D I M A C S

Center for Discrete Mathematics and Theoretical Computer Science

DIMACS SPECIAL YEAR ON LARGE SCALE DISCRETE OPTIMIZATION: 1999 ANNUAL REPORT

DIMACS Center • Rutgers University • 96 Frelinghuysen Road • Piscataway, NJ 08854-8018 • USA
TEL: 732-445-5928 • FAX: 732-445-5932 • EMAIL: center@dimacs.rutgers.edu
WWW: http://dimacs.rutgers.edu/ • TELNET: telnet info.rutgers.edu 90

DIMACS is a partnership of Rutgers University, Princeton University, AT&T Labs - Research, Bell Labs, NEC Research, and Telcordia Technologies (formerly Bellcore)

Overview:

The Special Year on Large Scale Discrete Optimization began in September 1998 and continues through much of 1999-00 at a considerably lower level than in 1998-99. Some pre-special year activities were held during 1997-98. The year builds on the 1992-93 Special Year on Combinatorial Optimization.

Organizers:

- Endre Boros, Rutgers University (local chair)
- William Pulleyblank, IBM
- Michael Trick, Carnegie Mellon University (general chair)
- Bob Vanderbei, Princeton University
- Vijay Vazirani, Georgia Tech

Sponsors and Collaborating Institutions:

The National Science Foundation
The New Jersey Commission on Science and Technology
AT&T Labs - Research
Bell Labs
Carnegie Mellon University
The Electric Power Research Institute (EPRI)
Georgia Institute of Technology
IBM Research
Institute for Advanced Study, Princeton
NEC Research Institute
Princeton University
Rutgers University
Telcordia Technologies

Theme of Program:

This special year is motivated by developments in the past ten years. There have been three simultaneous advances: at the algorithmic level, with new and interesting algorithms for solving large scale discrete problems both exactly and approximately; at the implementation level, providing codes and systems for solving such problems; and at the practical level, with interest in the fast solution of real world problems. Examples of such applications include airline crew scheduling, chemical process design, and telecommunication network design. Our goal in this year is to unite groups working at these three levels. The confluence of exciting recent work on approximation and other algorithms and greatly improved software/hardware for optimization makes large scale discrete optimization a practical, useful approach to solving problems of interest in applications.

The special year organizers have provided forums to increase research contacts between groups that do not frequently come in contact - for example practitioners and academics, operations researchers and computer science researchers, theoretically oriented researchers (in areas such as interior point IP methods, combinatorial geometry approaches, approximation methods, and group formulations) and computationally oriented researchers, including those interested in AI techniques such as constraint programming.

Workshops and Special Events February 1999 - January 2000

- **DIMACS Workshop on Large Scale Discrete Optimization in Logistics**
  - Dates: February 8-10, 1999
  - Location: Rutgers University
  - Attendance: 79
  - Organizers: George Nemhauser and Martin Savelsbergh, Georgia Tech
  - Comment: Joint with Special Year on Networks

  An efficient and effective logistics system is the key to being competitive in today's global marketplace. This realization has fueled an enormous interest in new technology for designing and operating logistics systems. Discrete optimization frequently is an integral part of such technology. In particular, in the last decade, the use of large-scale discrete optimization models in transportation and production planning has increased dramatically. This has been made possible by developments in modeling, algorithms, software, and hardware. This workshop brought together leading researchers to present state-of-the-art research activities and to discuss innovative future directions of research for this very active field with significant practical importance. The talks covered a wide range of topics including modeling, algorithms, and applications. Application domains that were covered include network design, production planning, vehicle routing, and crew scheduling. The workshop opened with a talk by William Cook titled "On the Solution of Traveling Salesman Problems". Research on the solution of the traveling salesman problem, the prototypical discrete optimization problem, has stimulated enormous interest in discrete optimization and has helped significantly in the growth of the field. Several other talks focused on methodology, such as "Approximation Algorithms via Linear Programming: Rounding Algorithms" by David Shmoys, "An Overview of Classical and Modern Heuristics for the Vehicle Routing Problem" by Gilbert Laporte, "GRASP: Greedy Randomized Adaptive Search Procedures" by Mauricio Resende, "AI Techniques for Large-Scale Discrete Optimization" by Jimmy Crawford, and "Very Large Scale Neighborhood Search" by Ravi Ahuja. Furthermore, there were several talks related to the main application domains of discrete optimization. In the topic of Airline Optimization, we had "Solving Aircraft Maintenance Scheduling Problems" by Daniel Bienstock, "Optimization Problems in Air Traffic Management" by Mike Ball, and "Airline Integrated Recovery: AIR and SimAIR" by Ellis Johnson. In the topic Manufacturing and Production Planning, we had "BC-PROD and LOTSIZELIB" by Laurence Wolsey, "Optimization in IBM Manufacturing" by Brenda Dietrich, and "Painting Cars and the Traveling Salesman Problem" by Tom Magnanti. In the topic Supply-Chain Management, there was "Optimal and Approximation Algorithms for Large Scale Supply Chain Management Problems" by David Simchi-Levi and "Dynamic Programming Approximations for the Discrete Multi-commodity Network Flow Problem over a Dynamic Graph" by Warren Powell. Finally, in the topic Telecommunication Network Design there was "The Prize Collecting Steiner Tree Problem" by David Johnson. There was also a full group discussion on future directions for the field. The group was very encouraged by the increasing use of discrete optimization methods in industry and the need for the solution of larger more complex models that, for example, include stochastic elements. In conclusion, the workshop provided an excellent environment in which leading researchers in the field could share their ideas not only with each other, but with junior researchers and students as well.

- **DIMACS Workshop on Large Scale Discrete Optimization in Robotics and Vision**
  - Dates: March 21-23, 1999
  - Location: Rutgers University
  - Attendance: 72
  - Organizers: Jitendra Malik, University of California - Berkeley, and Jean Ponce, University of Illinois - Urbana Champaign

The theme of the workshop was to address the fundamental mathematical and computational problems underlying computer vision and robotics. We were able to attract many of the major researchers in the field to help give a perspective on the leading issues today and the likely directions for future progress. Points of contact with computer science theory and algorithms, probabilistic inference, and other areas of applied mathematics in general were stressed. Biological vision is used to guide a number of different aspects of behavior–guiding locomotion, controlling manipulation and recognition of objects. Computer vision seeks to develop models and algorithms for these competences, drawing heavily on geometry, probability and statistics and computation. Likewise, robotics is concerned with controlling mechanisms such as manipulator arms, grippers and mobile platforms using sensory input as well as a priori knowledge about the geometry and physics of these mechanisms and their environment. This discipline is at the boundary between fields as diverse as mathematics, mechanics, control theory and computer science. The morning and early afternoon of the first day was devoted to the modeling and inference challenges in vision. David Mumford articulated the case for probabilistic modeling in a Bayesian framework as the natural way to incorporate both low level and high level knowledge for solving vision problems. He argued for particle filtering as a convenient formalism for representing and computing with probability distributions. David Forsyth demonstrated the power of sampling-based methods with algorithms for three different computer vision problems–color constancy, structure from motion, and finding people in images. Jitendra Malik and Ramin Zabih presented solutions to problems of image segmentation and stereoscopic correspondence in graph theoretic formulations. It turns out that spectral graph theory and various minimum cut formulations, which have been much studied by the theoretical computer science community, can provide practical and efficient algorithms for these problems. The morning session of the second day was again devoted to vision, primarily to object recognition. Sven Dickinson presented a theory for the generic representation of 2-D shape, in which structural descriptions are derived from the shocks (singularities) of a curve evolution process acting on a silhouette. David Kriegman described recognition algorithms that are based on characterizing the set of images of an object in fixed pose but under a range of lighting conditions, and demonstrated their performance in face recognition tasks. Martial Hebert presented a new approach to object recognition in range data that deals particularly well with the clutter present in complex industrial environments. The afternoons of both the first and second day had a series of talks on robotics. Jean-Paul Laumond gave an overview of classical configuration space concepts and modern approaches to motion planning, then focused on his recent work using probabilistic roadmaps methods to solve a variety of hard robotics problems such as obstacle avoidance for manipulators and navigation for non-holonomic mobile robots. Bud Mishra surveyed the field of grasp planning, presenting its geometric and algorithmic foundations before concentrating on his recent work aimed at using simple devices in efficient algorithms for grasping applications. Greg Hager investigated new theoretical aspects of visual servoing. In particular, he provided some answers to the question: "When is it possible to decide, on the basis of the information contained in the images of an imprecisely modeled two-camera vision system, whether or not a prescribed robot positioning task has been precisely accomplished?" There were four talks on the concluding day of the conference. Jean Ponce and Mark Overmars led off with the application of computational geometry to design of fixtures for part manipulation. Mark Overmars addressed the problem of constructing efficient algorithms for automatically designing conveyor part sorting systems using belts equipped with fences and vibratory bowl feeders. Jean Ponce concentrated on a new approach to manipulation planning based on the concept of capture regions, and illustrated it with experiments using a reconfigurable gripper and a part feeding mechanism. Andrew Zisserman gave a throrough survey of the state of the art in structure from motion–both the theory of the problem and some impressive reconstruction results from video sequences. The concluding talk of the workshop was given by Jan Koenderink, who has been the pioneer in the study of shape and the application of differential geometry and singularity theory to vision. In this talk he presented his experimental results on the perception of pictorial relief. In addition to the series of talks, we had two panel discussions, one on the first day "Can Probabilistic Inference in Vision be Made Computationally Tractable" and one on the second day "Vision and Robotics: Growing Apart or Coming Together?" Both were lively affairs with much participation from the audience. We also had a submitted poster session with a dozen participants. This gave an opportunity for various graduate student and postdoctoral fellows to present their research and get feedback.

- **Mini-Workshop on Mathematical and Computational Approaches to Elections**
  - Date: June 18, 1999
  - Location: Rutgers University
  - Attendance: 38
  - Organizer: Bruno Simeone, La Sapienza University, Rome, Italy

  In recent years, the debate on electoral reform has been very lively both in newborn and old democracies. Although the choice of an appropriate electoral system is affected by political and social factors, and as such it certainly belongs to the domain of social sciences, it is also evident that the formal aspects and procedures of voting naturally call for the use of quantitative methods for the analysis and design of such systems. Several quantitative approaches have been developed, such as axiomatics (in the more general context of social choice theory), game-theory, geometry, statistics, optimization, and discrete algorithms. The workshop aimed to discuss and compare these approaches and to pinpoint or foresee connections between them. Since electoral systems obviously involve only finite (although possibly quite large) sets and, on the other hand, they essentially provide procedures to transform votes into seats and to assign seats to elected candidates, combinatorial and algorithmic aspects are of paramount importance in their study. All of the speakers were well recognized for their individual contributions to voting systems, and altogether they indeed offered a comprehensive and broad view of this area. Fred Roberts opened the Workshop with a thorough discussion on the axiomatics of the mean and the median, two widespread social choice functions. He presented many elegant combinatorial results concerning these two functions in the areas of location theory, graph theory, and molecular biology. Steven Brams's talk, "Voting Models: Explaining Practice and Recommending Reform", dealt with three topics: resource-allocation models in US presidential campaigns, the establishment of approval voting procedures, and a new "paradox of multiple elections" occurring in referendums with multiple questions. His focus was on the implementation and validation of electoral procedures. Peter Fishburn talked about "Electoral Axiomatics and Combinatorics", demonstrating quite convincingly that elections and their axiomatics provide a rich source of intriguing combinatorial problems.

Michel Regenwetter, in his talk on "Probabilistic Models of Social Choice and Preference", presented some preference and choice models blending probabilistic tools with ordinal utility theory. Here too significant combinatorial results are obtained. Donald Saari's keen and stimulating presentation "From Arrow's and Sen's Theorems to Resolving all Vote Paradoxes" offered some fresh insights on classical Impossibility Theorems. A careful reading of these results suggests one possible way out of related paradoxes. Bruno Simeone, in his talk "Discrete Optimization Models for Electoral Systems", pointed out that discrete optimization, on the one hand, provides a broad conceptual framework for electoral procedures; on the other hand, it yields a powerful tool-kit for designing different components of an electoral system. He illustrated these two points showing that all proportional formulas are obtainable via the minimization of suitable Schur-convex functions, or generalizations thereof; and describing a multi-criteria graph partitioning model for political districting. The last talk, Michael Trick's "Computational Complexity and Algorithmic Aspects in Voting Theory", focused on computational complexity issues in determining winners, manipulating elections, and strategic voting. The participants found the workshop interesting and stimulating. There was general agreement that the piece de resistence of the workshop was the excellent mix of different approaches.

- **Discrete Optimization '99 (DO '99)**
  - Dates: July 25-30, 1999
  - Location: Rutgers University
  - Attendance: 102
  - Organizers: Peter Hammer and Endre Boros, Rutgers University
  - Co-Sponsor: RUTCOR (Rutgers Center for Operations Research)

Discrete Optimization underwent dramatic development in the last half century, with particularly spectacular growth in the last few decades. The main goal of this workshop was to survey the state of the art in D.O. There were a series of expository lectures presenting the major subareas of the field, including its theoretical foundations, its methodology and applications, with the surveys presented by some of the most prominent researchers in the field. There was also a forum for the presentation of new developments in D.O. through a series of contributed papers sessions. D.O.'99 was held 22 years after D.O.'77, which had very similar goals, and whose collection of surveys (*Discrete Optimization I and II*, Annals of Discrete Mathematics, vols. 4 and 5, ed. by P.L. Hammer, E.L. Johnson, and B.H. Korte, North-Holland, Amsterdam, 1979) is still frequently used. On the first day, Tom Magnanti presented a survey on network design, Egon Balas one on "Lift and Project," and Paolo Toth one on vehicle routing. Surveys on the second day were on complex scheduling problems by Peter Brucker, graph connectivity by Toshihide Ibaraki, greedy algorithms in linear programming problems by Alan Hoffman, bin packing problems in two and three dimensions by Silvano Martello, and ideal clutters by Gerard Corneujols. Day three featured surveys by David Williamson on approximation algorithms, Dorit Hochbaum on "A Unified Technique for Approximation Algorithms for Integer Programs with Three Variables per Inequality", Andras Recski on combinatorics of grid-like frameworks, Jim Orlin on neighborhood search, and Fred Glover on Tabu search and evolutionary developments. On Day 4, Henry Wolkowicz presented a survey on semidefinite programming for D.O., Yves Crama one on optimization models in production planning, and Rainer Burkard one on assignment problems. The fifth day saw surveys by Bill Pulleyblank on "Hilbert Bases, Caratheodory's Theorem and Integer

Programming," by Laurence Wolsey surveying mixed integer programming, Vasek Chvatal on the solution of the traveling salesman problem, Bernhard Korte on the subject "Gigahertz-processors Need Discrete Optimization," and Jakob Krarup on locational decisions. The final two surveys were given on day 6 by Endre Boros on pseudo Boolean optimization and Peter Hammer on Boolean optimization. Jakob Krarup led a lively roundtable discussion on perspectives and trends in D.O. The organization of the workshop was made possible by combining funding through a DIMACS allocation and an NSF grant to the Rutgers Center for Operations Research, the co-sponsor of the program, for the support of young participants. Arrangements have been made for the publication of two special issues of *Discrete Applied Mathematics*, one dedicated to the survey lectures, and the other to the contributed papers presented in the workshop. Moreover, the volume of survey lectures will be reprinted in book format and made available to the public at large.

- **DIMACS - IBM Conference on Theory and Practice of Integer Programming in Honor of Ralph Gomory on the Occasion of his 70th Birthday**

  - Dates: August 2-4, 1999
  - Location: IBM Watson Labs, Yorktown Heights, NY
  - Attendance: 120
  - Organizers: Bill Pulleyblank, IBM Watson Research Center, and William Cook, Rice University
  - Co-sponsor: IBM Research

  Integer Programming is a discipline which underlies much of modern Operations Research, including scheduling, logistics, resource allocation and routing problems. It is also the subject of much theoretical research, combining the study of linear convex polyhedra with integer lattices and group theory. Because of its importance to applications, it is also the focus of significant amounts of computational research, leading to software which is much more powerful than what was available a few years ago, but which still is unable to deal with real world problems of the required size and complexity. Recently we have seen significant advances in linear programming, an underlying problem, as well as the emergence of totally new methods, based on geometric principles. In addition, the broad availability of high performance scientific workstations is now making it possible for a much broader community to take advantage of integer programming software. The workshop brought together a who-is-who collection of international experts covering all aspects of this field. It provided a unique opportunity for a focused study of these advances and for exploration of the synergy possible by combining the various subfields. The attendance each day exceeded one hundred people, including a substantial number of graduate students and young researchers. The feedback we received from the participants was extremely positive. One of the senior participants wrote: "It was the best workshop I ever attended and the atmosphere was terrific." Another wrote: "It felt like a watershed event didn't it?" A lecturer in the final session stated that it was the first time in many years that he attended every lecture of a workshop. Vasek Chvatal opened the conference by talking about Gomory cuts and combinatorial reasoning. Herbert Scarf described homotopy methods for computing test sets. A featured talk by Egon Balas revisited disjunctive programming. George Nemhauser described large-scale integer programming in airline scheduling. Balanced matrices was the subject of the featured talk by Gerard Cornuejols. These and 26 other lectures covered a good mix of theory, computation, and applications. Theoretical topics included homotopy methods, the geometry of numbers, commutative algebra, Hilbert bases, and the structure of 0-1 polytopes. Computational lectures covered large-scale tests of cutting planes, tests of relaxation methods for linear programming, and branch-and-cut algorithms to solve combinatorial problems. The applications discussed included airline crew scheduling, the design of telecommunication networks, sports scheduling, and the design of computer processors. One of the important themes that emerged in the lectures is the remarkable performance of modern implementations of Gomory's mixed integer cutting plane method. This created a natural highlight to the meeting in Ralph Gomory's lecture on the corner polyhedron and its relationship to cutting planes. The meeting included a dinner in honor of Dr. Gomory. Speeches by Alan Hoffman, T.C. Hu, Ellis Johnson, William Baumol, Herb Scarf, and Ralph Gomory made this a memorable event.

- **Diagonal Matrix Scaling and its Generalizations and Their Applications in Convex Programming Over Cones**
  - Dates: August 25 - 27, 1999
  - Location: Rutgers University
  - Attendance: 14
  - Organizers: Bahman Kalantari, Rutgers University, Uriel Rothblum, Technion-Israel Institute of Technology, and Alex Samorodnitsky, DIMACS, Rutgers University

The diagonal matrix scaling problem and its generalizations form a fundamental set of problems in mathematical programming and computer science. On the one hand, for nonnegative matrices the doubly stochastic scaling is a fundamental problem in itself with diverse applications, e.g., in estimation of transition probability in Markov chains, network optimization, planning of traffic and transportation, matrix preconditioning, and image reconstruction. On the other hand, the doubly quasi-stochastic scaling of positive semidefinite matrices and its generalizations has tremendous theoretical and algorithmic applications in linear programming, semidefinite programming, and more generally self-concordant conic convex programming. Indeed diagonal matrix scaling has played a key role in Karmarkar's famous LP algorithm. The matrix scaling problem also finds application in combinatorial problems as well as in the approximation of NP-complete problems. The aim of this workshop was to bring together researchers working on theoretical and practical aspects of matrix scaling problems (including matrix balancing), and their generalizations. It was the hope of the organizers that the workshop would give convincing evidence of the broad significance of the matrix scaling problems to researchers within this area, as well as those interested in convex programming, self-concordance theory, duality theory, complexity of algorithms, combinatorial optimization, homogeneous inequalities, numerical analysis, etc. Moreover, the workshop was aimed at defining new research topics and describing open problems. The emphasis of the talks on the first day of the workshop was on topics related to the structural properties of problems pertaining to diagonal scaling. The emphasis of talks on the second day was on algorithms and applications in other fields such as convex programming. Lev Bregman spoke on generalized matrix scaling, combinatorial games, and other applications. Uriel Rothblum dealt with existence and computation in matrix scaling, with emphasis on complexity analysis of methods that compute approximate solutions. Michael Neumann talked about spectral radii of fixed Frobenius norm perturbations of nonnegative matrices. A talk about minimization of spectral radius and norm of a nonnegative matrix diagonal equivalence scaling to a multiple of a doubly stochastic matrix was given by Hans Schneider. Bit-Shun Tam spoke on matrices with cyclic structure. His talk was on structure of nonnegative matrices and generalizations. Alexander Barvinok talked about randomized polynomial time algorithms to approximate the permanent within a simply exponential factor. Leonid Gurvits presented the topic "Operator Matching, Scaling, and Sinkhorn's Iteration," dealing with various convergence properties of the operator Sinkhorn iteration and its applications, such as approximations of mixed determinants, non-convex optimization and generalized Alexandrov-Fenchel inequalities. He also introduced a certain operator/quantum analog of the usual permanents. Alex Samorodnitsky spoke on scaling in approximation algorithms and described two applications of scaling - in approximation of the permanent and in approximation of the mixed discriminant. Shmuel Friedland spoke on computing the Hausdorrf dimensions of subshifts using matrices. Bahman Kalantari concluded the meeting by speaking about matrix scaling over convex cones and dualities and their applications in mathematical programming.

- **Next Generation of Unit Commitment Models**
  - Dates: September 27 - 28, 1999
  - Location: Rutgers University
  - Attendance: 66
  - Organizers: Mike Rothkopf, Rutgers University, Benjamin F. Hobbs, Johns Hopkins University, and Hung-Po Chao, Electric Power Research Institute
  - Co-sponsor: Electric Power Research Institute (EPRI)

For many years, the electric power industry has been using a variety of technical methods to help them solve the unit commitment problem. The result has been savings of tens and perhaps hundreds of millions of dollars in fuel costs. However, things are changing. Optimization technology is improving and the industry is undergoing radical restructuring. Consequently, the role of unit commitment models is changing and the value of the improved solutions that better algorithms might yield is increasing. The purpose of this conference was to explore the technology and needs of the next generation of computer models for aiding unit commitment decisions. The unit commitment problem can be defined as the scheduling of production from power generation units over a daily to weekly time horizon in order to accomplish some objective. The problem solution must respect both generator constraints (such as ramp rate limits and minimum up or down times) and system constraints (reserve and energy requirements and, potentially, transmission constraints). The objective function should account for costs associated with energy production and start-up and shut-down decisions, along with possible revenue or customer impacts of those decisions. The resulting problem is a large scale nonlinear integer program. Because of the problem's size and complexity, and because of the large economic benefits that could result from its improved solution, considerable attention has been devoted to algorithm development. Heuristics such as "priority lists" have long been used by industry; but in the last three decades, more systematic procedures based on a variety of algorithms have been proposed and tested. These techniques have included dynamic programming, branch-and-bound mixed integer programming, linear and network programming approaches, and Benders decomposition methods, among others. Recently, metaheuristic methods have been tested, such as genetic programming and simulated annealing, along with expert systems and neural networks. However, the solution approach that has been most successful, and which is most widely used at present, is Lagrangian relaxation. This procedure decomposes the problem by multiplying constraints that couple different generators (such as energy demand and reserve constraints) by Lagrangian multipliers and placing them in the objective function. Given a set of multiplier values, the problem is then separable in the generating units, and a dynamic program of low dimension can be used to obtain a trial schedule for each unit. A process of multiplier adjustment is used to search for feasible near-optimal solutions. Lagrangian relaxation has proven useful for quick development of good, if not optimal, generator schedules. However, recent improvements in integer programming codes and other algorithms suggest that it may be possible to find better solutions more rapidly. Meanwhile, restructuring has sharpened the appetite of generation owners for more efficient operation. In the past, utilities have been regulated on a cost-plus basis, and so may not have put as much priority on squeezing out the last few percent improvements in the objective function. Furthermore, system operators realized that cost functions were approximate and so were perhaps more likely to be satisfied with good solutions that were technically suboptimal. Now, with restructuring, we have schedule coordinators making commitment decisions in a market environment, and ISOs dealing with bids. Bids are precise, and small improvements in solutions can result in significant changes in payments to bidders. Further, the fact that optimization models are, in some cases, being used to determine which generators will be operated and thus paid implies that there is a greater incentive to get exact answers to make the bidding process legitimate (and to disempower the bid taker). In other words, it's a whole new world, and a reconsideration of how unit commitment models are solved and the purposes for which they are used is needed. The goal of this workshop was to bring together those who know the problem and what is needed with those who know what improvements in algorithms are really possible. The workshop started with a series of in-depth presentations by top experts in algorithms and utility systems operations. Then working groups convened to define the new roles that unit commitment models may play and promising directions in algorithm development and application. Richard O'Neill, Director of the Office of Economic Policy of the Federal Energy Regulatory Commission, opened the workshop with a talk entitled "Why are We Here?" Andrew Ott from the Pennsylvania-NewJersey-Maryland ISO presented the point of view of an independent systems operator, Fred Lee, Art Breipohl, and Walter Hobbs spoke about unit commitment in a competitive environment, David Sun from ALSTOM ESCA spoke about current experience in solving unit commitment problems, and Joseph Graves from PHB Hagler Bailly spoke about simultaneous clearing and optimization of energy and ancillary service markets in the unit commitment context. There were presentations by Ross Baldick on alternative formulations, multiplicity of solutions, and unpriced products, by Robert Bixby from Rice and ILOG entitled "MIP: Closing the Gap between Theory and Practice," and Sebastian Ceria from Columbia and Dash Optimization, Inc. on new developments in integer programming. Panels discussed the definition of the problem to be solved and what we can expect from integer programming. Parallel sessions allowed many contributed talks. One highlight was the demonstration during the workshop by Bob Bixby of CPLEX of the capabilities of new integer programming codes to solve problems that were previously thought impossible to solve to optimality. A new theoretical result on the relation between solutions to the unit commitment model and the problem of deriving prices in daily power auctions was a direct result of the workshop and is described below. The workshop was partly supported by the Electric Power Research Institute (EPRI). Among the products of the workshop will be an EPRI report summarizing the deliberations and presenting the conclusions of the workshop and a book to be published by Kluwer or IEEE Press, with EPRI agreeing to purchase up to 1000 copies of the book for widespread distribution. The intent of the workshop sponsors is that this workshop will ultimately result in better algorithms yielding more value to their users.

Workshops and Miniworkshops Planned for the Rest of 1999-00

No workshops are planned for the rest of this academic year. However, the "Algorithm Implementation Challenge" associated with this special year, on semidefinite programming, will reach its culmination in a workshop in September:

- **Workshop: Computational Challenge on Semidefinite Programming**
  - Dates: September 13 - 15, 2000
  - Location: Rutgers University
  - Organizers: David Johnson, AT&T Labs - Research, Farid Alizadeh, Rutgers University, and Gabor Pataki, Columbia University

Special Year Tutorial Series
Tutorials February 1999 - January 2000

Metaheuristics
Fred Glover, Jim Kelly, Manuel Laguna,
University of Colorado, Boulder
April 29 and 30, 1999

An Introduction to Metaheuristics placed a primary focus on Tabu Search, from the standpoint of adaptive memory methods, and on Scatter Search and Path Relinking, from the standpoint of evolutionary methods. Comparisons and contrasts with Genetic Algorithms and other approaches were included. An Introduction to OptQuest highlighted demonstrations of real world applications, especially those applications that involve significant elements of uncertainty and nonlinearity, including discreteness.

Optimal Connected Partitions of Graphs
Bruno Simeone, University of Rome, La Sapienza, Italy
April 9, April 15, April 21, April 28 and May 5, 1999

This was a series of five lectures on the topics: Path equipartition: trimming the shortest path network; greedy algorithms, semimodular lattices, and supermodular functions; shifting algorithms for tree partitioning; how to make a polynomial movie out of pseudo-polynomially many frames: continuous shifting algorithms; and connected k-partitions in k-connected graphs: a topological approach.

Discrete Optimization in Telecommunications and VLSI Design
Michel Minoux, Universitè Paris 6, France
May 7 and May 13, 1999

This was a series of two lectures on the topics: generalized graph coloring and the frequency assignment problem; and some combinatorial problems in optimal telecommunication network design. The first lecture investigated generalized graph coloring problems arising in the context of operating large radio link telecommunication networks, or cellular radio communication networks. The second talk concentrated on minimum cost multicommodity flows as basic models frequently arising in the context of topological network optimization, optimal network design, and dimensioning.

Local Search Algorithms in Discrete Optimization
Toshihide Ibaraki, Kyoto University, Japan
June 1, 1999

Local search is a well-known basic tool for constructing heuristic algorithms for discrete optimization problems. It repeats the operation of replacing the current solution with a better solution in its neighborhood (if there is any) until no better solution can be found in the neighborhood (i.e., local optimality is attained). This lecture discussed generalizations, improvements, and ramifications of local search, including effect of neighborhood sizes, adaptive usage of neighborhoods of various sizes, and combination with metaheuristics.

Dynamic Load Balancing on Heterogenous Networks
Ewald Speckenmeyer, University of Kőln, Germany
June 10, 1999

This lecture presented a dynamic load balancing algorithm for a network of heterogeneous computers. The approach was based on a fast distributed precomputation step calculating the necessary load movements. When dealing with slow network connections or incomplete load information often a perfect balanced load is not necessary. Therefore the allowed degree of imbalance is adjustable. A detailed average case analysis was carried out concerning the total load transfer and the number of transfer rounds. Its performance was empirically demonstrated to solve the satisfiability problem on a heterogeneous network of workstations.

Selected Research Results

Connection between the Unit Commitment Problem and Daily Power Auctions The unit commitment problem was described above in the discussion of the workshop on Next Generation of Unit Commitment Models. As a direct result of that workshop, Michael Rothkopf (PM, RU), Ben Hobbs (V), Bill Stewart (V), and Richard O'Neill (V) demonstrated theoretically that the relaxed LP with cutting planes that solves certain forms of the unit commitment problem has a dual interpretation useful for deriving prices in daily power auctions. In particular, the mechanisms used by the New York and PJM system operators to "make whole" producers who don't recover their costs were shown to be closely related to the dual variables for cutting planes that force those producers into the optimal solution.

A Combinatorial Algorithm for Pfaffians The Pfaffian of an oriented graph was introduced by Kastelyn in 1967 to count the number of dimer coverings of a lattice graph. It is closely linked to Perfect Matching and one of the motivations for studying Pfaffians is to understand the complexity of Perfect Matching. Perfect matchings have many important applications in assignment problems, chemistry, etc. The Pfaffian is also naturally related to the determinant of an appropriately defined matrix. This relation between Pfaffian and determinant is usually exploited to give a fast algorithm for computing Pfaffians. Meena Mahajan (V), P. R. Subramanya and V. Vinay (V) found the first completely combinatorial algorithm for computing the Pfaffian in polynomial time. The algorithm works over arbitrary commutative rings. Over integers, it can be computed in the complexity class GapL, the class of all functions that are logspace reducible to computing the integer determinant of a matrix; this result was not known before. As a corollary, it follows that under reasonable encodings of a planar graph, Kasteleyn's algorithm for counting the number of perfect matchings in a planar graph is also in GapL. The combinatorial characterization of Pfaffian also makes it possible to directly establish several algorithmic and complexity theoretic results on Perfect Matching which otherwise use determinants in a roundabout way.

Unique Maximum Matching Algorithms Harold Gabow, Haim Kaplan (PD), and Robert Tarjan (PM, PU) considered the problem of testing the uniqueness of maximum matchings, both in the unweighted and in the weighted case. For the unweighted case, they obtained two results. First, given a graph with n vertices and m edges, they found a way to test whether the graph has a unique perfect matching, and find it if it exists, in $O(m \log^4 n)$ time. This algorithm uses a recent dynamic connectivity algorithm and an old result of Kotzig characterizing unique perfect matchings in terms of bridges. For the special case of planar graphs, they improved the algorithm to run in $O(n \log n)$ time. Second, given one perfect matching, they found a way to test for the existence of another in linear time. This algorithm is a modification of Edmonds' blossom-shrinking algorithm implemented using depth-first search. A generalization of Kotzig's theorem proved by Jackson and Whitty allowed them to give a modification of the first algorithm that tests whether a given graph has a unique f-factor, and find it if it exists. They also showed how to modify the second algorithm to check whether a given f-factor is unique. Both extensions have the same time bounds as their perfect matching counterparts. For the weighted case, they were able to test in linear time whether a maximum-weight matching is unique, given the output from Edmonds' algorithm for computing such a matching. The method is an extension of their algorithm for the unweighted case.

Improvements in Throughput Maximization for Real-Time Scheduling Piotr Berman and Bhaskar DasGupta (V) considered the problem of off-line throughput maximization for job scheduling on one or more machines, where each job has a release time, a deadline and a profit. They provided algorithms that do not use linear programming, are simple and much faster than the corresponding ones proposed earlier in the literature, while either having the same quality of approximation or improving it. More precisely, compared to the known results, their pseudo-polynomial algorithm for multiple unrelated machines and all of their strongly-polynomial algorithms have better performance ratios, all of their algorithms run much faster, are combinatorial in nature and avoid linear programming. Finally, they showed that algorithms with better performance ratios than 2 are possible if the stretch factors of the jobs are bounded.

An Artificial Intelligence Approach to Forecasting of the Event Rate for an Aircraft Operator The Event Rate is an aviation safety metric that records the number of occurrences, incidents, and accidents (or "mishaps") associated with the operations of an aircraft. The calculated value is the total number of these events for a one-month period divided by the aircraft's fleet size for that same period. The value is adjusted by the fleet size to consider that an operator's exposure to events depends on the number of aircraft in its fleet. Atul Deshmukh (GS, RU) and James Luxhoj (RU) investigated the use of an Artificial Intelligence technique, such as an artificial neural network, to develop a forecasting model for the Event Rate for a major aircraft operator. Inputs to the neural network included data on airworthiness (i.e., maintenance-related) and operations surveillance inspections as well as financial data, obtained through collaboration with a project sponsored by the Federal Aviation Administration. A variety of different learning architectures, such as the backpropagation, general regression neural network, etc. were evaluated across a variety of learning parameters. For the aircraft operator in this study, the preliminary conclusion was that the Multiple Regression (MR) models for Event Rate forecasting are sometimes more effective than the Neural Networks (NN) models.

DIMACS special foci cross-fertilize each other. There was considerable cross-fertilization between our Massive Data Sets Special Year and our Large Scale Discrete Optimization Special Year. The following two results are examples of such cross-fertilization.

Generating Transversals Generating transversals to a hypergraph (or the dual of a Boolean function) is a well-known problem, which became very important recently due to its data-mining applications ("association rules" or more specifically "frequent sets" are transversals to a certain hypergraph defined by the input data set.) This problem is also interesting since its true complexity (in the incremental sense) is not known. Due to an incrementally quasi-polynomial algorithm (Fredman and Khachiyan, 1996) it is very unlikely to be an NP-hard problem, but as of today we do not know any polynomial algorithm, either – it could even be the core to a separate new class, laying somewhere between P and NP-hard (unless of course these two coincide). Endre Boros (PM, RU), Vladimir Gurvich (V), Leo Khachiyan (PM, RU), and Kazuhisa Makino (V) showed that several open versions of transversal generation are in fact equivalent to Boolean dualization, and hence treatable with an incrementally quasi-polynomial algorithm. Their results are based on several new inequalities relating the number of maximal false points to the number of minimal true points of a Boolean function. For instance, they proved that the number of maximal infeasible binary vectors to a monotone system of linear inequalities is polynomially (almost linearly) bounded in the number of minimal feasible vectors. Based on this result, they showed how to generate all minimal feasible solutions to a monotone system of inequalities in incremental quasi-polynomial time.

Attribute Selection in Learning and Data-Mining Problems In the logical analysis of data, one of the first and most important steps is the selection of a small subset of binary attributes to a given input data set which will be used in later steps of the learning process. Such a reduction step, if done carefully, is useful to eliminate noise and to reduce dimensions, making possible the application of more demanding computational methods. Recently similar considerations became popular in the learning community – so called "attribute efficient learning." Endre Boros (PM, RU) and Toshihide Ibaraki (V) considered such an attribute selection problem for the case of binary data, and based on various measures of the "separating power" of the individual attributes and their small subsets. They showed that a wide variety of such models are in fact linearly constrained maximization problems with submodular objective functions, and hence efficient approximation algorithms with a guaranteed performance rate can be applied. They considered continuous relaxations of these problems, which turn out to be smooth convex minimizations over a convex set, and hence can also be solved efficiently.

Inapproximability Results for Bounded Occurrence Satisfiability Problems Piotr Berman and Marek Karpinski (V) obtained a number of improved inapproximability results, including the best up to date explicit approximation thresholds for bounded occurrence satisfiability problems like MAX-2SAT and E2-LIN-2, and the bounded degree graph problems, like MIS, NodeCover, and MAX CUT. They also proved for the first time the inapproximability of the problem of sorting by reversals and displayed an explicit approximation threshold for this problem.

Sequential Testing of Series-Parallel Systems of Small Depth Endre Boros (PM, RU) and Tonguc Unluyurt (GS, RU) considered the problem of testing sequentially the components of a multi-component system when the testing of each component is costly. They considered a polynomial time testing policy for series-parallel systems proposed in the literature. Generalizing earlier results, they proved that the considered policy is cost-minimal in the average case sense for two sub-families of series-parallel systems. They also demonstrated via examples that neither this algorithm nor some of its improved versions are optimal for general series-parallel systems, disproving some published claims.

Randomized Complexity of Linear Arrangements and Polyhedra Marek Karpinski (V) found new applications of recent results on the complexity of recognizing n-dimensional linear arrangements and convex polyhedra by randomized algebraic decision trees. In particular, he derived the first nontrivial, in fact quadratic, randomized lower bounds on the problems like Knapsack and Bounded

Integer Programming.

Partitionable Graphs. The famous Berge conjecture concerns perfect graphs, which have many important applications in biology, scheduling, traffic phasing, radio frequency assignment, etc. The nnclass of partitionable graphs includes all "counterexamples" to Berge's conjecture (if one exists), as was observed by Lovasz (1972) and by Padberg (1974). Still, there is very little known about this family. There are only a handful of infinite families known within this class, and even those are the ones with very high symmetries. Endre Boros (PM, RU), Vladimir Gurvich (V), and Stefan Hougardy found a recursion which generates several infinite families of partitionable graphs, based on the so called "critical edges" of some cliques of such a graph. For many of these generated families they were able to show that chordless odd cycles are inherited, hence proving that no counterexample to Berge's conjecture can exist within those.

The following six results connect up two special years, Large Scale Discrete Optimization and Networks.

The Element Connectivity Problem in Survivable Network Design The element connectivity problem falls in the category of survivable network design problems – it is intermediate to the versions that ask for edge-disjoint and vertex-disjoint paths. The edge version is by now well understood from the view-point of approximation algorithms, but very little is known about the vertex version. Vijay Vazirani (V) considered the problem where vertices are partitioned into two sets: terminals and non-terminals. Only edges and non-terminals can fail – we refer to them as *elements* – and only pairs of terminals have connectivity requirements, specifying the number of element-disjoint paths required. Vazirani found an algorithm that achieves an approximation guarantee of factor $2H_k$, where k is the largest requirement and $H_n = 1 + \frac{1}{2} + \cdots + \frac{1}{n}$. Besides providing possible insights for solving the vertex-disjoint paths version, the element connectivity problem is of independent interest, since it models a realistic situation.

Primitive Ring Generation and Ring Grooming in Design of Optimal Networks A central problem in the design of optical networks is the placement of the ADM (Add/Drop Multiplexes) devices. Many restricted versions of this general design problem have been studied in the literature. The general version can be formulated as a discrete optimization problem. Due to the complexity of the problem, the current methodology recommends dividing the general version of the design problem into two phases, called Primitive Ring Generation and Ring Grooming. The General Design Problem, Primitive Ring Generation, and Ring Grooming are NP-Hard. We have obtained theoretical results on the quality of approximation algorithms. Primitive Ring Generation has two variants, depending on whether or not traffic demands can be split. Gruia Calinescu (PD) and Peng-Jun Wan designed new algorithms for Primitive Ring Generation guaranteed to produce an output of cost at most 25% bigger than optimum if traffic demands can be split, and at most 50% bigger than optimum if traffic demands cannot be split. All previously known algorithms either have bigger worst-case guarantees or are provably not better than the new algorithms. For Ring Grooming, the traffic granularity parameter seems to influence the quality of the approximation algorithms. Gruia Calinescu (PD), Li-Wu Liu and Peng-Jun Wan provided a nontrivial analysis of a natural algorithm, proving performance guarantees of 1.5 and 2.5 respectively, for the important case when traffic granularity is 4 and 8 respectively.

Directed Multicut vs. Multicommodity Flow Given a graph $G = (V, E)$ and k pairs of vertices $(s_1, t_1), ..., (s_k, t_k)$, a *multicut* is a subset of edges $F \subseteq E$ such that if all edges in F are deleted, none of the pairs (s_i, t_i) are in the same connected component in the remaining graph $G' = (V, E \setminus F)$.

In the minimum multicut problem, we are also given a positive cost $c(e)$ for each edge $e \in E$ and we wish to find the multicut of minimal total cost. This problem is known to be NP-complete for $k \geq 3$, and therefore it is interesting to consider approximation algorithms for it. One way to do it is to state the multicut problem as an integer programming problem, relax it to a linear programming problem, and solve this problem or it's dual - the *maximum multicommodity flow problem.* The maximal multicommodity flow will always be at most as big as cost of the minimal multicut. On the other hand, it is easy to see that there are multicuts which cost at most k times the value of the maximal flow. Thus, the described approach gives a factor-k approximation algorithm for a minimal multicut. It turns out that when the underlying graph is *undirected* more can be said, and in fact the cut/flow ratio in this case can be at most as large as $O(\log(k))$. On the positive side, an algorithm of Garg, Vazirani and Yannakakis exists, which takes a maximal flow and constructs a multicut which costs at most as much as $O(\log(k))$ times the flow cost. Michael Saks (PM, RU), Alex Samorodnitsky (PD), and Leonid Zosin (PD) showed that in the more general *directed* case, the cut/flow ratio could be as large as the trivial upper bound of k. It seems, therefore, that to find a better than a factor-k approximation, one must look for other approaches than the simple linear relaxation technique.

A Unified Approach to Approximating Resource Allocation and Scheduling Amotz Bar-Noy (ATTL), Reuven Bar-Yehuda (V), Ari Freund, Joseph (Seffi) Naor, and Baruch Schieber developed a general framework for resource allocation and scheduling problems. Given a resource of fixed size, they developed algorithms that approximate the maximum throughput or the minimum loss by a constant factor. Their approximation factors apply to many prominent problems, among others: (i) Real-time scheduling of jobs on parallel machines. (ii) Bandwidth allocation for sessions between two endpoints. (iii) General caching. (iv) Dynamic storage allocation. (v) Bandwidth allocation on optical line and ring topologies. For many of the problems, this is the first constant factor approximation algorithm. The algorithms are extremely simple and efficient. They use the local-ratio technique but they can also be equivalently interpreted within the primal-dual schema.

Load Balancing in a Distributed Computing Environment Endre Boros (PM, RU) and Ewald Speckenmeyer (V) studied load balancing in a distributed computing environment. They showed that the minimum message length to re-balance the load of n processors can be computed in linear $O(n)$ time, even though the actual balancing may require $O(n^2)$ messages exchanged between the processors.

A Bottleneck Load Balancing Minimization Problem Consider a processor network with n overloaded and m underloaded processors. The overloaded and the underloaded nodes are connected by a completely bipartite network $K(n, m)$. Ewald Speckenmeyer (V) considered the problem of achieving a uniform load-distribution among all processors via sending appropriate amount of load from overloaded to underloaded nodes so that the maximum of load sent via a link is minimized. He showed that this problem can be solved in time $O(mn)$.

The Zero-Extension Problem. The Zero-Extension Problem, a natural generalization of multiway cut, is defined as follows: given a metric space (T, m) with at least three points called terminals, and a graph $G = (V, E)$ whose edges have nonnegative cost c, and such that T is a subset of V, find a mapping from V to T which keeps T invariant and minimizes the summation over all the edges of E of the product between the cost of an edge and the distance in between the terminals where the endpoints of the edge are mapped. Multiway Cut corresponds to the case when the metric space is uniform. We denote by k the number of terminals. Kleinberg and Tardos studied a generalization of Zero-Extension, called Metric Labeling, and provided an approximation algorithm with performance guarantee of the order of $\log k \log\log k$. The Zero-Extension problem has a natural linear programming relaxation. Karzanov gave a characterization of graph-induced metric spaces for which the linear programming relaxation equals the value of the Zero-Extension problem. Griua Calinescu (PD), Howard Karloff (PM, ATTL) and Yuval Rabani (V) found approximation algorithms for the Zero-Extension problem: a factor of $\log k$ rounding procedure for the linear programming relaxation, a $15/7 > 2$ lower bound for the integrality ratio of the linear program, and a dual-based constant-ratio rounding procedure for the case when G is a tree.

Formulating Semidefinite Programming Problems as Smooth Convex Nonlinear Optimization Problems Consider the diagonal entries d_j, $j = 1, 2, \ldots, n$, of the matrix D in an LDL^T factorization of an $n \times n$ matrix X. As a function of X, each d_j is well-defined on the closed domain of positive semidefinite matrices. Robert Vanderbei (PM, PU) and Hande Yurttan Benson (GS, PU) showed that these functions are twice continuously differentiable and concave throughout the interior of this domain. Using these facts, they showed how to formulate semidefinite programming problems as standard convex optimization problems that can be solved using an interior-point method for nonlinear programming. To summarize, these results allow one to marry implementations of SDP with implementations of general nonlinear programming solvers thereby creating the possibility to solve problems from both of these problem classes individually with one software system but more importantly to solve hybrid problems that contain a mix of traditional NLP stuff with some SDP components.

Approximation of the Permanent and of the Mixed Discriminant It is $\#P$ hard to compute the permanent of a nonnegative matrix, but it seems like a realistic goal to approximate the permanent efficiently. A strongly polynomial algorithm achieving approximation up to a multiplicative factor of e^n was given in 1998 by Linial, Samorodnitsky, and Wigderson. Alex Samorodnitsky (PD) has extended these results to some families of $0-1$ matrices, such as *dense matrices* (more than half the entries in each row and column are 1) and *random matrices* (each entry is assigned zero or one randomly and independently). As it turns out, the earlier results may also be generalized to approximation of a *mixed discriminant* of an n-tuple of positive semidefinite (PSD) matrices. This line of research has been pursued by Leonid Gurvits (NEC) and Alex Samorodnitsky (PD). They defined notions of *double stochasticity* and of a *doubly stochastic scaling* (generalizing the usual matrix scaling) for an n-tuple of PSD matrices, and constructed a polynomial-time algorithm for such a generalized scaling. This together with bounds on the mixed discriminant of doubly stochastic n-tuples leads to a deterministic polynomial-time approximation algorithm for a mixed discriminant. It also provides a deterministic polynomial-time approximation algorithm for the mixed volume of n convex bodies, answering a question of Dyer, Gritzmann and Hufnagel. Following this work, very recently Leonid Gurvits proved a generalization of the Van der Waerden conjecture for permanents to the mixed discriminant case. This result implies, in particular, that the Gurvits-Samorodnitsky algorithm gives a polynomial-time deterministic approximation algorithm for the mixed discriminant whose performance is of the same order as of the best randomized algorithms (and also a near-best possible approximation factor for the mixed volume).

Linearity Tests and Probabilistically Checkable Proofs Probabilistically checkable proofs (PCPs) are a model of computation that helps in understanding the approximability properties of combinatorial optimization problems. In 1998, Luca Trevisan (PD) developed a general analysis technique and a linearity test (a probabilistic algorithm to verify whether a given Boolean function is linear) achieving *amortized query complexity* of 1.5 bits. (A testing algorithm uses q amortized bits if, for some constant k it reads qk bits and has error probability at most 2^{-k}.) (One should mention that amortized query complexity couldn't go below 1, unless $P = NP$.) Linearity testing is interesting because of its applications in the PCPs. Sudan and Trevisan had succeeded in 1998 in adapting the 1998 approach of Trevisan to the PCP setting, obtaining a PCP characterization of NP using, essentially, 1.5 amortized query bits. Trevisan conjectured in 1998 that for any $\epsilon > 0$ there exists a linearity testing algorithm achieving amortized query complexity of $1+\epsilon$ bits, and that NP can be (PCP)-characterized using $1+\epsilon$ amortized query bits. He made the following conjecture, which implies the possibility of $1+\epsilon$ amortized query bits linearity testing: For any $f : \{-1,1\}^n \to \{-1,1\}$, for any nonempty graph $G = ([k], E)$,

$$\mathcal{E}_{x_1,\ldots,x_k} \prod_{(i,j)\in E} f(x_i)f(x_j)f(x_ix_j) \leq max_\alpha |\hat{f}(\alpha)|,$$

where $\hat{f}$ is the Fourier transform of f. Alex Samorodnitsky (PD) and Trevisan (PD) proved this conjecture and showed, moreover, that it can be adopted to the PCP setting, implying that NP can be (PCP)-characterized using $1+\epsilon$ amortized query bits. In particular this provides a significantly easier proof of a famous result of Hastad on hardness of approximation of the size of the maximal clique in a graph.

Generalized Matrix Scaling and Strategy Properties in Combinatorial Games Lev Bregman (V) considered a game with the set $\{X_1, X_2, .., X_p\}$ as the pure strategies set for the first player and the set $\{Y_1, Y_2, .., Y_q\}$ as the pure strategies set for the second player (X_i are non-negative vectors in R^k and Y_j are non-negative vectors in R^l). Let H be a $k * l$- matrix and define the payoff for the first player as (X_i, HY_j), if the players use their strategies X_i and Y_j respectively. We usually apply this framework to games with strategy sets of a combinatorial nature and thus term such games “combinatorial games”. The numbers p and q are assumed to be much bigger than k and l. We are interested in finding optimal mixed strategies of the game which may be described by a small number of parameters. In particular, a product form is considered when the probability of a pure strategy is represented as a product of factors dependent on only one component of the strategy. To get such a representation, Bregman used a generalized matrix scaling and generalized permanents. The necessary and sufficient condition for the existence of such representations was established.