05031 Abstracts Collection
Algorithms for Optimization with Incomplete Information
— Dagstuhl Seminar —

Susanne Albers[1], Rolf H. Möhring[2], Georg Ch. Pflug[3] and Rüdiger Schultz[4]

[1] Universität Freiburg, DE
salbers@informatik.uni-freiburg.de
[2] TU Berlin, DE
[3] Universität Wien, AT
georg.pflug@univie.ac.at
[4] Universität Duisburg-Essen, DE
schultz@math.uni-duisburg.de

Abstract. From 16.01.05 to 21.01.05, the Dagstuhl Seminar 05031 "Algorithms for Optimization with Incomplete Information" was held in the International Conference and Research Center (IBFI), Schloss Dagstuhl. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Online optimization, robust optimization, stochastic programming, stochastic scheduling

05031 Summary – Algorithms for Optimization with Incomplete Information

The purpose of this Seminar was to bring together top specialists working in algorithms for optimisation when the decision maker has only partial information. While problem descriptions in the different approaches to optimisation with incomplete information are quite similar, solution concepts and methods of solution may be quite different. Traditionally, the *stochastic programming* community has focussed on problems, where all uncertainty is due to the fact that concrete realizations are unknown, but the probability distributions from which they stem are fully known. The quality of the solution is typically measured in average case sense. In contrast, the *online optimisation* community assumes no particular probability model. Therefore the focus is traditionally on worst-case analysis. Recently, new developments made the gap between the two communities smaller. *Robust optimisation* replaces the assumption of a known probability distribution

by an assumption about the range of possible values. *Stochastic scheduling* incorporates ideas of the competitiveness of algorithms with stochastic models for the demands. The typical assumption in stochastic programming that decisions do not influence the underlying probability distribution can usually no longer be maintained in stochastic scheduling.

To facilitate familiarizing of the different communities with each others ways of thinking, basic concepts, and basic research questions the seminar was started by four one-hour overview talks. These were delivered by Jiri Sgall (*Online optimisation*), Andrzej Ruszczynski (*Stochastic Programming*), Garud Iyengar (*Robust Optimisation*), and Marc Uetz (*Stochastic Scheduling*).

The regular program consisted of 38 thirty minutes talks, which could be classified into the following subgroups: *Robust and minimax optimisation* (Sim, Dupacova); *Two- and multistage stochastic optimisation* (Hochreiter, Dye, Sen, Stougie, Tomasgard); *Assessing quality of solution* (Morton, Rambau); *Approximation* (Higle, Swamy, van der Vlerk); *Algorithmic approaches using game theory and nonlinear programming* (Lorenz, Steinbach, Kleywegt, Bastin, Norkin); *Applications in Communications and Robotics* (Erlebach, Fekete, Epstein, Richter); *Dynamic stochastic optimisation* (Weiss, Philpott, Nino-Mora); *Average case competitive analysis* (Fujiwara, Vredefeld); *Competitiveness Analysis* (van Stee, Schäfer, Ebenlendr, Zhang, Skutella); Risk issues (Dentcheva, Eichhorn); *Stochastic online scheduling* (Megow, Schulz, Krumke); *Probabilistic criteria* (Henrion, Hoogeveen).

To assess the results of this seminar, on Thursday afternoon an open discussion was held about different views and perceptions on optimisation with incomplete information. The results of this discussion can be summarized as follows:

What the communities have in common is:

- The desire for optimality.
- The desire for more efficient algorithms, i.e. better/faster results.
- The fact that solutions, which require clairvoyance are not implementable.
- The necessity of comparing the non-clairvoyant solution to the ideal clairvoyant solution by either taking differences (value of perfect information) or ratios (competitive ratio).
- The distinction between individual solutions and solution rules (policies).
- The necessity of approximation.
- The interest in complexity issues.

What distinguishes the communities is:

- The way uncertainty is modelled (from sets of possible values via probability distributions to families of probability distributions).
- The frequency of decision making (once in a while versus online).
- The objective (to look for worst cases, average cases, include risks , chance constraints etc.).
- The class of problems (general as multistage LP, QP, MIP or specialised as scheduling, packing, sequencing).

- The way information is revealed (fixed observation times versus uncertainty about when and if ever information will be available).
- The view on risk.

Some participants brought up their individual views on the topic. It was felt that the advantage of probabilistic modelling lies in the sound concept of probability, developed over centuries, and the clear way of how to obtain and process information (samples). On the other hand, the assumption that a probability model is governing the data process is not always fulfilled, or information is so poor that range sets is all we have. Also, in long term models it is unrealistic to assume that probability distributions do not change over time. Adaptive algorithms in the broad sense are a way to circumvent this difficulty.

This inspired a discussion about bridges and possible collaboration between the communities. As already existing bridges were cited: *Minimax and robust approaches, stochastic competitiveness analysis, certain stochastic dynamic models, complexity studies in stochastic optimisation.* The need for more real world data and problems was expressed as well as the unanimous wish to study special problem classes which were presented at this seminar in more detail.

Joint work of: S. Albers (Universität Freiburg, DE), R. H. Möhring (TU Berlin, DE), G. Ch. Pflug (Universität Wien, AT), R. Schultz (Universität Duisburg-Essen, DE)

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/113

An adaptive trust-region approach for nonlinear stochastic optimisation with an application in discrete choice theory

Fabian Bastin (Cerfacs - Toulouse, F)

We consider stochastic nonlinear programs, restricting ourself to differentiable, but possibly non-convex, problems. The non-convexity leads us to consider nonlinear approaches, designed to find second-order critical solutions. We focus here on the use of trust-region approaches when solving a sample average approximation, and adapt the technique to only use sub-samples when possible, adjusting the sample size at each iteration. We show that under reasonable assumptions, we solve the original SAA problem. We also consider an extension to the estimation of mixed logit models, that are popular in discrete choice theory when the population heterogeneity is taken into account. We present numerical experimentations underlining the practical interest of the method. We finally examine some avenues and preliminary experimentations for future research.

Keywords: Nonlinear stochastic programming, Monte-Carlo, mixed logit, discrete choice, trust-region

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/55

Risk Shaping by Stochastic Dominance Constraints

Darinka Dentcheva (Stevens Institute of Technology, USA)

We consider a novel model for decision making under uncertainty and risk. Our main objective is to develop a new theory that would allow to shape the risk associated with random outcomes. We propose a model involving stochastic dominance constraints with respect to random benchmarks. We develop optimality and duality theory for these models. In this talk we present a new form of the conditions which do not use splitting as in our earlier presentations.

The analysis demonstrates that the expected utility model of von Neumann and Morgenstern is dual to the model with stochastic dominance constraints. We also show that second order stochastic dominance constraint may be regarded as a convexification of the first order constraint.

Keywords: Stochastic programming, risk, stochastic dominance, semi-infinite optimization, optimality conditions, duality, utility, convexification

Joint work of: Dentcheva, Darinka; Ruszczynski, Andrzej

Uncertainties in stochastic programming models: The minimax approach

Jitka Dupacová (Charles University - Prague, CZ)

50 years ago, stochastic programming was introduced to deal with uncertain values of coefficients which were observed in applications of mathematical programming. These uncertainties were modeled as random and the assumption of complete knowledge of the probability distribution of random parameters became a standard. Hence, there is a new type of uncertainty concerning the probability distribution.

Using a hypothetical, ad hoc distribution may lead to bad, costly decisions. Besides of a subsequent output analysis it pays to include the existing, possibly limited information into the model, cf. the minimax approach which will be the main item of this presentation. It applies to cases when the probability distribution is only known to belong to a specified class of probability distributions and one wishes to hedge against the least favorable distribution. The minimax approach has been developed for special types of stochastic programs and special choices of the class of probability distributions and there are recent results aiming at algorithmic solution of minimax problems and on stability properties of minimax solutions.

Keywords: Stochastic programming models, minimax approach

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/45

Subtree decomposition for multistage stochastic programs

Shane Dye (University of Canterbury - Christchurch, NZ)

A class of algorithms for solving multistage stochastic recourse problems is described. The scenario tree is decomposed using a covering collection of subtrees. The approach is illustrated with two examples: adapting the diagonal quadratic approximation algorithm and adapting nested Bender's decomposition. The approach leads to a class of methods based on the subtree cover chosen (including the original implementation of the algorithm adapted).

This approach increases flexibility in the size, number and structure of subproblems for multistage stochastic programming decomposition methods.

Keywords: Stochastic programming, scenario tree, decomposition

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/56

Optimal and Online Scheduling on Related Machines

Tomas Ebenlendr (Academy of Sciences - Prague, CZ)

We consider the problem of preemptive scheduling on uniformly related machines. We present a semi-online algorithm which, if the optimal makespan is given in advance produces an optimal schedule. Using the standard doubling technique, this yields a 4 competitive deterministic and e 2.71 competitive randomized online algorithms. In addition it matches the performance of previously known algorithms for the offline case, with considerably simpler proof. Finaly we study the performance of greedy heuristics for the same problem.

Keywords: Online, preemption, scheduling

Polyhedral Risk Measures and Lagrangian Relaxation in Electricity Portfolio Optimization

Andreas Eichhorn (HU Berlin, D)

We present a multistage stochastic programming model for mean-risk optimization of electricity portfolios containing physical components and energy derivative products. Stochasticity enters the model via the uncertain (time-dependent) prices and electricity demand. The objective is to maximize the expected overall revenue and, simultaneously, to minimize a multiperiod risk measure, i.e., a risk measure that takes into account the intermediate time cash values. We compare the effect of different multiperiod risk measures taken from the class of polyhedral risk measures which was suggested in our earlier work.

Furthermore, we discuss how such a mean-risk optimization problem can be solved by dual decomposition techniques (Lagrangian relaxation).

Keywords: Stochastic programming; mean-risk optimization; risk measure; lagrangian relaxation; electricity

Joint work of: Eichhorn, Andreas; Römisch, Werner; Wegner, Isabel

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/57

Tracking mobile users

Leah Epstein (The Interdisciplinary Center - Herzliya, IL)

Cellular telephony systems, where locations of a mobile users may be unknown at some times, are becoming more common. Mobile users are roaming in a zone. A user reports its location only if it leaves the zone entirely. We consider cellular zones with n cells and m mobile users roaming among the cells. The location of the users is uncertain and is given by m probability distribution vectors. The Conference Call Search problem (CCS) deals with tracking a set of mobile users, in order to establish a call between all of them. The search is performed in a limited number of rounds, and the goal is to minimize the expected search cost. In the "unit cost model", a single query for a cell outputs a list of users located in that cell. The "bounded bandwidth" model allows a query for a single user per cell in each round. We discuss three types of protocols; oblivious, semi-adaptive and adaptive search protocols. An oblivious search protocol decides on all requests in advance, and stops only when all users are found. A semi-adaptive search protocol decides on all the requests in advance, but it stops searching for a user once it is found. An adaptive search protocol stops searching for a user once it has been found (and its search strategy may depend on the subsets of users that were found in each previous round). We establish the differences between the distinct protocol types and answer some open questions which were posed in previous work on the subject.

Keywords: Mobile users, PTAS

Joint work of: Epstein, Leah; Levin, Asaf

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/58

Network Discovery and Verification

Thomas Erlebach (University of Leicester, GB)

Consider the problem of discovering (or verifying) the edges and non-edges of a network, modelled as a connected undirected graph, using a minimum number of queries. A query at a vertex v discovers (or verifies) all edges and non-edges whose endpoints have different distance from v. In the network discovery problem, the edges and non-edges are initially unknown, and the algorithm must select

the next query based only on the results of previous queries. We study the problem using competitive analysis and give a randomized on-line algorithm with competitive ratio O(sqrt(n*log n)) for graphs with n vertices. We also show that no deterministic algorithm can have competitive ratio better than 3. In the network verification problem, the graph is known in advance and the goal is to compute a minimum number of queries that verify all edges and non-edges. This problem has previously been studied as the problem of placing landmarks in graphs or determining the metric dimension of a graph. We show that there is no approximation algorithm for this problem with ratio o(log n) unless P=NP. Furthermore, we prove that the optimal number of queries for d-dimensional hypercubes is Theta(d/log d).

Keywords: On-line algorithms; set cover; landmarks; metric dimension

Joint work of: Beerliova, Zuzana; Eberhard, Felix; Erlebach, Thomas; Hall, Alexander; Hoffmann, Michael; Mihalak, Matus; Ram, L. Shankar

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/59

Note on Negative Probabilities and Observable Processes

Ulrich Faigle (Universität Köln, D)

A mathematical framework for observable processes is introduced via the model of systems whose states may be time dependent and described by possibly "negative probabilities". The model generalizes and includes the linearly dependent models or observable operator models for classical discrete stochastic processes. Within this model a general convergence result for finite-dimensional processes, which generalize finite state (hidden) Markov models, is derived. On the philosophical side, the model furthermore offers an explanation for Bell's inequality in quantum mechanics.

Keywords: Negative Probability, Observable Process, Markov Chain, Stochastic Process, Bell's Inequality, HHM, LDP, OOM

Joint work of: Faigle, Ulrich; Schoenhuth, Alexander

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2005/108

Searching with an Autonomous Robot

Sándor Fekete (TU Braunschweig, D)

We discuss online strategies for visibility-based searching for an object hidden behind a corner, using Kurt3D, a real autonomous mobile robot. This task is closely related to a number of well-studied problems.

Our robot uses a three-dimensional laser scanner in a stop, scan, plan, go fashion for building a virtual three-dimensional environment.

Besides planning trajectories and avoiding obstacles, Kurt3D is capable of identifying objects like a chair.

We derive a practically useful and asymptotically optimal strategy that guarantees a competitive ratio of 2, which differs remarkably from the well-studied scenario without the need of stopping for surveying the environment. Our strategy is used by Kurt3D, documented in a separate video.

Keywords: Searching, visibility problems, watchman problems, online searching, competitive strategies, autonomous mobile robots, three-dimensional laser scanning, Kurt3D.

Joint work of: Fekete, Sándor; Klein, Rolf; Nüchter, Andreas

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/191

Average-Case Competitive Analyses for Ski-Rental Problems

Hiroshi Fujiwara (Universität Freiburg, D)

Let s be the ratio of the cost for purchasing skis over the cost for renting them. Then the famous result for the ski-rental problem shows that skiers should buy their skis after renting them $(s-1)$ times, which gives us an optimal competitive ratio of $2-\frac{1}{s}$. In practice, however, it appears that many skiers buy their skis before this optimal point of time and also many skiers keep renting them forever. In this paper we show that these behaviors of skiers are quite reasonable by using an *average-case competitive ratio*. For an exponential input distribution $f(t)=\lambda e^{-\lambda t}$, optimal strategies are (i) if $\frac{1}{\lambda} \leq s$, then skiers should rent their skis forever and (ii) otherwise should purchase them after renting approximately $s^2\lambda \quad (< s)$ times. Thus average-case competitive analyses give us the result which differs from the worst-case competitive analysis and also differs from the traditional average cost analysis. Other distributions and related problems are also discussed.

Keywords: Online algorithm, competitive analysis

Joint work of: Fujiwara, Hiroshi;Iwama, Kazuo

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/60

Properties and Calculation of Singular Normal Distributions

René Henrion (Weierstraß Institut - Berlin, D)

The need for calculating and characterizing singular normal distributions arises in a natural way when considering chance constraints of the type P(Az <= b(x) >= p, where A is a rectangular matrix having more rows than columns, b is some function and z is a random vector having some nondegenerate multivariate normal distribution. Such situation is typical, for instance, in stochastic networks, where a comparatively small random vector may induce a possibly large number of linear inequality constraints.

Passing to the transformed random variable q:=Az, the constraint can be equivalently rewritten as F(b(x))>= p, where F is the distribution function of q. In contrast to the original random vector z, the transformed vector q has a singular normal distribution. The talk demonstrates how to get back from here to (a sum of) regular normal distributions under a full rank regularity condition. This allows for an efficient calculation of singular normal distributions and provides a numerical method which outperforms competing procedures in moderate dimensions. Computational results for test examples are provided for the sake of comparison.

In general, if the mentioned regularity condition is violated, then the singular normal distribution function F may even lack continuity.

The talk provides an equivalent criterion for Lipschitz continuity of F and characterizes differentiability and subdifferentiability of F.

Keywords: Singular normal distribution, chance constraints, normal probability of polyhedra

Joint work of: Henrion, René; Szantai, Tamas

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/109

Solution Validation in Multistage Stochastic Linear Programs

Julie L. Higle (University of Arizona, USA)

In this presentation, we investigate solution validation techniques for multistage stochastic linear programs based on estimates of error bounds.

Different approaches involving manipulations of the nonanticipativity constraints are studied both analytically and computationally.

Keywords: Solution validation, error bounds, stochastic programming

Joint work of: Higle, Julia L.; Zhao, Lei

Scenario Optimization for Multi-Stage Stochastic Programming Problems

Ronald Hochreiter (Universität Wien, A)

The field of multi-stage stochastic programming provides a rich modelling framework to tackle a broad range of real-world decision problems. In order to numerically solve such programs - once they get reasonably large - the infinite-dimensional optimization problem has to be discretized. The stochastic optimization program generally consists of an optimization model and a stochastic model. In the multi-stage case the stochastic model is most commonly represented as a multi-variate stochastic process. The most common technique to calculate an useable discretization is to generate a scenario tree from the underlying stochastic process. In the first part of the talk we take a look at scenario optimization from the viewpoint of a decision taker, to provide rather non-technical insights into the problem. In the second part of the talk we examplify scenario tree generation by reviewing one specific algorithm based on multi-dimensional facility location applying backward stagewise clustering. An example from the area of financial engineering concludes the talk.

Keywords: Stochastic programming, scenario generation, facility location, financial engineering

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/61

Getting rid of stochasticity: applicable sometimes

Han Hoogeveen (Utrecht University, NL)

We consider the single-machine scheduling problem of minimizing the number of late jobs. This problem is well-studied and well-understood in case of deterministic processing times. We consider the problem with stochastic processing times, and we show that for a number of probability distributions the problem can be reformulated as a deterministic problem (and solved by the corresponding algorithm) when we use the concept of minimum success probabilities, which is, that we require that the probability that a job complete on time is 'big enough'. We further show that we can extend our approach to the case of machines with stochastic output.

Keywords: Scheduling, sequencing, single machine, number of late jobs, stochastic processing times, minimum success probability, dynamic programming, unreliable machines

Joint work of: Van den Akker, Marjan; Hoogeveen, Han

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/192

A survey of robust optimization

Garud Iyengar (Columbia University, USA)

In this talk we will review the literature, the main models and applications of robust optimization. The goal is to place robust optimization in context by highlighting the strengths of this formalism and also some of its shortcomings. We will also briefly mention possible research questions and connections to other methods of addressing uncertain optimization.

A Derivative-Free Algorithm for Stochastic Optimization

Anton J. Kleywegt (Georgia Institute of Technology, USA)

We propose an algorithm for a stochastic optimization problem that does not require derivatives of the expected objective value function f(x) := E[F(x,w)] or the sample function F(x,w) to be computed. For a given point x, the algorithm only estimates the expected objective value f(x), for example with a sample average approximation. It is required that the expected value function f be smooth, but the sample function F need only be Lipschitz continuous. Such algorithms are appealing when the approximation of the objective function f is not differentiable everywhere, or when the practitioner cannot or does not want to implement the code to compute the derivatives. We also present some numerical results on a set of test problems.

Keywords: Stochastic optimization, derivative free optimization

Joint work of: Kleywegt, Anton J.; Bharadwaj, Vijay

Probabilistic Competitive Analysis of an Online-Dial-a-Ride Problem

Sven O. Krumke (TU Kaiserslautern, D)

Competitive analysis has become one of the standard methods to evaluate online algorithms: an algorithm is called c-competitive if for all input instances the online algorithm returns a solution of cost at most c-times the optimal offline cost.

Competitive analysis is a worst-case analysis and has often been criticized as being too pessimistic. Often the offline adversary is simply too powerful to allow for meaningful competitiveness results.

We consider the following online-dial-a-ride-problem: a server moves in a tree network and transportation requests become known (released) over time. The goal is to minimize the total distance travelled. For this problem standard

competitive analysis fails in the respect that we hit the so called triviality-barrier: no algorithm can be competitive at all!

However, if we consider random inputs (generated from a wide class of distributions) we are able to show that a simple and natural online-strategy actually achieves a (probabilistic) competitive ratio of 1.

Keywords: Competitive analysis, probabilistic analysis, average case analysis, triviality barrier

Joint work of: Krumke, Sven O.; Hiller, Benjamin

Disruption Management and Planning with Uncertainties in Aircraft Planning

Ulf Lorenz (Universität Paderborn, D)

An important insufficiency of modern industrial plans still is their lack of robustness. Disruptions prevent companies from operating as planned before and induce high costs for trouble shooting. The main reason for the severe impact of disruptions stems from the fact that planners do traditionally consider deterministic input data to be available at planning time. In practice, there are often only distributions over the possible input data available.

The Repair Game is a formalization of a planning task, which brings two branches of computer science — game tree search and logistic planning optimization with OR tools — together. Playing it performs disruption management and generates robust plans with the help of game tree search. Our method significantly outperformed the traditional one by means of simulations.

Keywords: Uncertainty, planning, game playing, aviation application

Joint work of: Ehrhoff, Jan; Grothklags, Sven; Lorenz, Ulf

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/62

Models and Algorithms for Stochastic Online Scheduling

Nicole Megow (TU Berlin, D)

We introduce a model for non-preemptive scheduling under uncertainty. In this model, we combine the main characteristics of online and stochastic scheduling in a simple and natural way. Job processing times are assumed to be stochastic, but in contrast to the classical stochastic scheduling models, we assume that jobs arrive online over time, and there is no knowledge about the jobs that will arrive in the future. The model incorporates both, stochastic scheduling and online scheduling as a special case. The particular setting we analyze is parallel machine scheduling, with the objective to minimize the total weighted completion times of jobs. We propose simple, combinatorial online scheduling policies

for that model, and derive performance guarantees that match the currently best known performance guarantees for stochastic parallel machine scheduling. For processing times that follow NBUE distributions, we improve upon previously best known performance bounds, even though we consider a more general setting.

Keywords: Stochastic scheduling, online optimization, weighted completion time

Joint work of: Megow, Nicole; Uetz, Marc; Vredeveld, Tjark

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/110

Assessing Solution Quality in Stochastic Programs

David P. Morton (University of Texas - Austin, USA)

Assessing whether a solution is of high quality (optimal or near optimal) is a fundamental question in optimization. We develop Monte Carlo sampling-based procedures for assessing solution quality in stochastic programs. Quality is defined via the optimality gap and our procedures' output is a confidence interval on this gap. We review a multiple-replications procedure and then present a result that justifies a computationally simplified single-replication procedure. Even though the single replication procedure is computationally significantly less demanding, the resulting confidence interval may have low coverage for small sample sizes on some problems. We provide variants of this procedure and provide preliminary guidelines for selecting a candidate solution.

Both are designed to improve the basic procedure's performance.

Keywords: Stochastic programming, Monte Carlo simulation

Joint work of: Morton, David P.; Bayraksan, Guzin

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/63

Marginal productivity index policies for scheduling restless bandits with switching penalties

José Niño-Mora (Univ. Carlos III de Madrid, E)

We address the problem of designing a tractable, well-grounded policy for the dynamic allocation of effort to a collection of restless bandit projects, i.e. binary-action (active/passive) Markov decision processes, in which sequence-independent switching penalties (costs or delays) are incurred when switching from one project to another. We deploy the framework of partial conservation laws, introduced by Niño-Mora (2001, 2002), to establish the existence of and calculate a marginal productivity index (MPI), under certain conditions. The MPI, which extends earlier indices proposed by Gittins (1979) and Whittle (1988), yields a corresponding MPI policy, which prescribes to dynamically engage the project with larger index.

Keywords: Stochastic scheduling; restless bandits; index policies; switching penalties

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/64

Minorant methods for stochastic global optimization

Vladimir I. Norkin (Ukrainian Acad. of Sciences, RUA)

We develop numerical methods for solution of stochastic global optimization problems: min$[F(x) = Ef(x,)|x \in X]$ and $min[F(x) = Pf(x,)0|x \in X]$, where x is a finite dimensional decision vector with possible values in the set X, ęØ is a random variable, $f(x,)$ is a nonlinear function of variable x, E and P denote mathematical expectation and probability signs respectively.

These methods are based on the concept of stochastic tangent minorant, which is a random function $(x, y,)$ of two variables x and y with expected value $(x, y) = E(x, y,)$ satisfying conditions: (i) $(x, x) = F(x)$, (ii) $(x, y)F(x)$ for all x,y. Tangent minorant is a source of information on a function global behavior. We develop a calculus of (stochastic) tangent minorants.

We develop a stochastic analogue of Pijavskiąś global optimization method and a branch and bound method with stochastic minorant bounds.

Applications to optimal facility location and network reliability optimization are discussed

Keywords: Stochastic global optimization, stochastic tangent minorant, branch and bound method

Joint work of: Norkin, Vladimir; Onischenko, Boris

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/211

Rowing to Barbados

Andy Philpott (Univ. of Auckland, NZ)

In October 2003, sixteen boats set off from La Gomera in the Canary Islands headed for Barbados 4800 km distant. Each boat was manned by two oarsmen who were competing in the Transatlantic Challenge, an ocean rowing endurance event. This paper describes an optimization model developed for route planning in this event. It was used successfully by the Holiday Shoppe team to win the race in world record time. We describe the tool, its history, and the way it was used in the race.

Keywords: Ocean rowing, weather routing, dynamic programming, isochrones

Joint work of: Philpott, Andy; Leyland, Geoff

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2005/65

Deferment Control for Reoptimization – How to Find Fair Reoptimized Dispatches

Jörg Rambau (Universität Bayreuth, D)

The german automobile association ADAC maintains a fleet of 1700 vehicles and has agreements with around 5000 service contractors. With these ressources, they help people whose cars have broken down on the road. Those people can call an ADAC help center, and within 10 seconds, an assignment of a service ressource to their request is made. At the same time, for all service vehicles, tours through the assigned requests have to be planned so as to minimize a certain (complicated) cost function for this so-called dispatch.

No usefule knowledge about future requests is available at the time being.

Therefore, the current policy of the automated system, developed in joint work with Sven O. Krumke, is to reoptimize the whole dispatch upon the occurrence of each relevant event, like the arrival of a new request. A similar online-optimization problem appears in the pallet elevator group control in a large distribution center of Herlitz PBS AG in Falkensee near Berlin.

The problem with reoptimization policies in general is that, depending on the reoptimization objective, an arbitrarily large deferment of individual requests can be observed. In a way, individual requests are sacrificed in favor of a good performance according to the reoptimization objective.

Nevertheless, w.r.t. the reoptimization objective, the reoptimization policies in the long run usually perform much better than the currently known policies that can not cause infinite deferment. Therefore, the goal is to modify reoptimization policies so as to prevent deferment.

Sometimes deferment can be almost eliminated by enhancing the reoptimization objective with some terms that penalize waiting, but service in a fixed time can still not be guaranteed, and this kind of objective function engineering is a very time consuming tuning issue, interfering with the orgininal management objective.

In this talk, the new policy of flow and makespan constrained reoptimization with reoptimization admission control is introduced. The main result is that, under d-reasonable load, for any reoptimization model, this policy yields a maximal flow time that is bounded by a constant 2d, depending only on the system load parameter d, not on the instance. In simulation experiments for the elevator group control problem we still obtain a very satisfactory average performance w.r.t. the reoptimization objective.

Keywords: Online optimization, dynamic vehicle dispatching, reoptimization, integer linear program, dynamic column generation, infinite deferment, reasonable load, flow time and makespan constrained reoptimization

Joint work of: Rambau, Jörg

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/66

An improved algorithm for CIOQ switches

Yossi Richter (Tel Aviv University, IL)

The problem of maximizing the weighted throughput in various switching settings has been intensively studied recently through competitive analysis. To date, the most general model that has been investigated is the standard CIOQ (Combined Input and Output Queued) switch architecture with internal fabric speedup $S \geq 1$. CIOQ switches, that comprise the backbone of packet routing networks, are $N \times N$ switches controlled by a switching policy that incorporates two components: Admission control and scheduling. An admission control strategy is essential to determine the packets stored in the FIFO queues in input and output ports, while the scheduling policy conducts the transfer of packets through the internal fabric, from input ports to output ports. The online problem of maximizing the total weighted throughput of CIOQ switches was recently investigated by Kesselman and Rosén [SPAA03]. They presented two different online algorithms for the general problem that achieve non-constant competitive ratios (linear in either the speedup or the number of distinct values, or logarithmic in the value range).

We introduce the first constant-competitive algorithm for the general case of the problem, with arbitrary speedup and packet values. Specifically, our algorithm is 8-competitive, and is also simple and easy to implement.

Keywords: On-line algorithms, competitive ratio, CIOQ switch, packet switching, buffer management, quality of service.

Joint work of: Azar, Yossi; Richter, Yossi

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/67

Recent Advances in Stochastic Programming

Andrzej Ruszczynski (Rutgers Univ. - Piscataway, USA)

We review most important models and methods for decision-making under uncertainty, where the uncertainty is described by stochastic models. We focus on modeling risk aversion. We present new approaches to probabilistic programming, mean-risk models, general risk measures, and stochastic dominance constraints. We also outline cutting-edge research directions.

Keywords: Stochastic programming, risk

Joint work of: Ruszczynski, Andrzej

New Old Algorithms for Stochastic Scheduling

Andreas S. Schulz (MIT - Cambridge, USA)

We consider the stochastic identical parallel machine scheduling problem and its online extension, when the objective is to minimize the expected total weighted completion time of a set of jobs that are released over time. We give randomized as well as deterministic online and offline algorithms that have the best known performance guarantees in either setting, online or offline and deterministic or randomized. Our analysis is based on a novel linear programming relaxation for stochastic scheduling problems that can be solved online.

Keywords: Stochastic scheduling; online algorithms; competitive analysis; approximation algorithms; linear programming relaxations

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/193

Topology Matters: Smoothed Competitiveness of Metrical Task Systems

Guido Schäfer (Università di Roma I, I)

We consider online problems that can be modeled as metrical task systems:

An online algorithm resides in a graph of n nodes and may move in this graph at a cost equal to the distance. The algorithm has to service a sequence of tasks that arrive over time; each task specifies for each node a request cost that is incurred if the algorithm services the task in this particular node. The objective is to minimize the total request plus travel cost. Borodin, Linial and Saks gave a deterministic work function algorithm (WFA) for metrical task systems having a tight competitive ratio of 2n-1.

We present a smoothed competitive analysis of WFA.

Given an adversarial task sequence, we add some random noise to the request costs and analyze the competitive ratio of WFA on the perturbed sequence. We prove upper and matching lower bounds. Our analysis reveals that the smoothed competitive ratio of WFA is much better than its (worst case) competitive ratio and that it depends on several topological parameters of the graph underlying the metric, such as maximum degree, diameter, etc. For example, already for moderate perturbations, the smoothed competitive ratio of WFA is $O(log(n))$ on a clique and $O(\sqrt{n})$ on a line.

We also provide the first average case analysis of WFA.

For a large class of probability distributions, we prove that WFA has O(log(D)) expected competitive ratio, where D is the maximum degree of the underlying graph.

Keywords: Online algorithm, metrical task systems, work function algorithm, smoothed competitive analysis

Joint work of: Schäfer, Guido; Sivadasan, Naveen

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/68

Network Design in the presence of a disruptive opponent

Suvrajeet Sen (University of Arizona, USA)

Network design problems have a variety of applications in computer science, manufacturing, military operations, and of course, telecommunications, and transportation. Traditionally, these models have been treated as deterministic optimization problems in which all data are assumed to be known. Over the last decade or so, authors have started considering the network design problem under uncertainty, and some stochastic programming models have been proposed. Such stochastic programming models are based on the assumption that the decisions themselves do not have any impact on the random variables. However, in the presence of a disruptive opponent, it is difficult to justify the assumption that decisions do not affect the randomn variables, because the latter are used to model the opponent's actions. For instance, the decision to build certain facilities (links or nodes) entices the opponent to choose strategies that cause a disruption of these facilities.

In this lecture we demonstrate that as long as we allow the inclusion of binary variables in the stochastic programming model of network design, we are able to include the possibility of disruptions which are dependent on the decisions made in the network design problem. Thus stochastic programming can be used to develop designs that are more resilient to the opponent's actions.

Keywords: Robust network design, stochastic programming

Online Scheduling

Jiri Sgall (Academy of Sciences - Prague, CZ)

We survey some recent results on scheduling unit jobs. The emphasis of the talk is both on presenting some basic techniques and providing an overview of the current state of the art. The techniques presented cover charging schemes, potential function arguments, and lower bounds based on Yao's principle.

The studied problem is equivalent to the following buffer management problem: packets with specified weights and deadlines arrive at a network switch and need to be forwarded so that the total weight of forwarded packets is maximized. A packet not forwarded before its deadline brings no profit. The presented algorithms improve upon 2-competitive greedy algorithm, the competitive ratio is 1.939 for deterministic and 1.582 for randomized algorithms.

Keywords: Online algorithms; scheduling

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/69

A Robust Optimization Perspective to Stochastic Models

Melvyn Sim (Nat. Univ. - Singapore, SGP)

We propose a robust optimization model with refined uncertainty set that captures random data uncertainties that are potentially distributed asymmetrically. With the new uncertainty set, the robust model has the flexibility of incorporating information on forward and backward deviations of the underlying data uncertainties in achieving less conservative robust solutions while maintaining reasonable probabilistic guarantee against the violation of constraints. We also propose a tractable robust optimization approach for obtaining a feasible solution to stochastic linear optimization with fixed recourse and chance constraints in which we can obtain an upper bound to the minimum objective for all underlying distributions that satisfy the parameters of deviations. An attractive feature of the framework is the computational scalability to multiperiod models. We show an application of the framework to project management.

Joint work of: Sim, Melvyn; Chen, Xin; Sun, Peng

Online Scheduling with Bounded Migration

Martin Skutella (Universität Dortmund, D)

Consider the classical online scheduling problem where jobs that arrive one by one are assigned to identical parallel machines with the objective of minimizing the makespan. We generalize this problem by allowing the current assignment to be changed whenever a new job arrives, subject to the constraint that the total size of moved jobs is bounded by β times the size of the arriving job.

Our main result is a linear time 'online approximation scheme', that is, a family of online algorithms with competitive ratio $1+\epsilon$ and constant migration factor $\beta(\epsilon)$, for any fixed $\epsilon > 0$. This result is of particular importance if considered in the context of sensitivity analysis: While a newly arriving job may force a complete change of the entire structure of an optimal schedule, only very limited 'local' changes suffice to preserve near-optimal solutions. We believe that this concept will find wide application in its own right.

We also present simple deterministic online algorithms with migration factors $\beta = 2$ and $\beta = 4/3$, respectively. Their competitive ratio $3/2$ beats the lower bound on the performance of any online algorithm in the classical setting without migration. We also present improved algorithms and similar results for closely related problems. In particular, there is a short discussion of corresponding results for the objective to maximize the minimum load of a machine. The latter problem has an application for configuring storage servers that was the original motivation for this work.

Keywords: Scheduling, sensitivity analysis, online algorithm

Joint work of: Sanders, Peter; Sivadasan, Naveen; Skutella, Martin

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/70

Tree-Sparse Modeling and Solution of Multistage Stochastic Programs

Marc Steinbach (Zuse Institute Berlin, D)

Multistage stochastic programs are prototypical for nonlinear programs with an inherent tree structure inducing characteristic sparsity patterns in the KKT systems of interior methods. We present an integrated modeling and solution approach for such tree-sparse programs. Three closely related natural formulations having desirable control-theoretic properties lead to KKT system solution algorithms with linear complexity. Application examples from computational finance and process engineering demonstrate the efficiency of the approach.

Keywords: Tree-sparse programs, multistage stochastic optimization, KKT systems, hierarchical sparsity

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/71

Computational complexity of stochastic programming problems

Leen Stougie (TU Eindhoven, NL)

Stochastic programming is the subfield of mathematical programming that considers optimization in the presence of uncertainty. During the last four decades a vast quantity of literature on the subject has appeared.

Developments in the theory of computational complexity allow us to establish the theoretical complexity of a variety of stochastic programming problems studied in this literature. Under the assumption that the stochastic parameters are independently distributed, we show that two-stage stochastic programming problems are ♯P-hard. Under the same assumption we show that certain multistage stochastic programming problems are PSPACE-hard. The problems we consider are non-standard in that distributions of stochastic parameters in later stages depend on decisions made in earlier stages.

Joint work of: Dyer, Martin, Stougie, Leen

Approximation Algorithms for 2-stage and Multi-stage Stochastic Optimization

Chaitanya Swamy (California Institut of Technology - Pasadena, USA)

Stochastic optimization problems attempt to model uncertainty in the data by assuming that (part of) the input is specified by a probability distribution. We consider the well-studied paradigm of stochastic recourse models, where the uncertainty evolves through a series of stages and one can take decisions in each stage in response to the new information learned.

We obtain the first approximation algorithms for a variety of 2-stage and k-stage stochastic integer optimization problems where the underlying random data is given by a "black box" and no restrictions are placed on the costs of the two stages: one can merely sample data from this distribution, but no direct information about the distributions is given.

Our results are based on two principal components. First, we show that for a broad class of 2-stage and k-stage linear programs, where k is not part of the input, given only a "black box" to draw independent samples from the distribution, one can, for any $\epsilon > 0$, compute a solution of cost guaranteed to be within a $(1 + \epsilon)$ factor of the optimum, in time polynomial in $1/\epsilon$, the size of the input, and a parameter λ that is the ratio of the cost of the same action in successive stages which is a lower bound on the sample complexity in the "black-box" model. This is based on reformulating the stochastic linear program, which has both an exponential number of variables and an exponential number of constraints, as a compact convex program, and adapting tools from convex optimization to solve the resulting program to near optimality. In doing so, a significant difficulty that we must overcome is that even evaluating the objective function of this convex program at a given point may be quite difficult and provably hard. To the best of our knowledge, this is the first such result for multi-stage stochastic programs.

Second, we give a rounding approach for stochastic integer programs that shows that approximation algorithms for a deterministic analogue yields, with a small constant-factor loss, provably near-optimal solutions for the stochastic generalization. Thus we obtain approximation algorithms for several stochastic problems, including the stochastic versions of the set cover, vertex cover, facility location, multicut (on trees) and multicommodity flow problems.

Keywords: Algorithms, Approximation Algorithms, Optimization, Convex Optimization, Stochastic Optimization

Joint work of: Swamy, Chaitanya; Shmoys, David

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/72

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2005/72

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2005/72

Facility location with uncertain demand and economies of scale

Asgeir Tomasgard (SINTEF - Trondheim, N)

This paper adresses facility location under uncertain demand. The problem is to determine the optimal location of facilities and allocation of uncertain customer demand to these facilities. The costs of operating the facilities are subject to economies of scale. The objective is to minimize the total expected costs. These costs can be split into two parts: firstly the costs of investing in a facility as well as maintaining and operating it with strictly diminishing average costs, and secondly linear transportation cost. We formulate the problem as a two-stage stochastic programming model and present a solution method based on Lagrangian Relaxation. We also show some computional results based on data from the Norwegian meat industry regarding the location of slaughterhouses.

Keywords: Facility location;stochastic;economies of scale

Joint work of: Schütz, Peter;Stougie, Leen;Tomasgard, Asgeir

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2005/111

Approximation in Stochastic Scheduling: Models and Algorithms

Marc Uetz (Maastricht University, NL)

We consider scheduling problems where the processing times of jobs are not known in advance, but are realized according to probability distributions. Assuming that this source of uncertainty is purely exogenous, one possible option is to optimize the expected performance of an algorithm. We discuss the stochastic scheduling model, and also briefly contrast the model with other models for optimization under uncertainty, such as online optimization or stochastic programming.

The talk then gives an overview of the results and main ideas in obtaining approximation algorithms for such problems. We end with a brief outlook on a new model, and results, combining the features of online and stochastic scheduling. (Based on joint work with N Megow, RH Möhring, AS Schulz, M Skutella, T Vredeveld)

Average Case and Smoothed Competitive Analysis of the Multi-Level Feedback Algorithm

Tjark Vredeveld (Maastricht University, NL)

In this paper we introduce the notion of smoothed competitive analysis of online algorithms. Smoothed analysis has been proposed by Spielman and Teng to explain the behaviour of algorithms that work well in practice while performing very poorly from a worst case analysis point of view.

We apply this notion to analyze the Multi-Level Feedback (MLF) algorithm to minimize the total flow time on a sequence of jobs released over time when the processing time of a job is only known at time of completion.

The initial processing times are integers in the range $[1, 2^K]$. We use a partial bit randomization model, i.e., the initial processing times are smoothed by changing the k least significant bits under a quite general class of probability distributions.

We show that MLF admits a smoothed competitive ratio of $O((2^k/\sigma)^3 + (2^k/\sigma)^2 2^{K-k})$, where σ denotes the standard deviation of the distribution.

In particular, we obtain a competitive ratio of $O(2^{K-k})$ if $\sigma = \Theta(2^k)$.

We also prove an $\Omega(2^{K-k})$ lower bound for any deterministic algorithm that is run on processing times smoothed according to the partial bit randomization model.

For various other smoothing models, including the additive symmetric smoothing one, which is a variant of the model used by Spielman and Teng, we give a higher lower bound of $\Omega(2^K)$.

A direct consequence of our result is also the first average case analysis of MLF. We show a constant expected ratio of the total flow time of MLF to the optimum under several distributions including the uniform one.

Keywords: Competitive analysis, average case analysis, smoothed analysis, scheduling

Joint work of: Becchetti, Luca; Leonardi, Stefano; Marchetti-Spaccamela, Alberto; Schaefer, Guido; Vredeveld, Tjark

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/75

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2005/75

A simplex algorithm for the solution of separated continuous linear programs.

Gideon Weiss (Haifa University, IL)

Bellman has introduced continuous linear programming problems in the context of economic modeling in 1953. Anderson defined a subset of these as separated continuous linear programs, to approximate some job shop scheduling problems. CLP and SCLP are somewhat in between control problems and linear programs. In spite of continuous effort to solve them, by Dantzig and his students Perold and Anstreicher, and by Anderson, Philpott and Nash, and more recently by Pullan, much remained unknown in the theory of CLP/SCLP, and algorithms were based on solving multiperiod discrete time LP approximations.

We focus on SCLP of the form:
$\max int_0^T((\gamma' + (T-t)c')u(t) + d'x(t))dt$
s.t. $int_0^t Gu(s)ds + Fx(t) <= \alpha + at$,
$Hu(t) <= b$,
$x(t), u(t) >= 0, 0 <= t <= T$.

We shall describe our algorithm which solves this SCLP exactly in a finite number of steps. This is a simplex type algorithm based on:

- SCLP has a symmetric dual SCLP problem running in reverse time.
- Extreme points in functional space are characterized by a sequence of adjacent bases of the rates LP of the problem.
- Extreme points are primal/dual complementary slack optimal in a convex polyhedral cone of primal initial values, dual initial values and time horizons. This vone is called the validity region of the extreme point.
- Neighboring extreme points have common boundary points of their validity regions. We can perform a 'pivot' operation from an extreme point to a neighboring extreme point.
- We can solve SCLP by a finite sequence of such pivots which takes us along a parametric path to the optimum.

Time permitting we will discuss implementation, applications, and extensions.

Online removable square packing

Guochuan Zhang (Zheiiang University - Hangzhou, PRC)

The online removable square packing problem is a two-dimensional version of the online removable Knapsack problem. For a sequence of squares with side length at most one, we are requested to pack a subset of them into a unit square in an online fashion where the online player can decide whether to take the current square or not and which squares currently in the unit square to remove. The goal

is to maximize the total packed area. We first show that any online algorithm cannot achieve a better competitive ratio than 2.618 (even when repacking is allowed). Then the matching upper bound is achieved by a relatively simple online algorithm if repacking is allowed. Without repacking, we can achieve an upper bound of 3 with the help of the concept of bricks by Januszewski and Lassak (1997).

Keywords: Online algorithms; square packing

Joint work of: Han, Xin; Iwama, Kazuo; Zhang, Guochuan

Online scheduling of splittable tasks

Rob van Stee (Universität Karlsruhe, D)

We consider online scheduling of splittable tasks on parallel machines. In our model, each task can be split into a limited number of parts, that can then be scheduled independently.

We consider both the case where the machines are identical and the case where some subset of the machines have a (fixed) higher speed than the others.

We design a class of algorithms which allows us to give tight bounds for a large class of cases where tasks may be split into relatively many parts.

For identical machines we also improve upon the natural greedy algorithm in other classes of cases.

Keywords: Online scheduling; splittable tasks; parallel machines; greedy algorithm

Joint work of: Epstein, Leah; van Stee, Rob

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/74

Modification of Recourse Data for Mixed-Integer Recourse Models

Maarten H. van der Vlerk (University of Groningen, NL)

We consider modification of the recourse data, consisting of the second-stage parameters and the underlying distribution, as an approximation technique for solving two-stage recourse problems. This approach is applied to several specific classes of mixed-integer recourse problems; in each case, the resulting recourse problem is much easier to solve.

Keywords: Stochastic programming, integer programming, approximation

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/73