisenbud, H. W. Lenstra, B. Trager (editors):
Computational Aspects of Commutative Algebra and Algebraic Geometry, Dagstuhl-Seminar-Report; 179; 26.05.-30.05.97 (9722)

J. Blazewicz, K. Ecker, W. Kubiak, D. Trystram (editors):
Scheduling in Computer & Manufacturing Systems, Dagstuhl-Seminar-Report; 180; 02.06.-06.06.97 (9723)

H. Hagen, G. Nielson, F. Post (editors):
Scientific Visualization, Dagstuhl-Seminar-Report; 181; 09.06.-13.06.97 (9724)

A. Danthine, W. Effelsberg, D. Ferrari, O. Spaniol (editors):
High-Performance Networks for Multimedia Applications, Dagstuhl-Seminar-Report; 182; 16.06.-20.06.97 (9725)

O. Etzion, S. Jajodia, S. Sripada (editors):
Temporal Databases, Dagstuhl-Seminar-Report; 183; 23.06.-27.06.97 (9726)

N. Georganas, T. Little, K. Rothermel, R. Steinmetz (editors):
Multimedia Syncronization and Resource Management in Advanced Multimedia Environments, Dagstuhl-Seminar-Report; 184; 30.06.-04.07.97 (9727)

P. Flajolet, R. Kemp, H. Mahmoud, H. Prodinger (editors):
`Average-Case'-Analysis of Algorithms, Dagstuhl-Seminar-Report; 185; 07.07.-11.07.97 (9728)

E. Coffman, E.W. Mayr, U. Schwiegelshohn (editors):
Parallel Scheduling, Dagstuhl-Seminar-Report; 186; 14.07.-18.07.97 (9729)

Y. Rabani, D. B. Shmoys, G. Woeginger (editors):
Combinatorial Approximation Algorithms, Dagstuhl-Seminar-Report; 187; 18.08.-22.08.97 (9734)

E.W. Mayr, F. Meyer a.d. Heide, L. Rudolph (editors):
Parallel and Distributed Algorithms, Dagstuhl-Seminar-Report; 188; 08.09.-12.09.97 (9737)

G. Haring, Ch. Lindemann, M. Reiser (editors):
Performance Evaluation: Origins and Directions, Dagstuhl-Seminar-Report; 189; 15.05.-19.09.97 (9738)

A. Odlyzko, C.P. Schnorr, A. Shamir, J. Stern (editors):
Cryptography, Dagstuhl-Seminar-Report; 190; 22.09.-26.09.97 (9739)

Ph. Codognet, U. Montanari, A. Podelski, V. Saraswat (editors):
Concurrent Constraint Programming, Dagstuhl-Seminar-Report; 191; 06.-10.10.97 (9741)

F. Mattern, K. Rothermel, F. Schneider (editors):
Mobile Software-Agents, Dagstuhl-Seminar-Report; 192; 13.10.-17.10.97 (9742)

H. Comon, D. Kozen, H. Seidl, M. Vardi (editors):
Applications of Tree Automata in Rewriting, Logic and Programming, Dagstuhl-Seminar-Report; 193; 20.10.-24.10.97 (9743)

D. Mundici, P. Schmitt, L. Zadeh (editors):
Multiple Valued Logic, Dagstuhl-Seminar-Report; 194; 27.10.-31.10.97 (9744)

G. Alefeld, O. Mahrenholtz, R. Vollmar (editors):
Parallel Processing in the Engineering Sciences - Methods and Applications, Dagstuhl-Seminar-Report; 195; 03.11.-07.11.97 (9745)

H. I. Christensen, D. Hogg, B. Neumann (editors):
Knowledge-Based Computer Vision, Dagstuhl-Seminar-Report; 196; 08.12.-12.12.97 (9750)

M. Jünger, G. Reinelt, H. Rieger, G. Rinaldi (editors):

Algorithmic Techniques in Physics, Dagstuhl-Seminar-Report; 197; 15.12.-19.12.97 (9751)

P. Gloor (Zürich), R. Harding (Cambridge, UK), Th. Ottmann (Freiburg) (editors):
New Media in Computer-Science Teaching at University Level; Dagstuhl-Seminar-Report; 198; 02.02.1998 - 07.02.1998 (98051)

X.T. Bui (Hawaii, USA), J.M. Carroll (Blacksburg, USA), M. Jarke (Aachen) (editors):
Szenario Management; Dagstuhl-Seminar-Report; 199; 08.02.1998 - 13.02.1998 (98061)

H.-D. Ehrich (Braunschweig), U. Goltz (Hildesheim), J. Meseguer (SRI, Menlo Park) (editors):
Information Systems as Reactive Systems; Dagstuhl-Seminar-Report; 200; 15.02.1998 - 20.02.1998 (98071)

K.H. Brenner (Mannheim), H. ElGindy (Callghan, AUS), H. Schmeck (Karlsruhe), H. Schröder (Loughborough, UK) (editors):
Dynamically Reconfigurable Architectures; Dagstuhl-Seminar-Report; 201; 22.02.1998 - 27.02.1998 (98081)

S. Näher (Halle), P. Widmayer (ETH Zürich), I. Munroe (Waterloo) (editors):
Data Structures; Dagstuhl-Seminar-Report; 202; 01.03.1998 - 06.03.1998 (98091)

H. Müller (Victoria, CDN), H. Weber (TU-Berlin) (editors):
Continous Engineering for Industrial Scale Software Systems; Dagstuhl-Seminar-Report; 203; 01.03.1998 - 06.03.1998 (98092)

H. Müller (Victoria, CDN), T. Reps (Wisconsin), G. Snelting (Braunschweig), (editors):
Program Comprehension and Software Reengineering; Dagstuhl-Seminar-Report; 204; 08.03.1998 - 13.03.1998 (98101)

R. Haralick (Seattle), R. Klette (Auckland), S. Stiehl (Hamburg), M. Viergever (Delft) (editors):
Evaluation and Validation of Computer Vision Algorithms; Dagstuhl-Seminar-Report; 205; 15.03.1998 - 20.03.1998 (98111)

U. Herzog (Erlangen), G. Latouche (Bruxelles), V. Ramaswami (Bell, Morristown), P. Tran-Gia (Würzburg) (editors):
Advanced Stochastic Modelling Applied to Telecommunication-Networks and Distributed Systems; Dagstuhl-Seminar-Report; 206; 29.03.1998 - 03.04.1998 (98131)

M. Schwartzbach (Aarhus), M. Sagiv (Tel Aviv),K. Weihe (Konstanz) (editors):
Programs with Recursively Defined Data Structures; Dagstuhl-Seminar-Report; 207; 19.04.1998 - 24.04.1998 (98161)

M. Jazayeri (Wien), R. Loos (Tübingen), D. Musser (Troy), A. Stepanov (SGI, Mountain View) (editors):
Generic Programming; Dagstuhl-Seminar-Report; 208; 26.04.1998 - 30.04.1998 (98171)

S. Brookes (CMU Pittsburgh), M. Droste (Dresden), M. Mislove (Tulane) (editors):
Domain Theory and Its Applications; Dagstuhl-Seminar-Report; 209; 03.05.1998 - 08.05.1998 (98181)

T. Beth (Karlsruhe), G. Brassard (Montreal) (editors):
Quantum Algorithms; Dagstuhl-Seminar-Report; 210; 10.05.1998 - 15.05.1998 (98191)

P. Mathe (WIAS, Berlin), K. Ritter (Erlangen), J. Traub (Columbia NY) (editors):
Algorithms and Complexity for Continuous Problems; Dagstuhl-Seminar-Report; 211; 17.05.1998 - 22.05.1998 (98201)

M. Gross (ETH Zürich), H. Müller (Dortmund), P. Schröder (CalTech), H.-P. Seidel (Erlangen) (editors):
Hierarchial Methods in Computer Graphics; Dagstuhl-Seminar-Report; 212; 24.05.1998 - 29.05.1998 (98211)