

KEYNOTE PRESENTATIONS

Relaxations of Approximate Linear Programs

Francois Margot (CMU)

The management of facilities for storing and distributing natural gas gives rise to intractable Markov decision processes (MDPs). This intractability is due primarily to the need to represent future price evolution, as the models commonly used in practice for this are high dimensional. We use an Approximate Linear Programming (ALP) approach to these problems and point to one of its major deficiency. To address this weakness, we propose a framework that can be used to derive ALP relaxations and show that several existing approaches are specific realizations of the framework.

We compare empirical lower and upper bounds on the value of an optimal policy on natural gas storage instances for a few ALP relaxations. They significantly outperform their corresponding ALPs. The best ALP relaxation among those tested is competitive with state-of-the-art methods for computing lower bounds and is much more efficient for upper bound estimation.

The proposed framework is potentially relevant for the approximate solution of other MDPs that arise in the real option management, as well as the valuation of real and financial options that depend on forward curve dynamics.

Joint work with S. Nadarajah and N. Secomandi

Facility Location in Evolving Metrics

Claire Mathieu (ENS Paris)

Understanding the dynamics of evolving social or infrastructure networks is a challenge in applied areas such as epidemiology, viral marketing, and urban planning. During the past decade, data has been collected on such networks but has yet to be analyzed fully. We propose to use information on the dynamics of the data to find stable partitions of the network into groups. For this purpose, we introduce a time-dependent, dynamic version of the facility location problem, which includes a switching cost when a client's assignment changes from one facility to another. This might provide a better representation of an evolving network, emphasizing the abrupt change of relationships between subjects rather than the continuous evolution of the underlying network. We show for some realistic examples that this model indeed yields better hypotheses than its counterpart without switching costs, where each snapshot can be optimized independently. We present approximation algorithms and hardness results

This is joint work with David Eisenstat and Nicolas Schabanel

Optimizing Air Liquide's Supply Chains

Nicoleta Neagu (Air Liquide)

Air Liquide, world leader in gases for industry, health, and environment, produces and distributes gases worldwide. In this talk I will present a few supply chain problems that arise in this context.

I will first present problems related to the bulk supply chain. Specifically, we seek to optimize gas production planning, inventory management, and distribution. Secondly, I will present problems occurring in the packaged gases supply chain; how to jointly optimize facility locations, filling tools allocation, inventory management, and transportation.

I will show how applications of OR methods allowed to substantially improve Air Liquide's supply chain efficiency and decision making quality.

The Decision Rule Approach to Stochastic Programming

Daniel Kuhn (EPFL)

Stochastic programming is a dominant approach to address practical optimization problems affected by uncertainty. Most of the existing solution techniques rely on some form of discretization of the underlying probability distribution or state space. While these discretization schemes theoretically achieve any desired level of accuracy, they may suffer from a curse of dimensionality that restricts their application to small and medium-sized problem instances. Recently, an alternative solution paradigm has gained popularity, which preserves the exact distribution of the uncertain parameters but restricts the set of feasible recourse decisions to those possessing a simple functional form, such as linear, piecewise linear or polynomial decision rules. Under this approximation, linear stochastic programs can be reformulated as standard conic optimization problems by using duality techniques. An attractive feature of this decision rule approach is that it typically leads to polynomial-time solution schemes. This talk surveys some of the main theoretical results on decision rules such as their scalability properties, bounds on their approximation quality as well as systematic techniques to improve their approximation quality by augmenting the dimension of the underlying probability space. Even though decision rule approximations have gained broader attention only since 2004, they have already found successful use in a variety of application areas ranging from operations management, portfolio optimization and automatic control to network design and energy planning. By taking a closer look at numerical results obtained in some of these application areas, we will showcase the potential of decision rule-based solution schemes to enable scalability to industry-sized problem instances.

CONTRIBUTED PRESENTATIONS

A polyhedral Frobenius theorem and applications to non-linear variable-dimension integer optimization

David Adjiashvili (ETHZ)

We prove a new representation theorem of projections of sets of integer points by an integer matrix W . Our result can be seen as a polyhedral analogue of several classical and recent results related to the Frobenius problem. Our result is motivated by a large class of non-linear integer optimization problems in variable dimension. Concretely, we aim to optimize $f(Wx)$ over a set $F = P \cap \mathbb{Z}^d$, where f is a non-linear function, $P \subset \mathbb{R}^d$ is a polyhedron and $W \in \mathbb{Z}^{n \times d}$. As a consequence of our representation theorem, we obtain a general efficient transformation from the latter class of problems to integer linear programming. Our bounds depends polynomially on various important parameters of the input data leading, among others, to first polynomial time algorithms for several classes of non-linear optimization problems. This is joint work with Timm Oertel and Robert Weismantel

Generation and evaluation of passenger-oriented railway disposition timetables

Stefan Binder (EPFL)

As delays are one of the main sources of passenger dissatisfaction in public transportation networks, railway companies put considerable effort in trying to avoid them. However, on a daily basis, delays occur for a number of reasons, e.g., a jammed door at a station or a temporarily unavailable track. Once an initial delay has occurred, the original timetable needs to be updated to a so-called disposition timetable. The disposition management problem solves the decision whether a train should wait for a delayed feeder train or not. This new timetable has to be conflict-free in terms of operational constraints (e.g., trains cannot use the same track section at the same time) and as convenient as possible for the passengers. The problem of finding a conflict-free disposition timetable has received a huge attention in the scientific literature in recent years, both at a macroscopic and at a microscopic level. The literature focusing on minimizing the negative effects of delays for passengers has also been developing but is much sparser. We propose a model that describes the railway recovery problem in a way that takes passengers explicitly into account. Our model focuses mainly on severe disruptions (as opposed to minor disturbances that can be handled by slight modifications to the timetable) and proposes and evaluates several recovery strategies. The evaluation of the recovery strategies (e.g., partial train cancellation, complete train cancellation, train addition, train replacement) provided by the disposition timetable is based on a number of passenger satisfaction indicators, such as total arrival delay, number of maintained connections and overall number of connections to be made. This model will assist train operating companies when evaluating the trade-off between economic and infrastructural feasibility of recovery schemes on the one hand side and passenger satisfaction on the other.

Steiner Trees for Functional analysis in Cancer System Biology

Murodzhon Akhmedov (IDSIA)

We aim to develop and improve computational approaches for functional analysis of cancer genomic data based on Steiner Tree approach. In recent years there has been important progress in new high-throughput technologies such as microarrays and next generation sequencers. It is assumed that cancer initiation and progression is not caused by single gene, but due to accumulation of aberrations in different genes. System biology tries to understand the biological functions of genes by studying the interaction between the genes in biological networks. Prize-collecting Steiner tree (PCST) is well-studied and popular problem in Graph Theory and Combinatorial Optimization. Given an undirected graph G , whose nodes are associated with prize and arcs with positive cost, PCST finds a tree by minimizing cost of arcs in a tree and prize of nodes that are left out of a tree. Recently a connection between PCST and biological networks has been identified. Briefly, in a large interaction network PCST attempts to find a neighborhood or sub-network where genetic aberrations are most concerned. For example using gene expression data, one can find a connected neighborhood where many genes are differently expressed. The PCST problem is NP-hard, and existing exact approaches suffer from solving the large network instances. However, the size of biological networks (e.g. gene- gene network) can be very large. Thus, we need fast and efficient heuristics to solve PCST and interpret the biological network. There exists a fast heuristic algorithm for Steiner Tree in the literature. We propose to extend the existing heuristic to solve PCST and enrich the algorithm with an effective leaf pruning procedure. We test our heuristic on real world instances as well as randomly generated networks, and compare the running time of the heuristic with existing approaches. The results show that the extended heuristic has a fast running time that can shed light on large biological networks.

Compact Representations of All Members of an Independence System

Utz-Uwe Haus (ETHZ)

It is well known that (reduced, ordered) binary decision diagrams (BDDs) can sometimes be compact representations of the full solution set of Boolean optimization problems. Recently they have been suggested to be useful as discrete relaxations in integer and constraint programming. We show that for every independence system there exists a top-down (i.e., single-pass) construction rule for the BDD. Furthermore, for packing and covering problems on n variables whose bandwidth is bounded by $O(\log n)$ the maximum width of the BDD is bounded by n . We also characterize minimal widths of BDDs representing the set of all solutions to a stable set problem for various basic classes of graphs. Besides counting all solutions and optimizing a class of nonlinear objective functions that includes separable functions, the results can be applied for effective evaluation of generating functions. (This is joint work with Carla Michini)

Finding small stabilizers for unstable graphs

Adrian Bock (EPFL)

An undirected graph $G = (V, E)$ is stable if its inessential vertices (those that are exposed by at least one maximum matching) form a stable set. We call a set of edges $F \subseteq E$ to be a stabilizer if its removal from G yields a stable graph. In this talk we study the following natural edge-deletion question: given a graph $G = (V, E)$, can we find a minimum-cardinality stabilizer? The problem is shown to be vertex-cover hard. We then prove that there is a minimum-cardinality stabilizer that avoids some maximum matching of G . This insight is used to give efficient approximation algorithms for sparse graphs and for regular graphs.

Hardness of approximation for preemptive job scheduling on parallel machines with precedence constraints.

Ashkan Norouzi Fard (EPFL)

We will present 1.5 Unique Game hardness for preemptive job scheduling on parallel machine with precedence constraints. We also provide some intuition why we believe that this problem is hard to approximate within a factor of 2 . In this problem, we have n jobs and m machines. Each job can be scheduled on any machine, and preemption is allowed. There are some precedence constraints between the jobs and the goal is to minimize the makespan of the schedule.

The Ideal Train Timetabling Problem

Tomas Robenek (EPFL)

The aim of this talk is to analyze and to improve the current planning process of the passenger railway service. In the first part, the state-of-the-art in research is presented. However, given the recent changes in legislature allowing competitors in the railway industry, the current way of planning is not sufficient anymore. The original planning is based on the accessibility/mobility concept provided by one carrier, whereas the competitive market consists of several carriers that are driven by the profit. Moreover, the current practice does not define the ideal timetables and thus it is assumed that they evolve incrementally, based on a historical data (train occupation, ticket sales, etc.). And thus, we introduce a definition of an ideal timetable and apply the scheduled delay concept (to cover the elasticity of the demand). In order to create such timetables, we propose to insert the Ideal Train Timetabling Problem (ITTP) that is solved for each Train Operating Company (TOC) separately, into the planning process. The ITTP approach incorporates the passenger demand in the planning and its aim is to minimize the passengers' total travel time (weighted by the demand). The outcome of the ITTP is the ideal timetables (including connections between the trains, based on the demand), which then serve as

an input for the traditional Train Timetabling Problem (TTP). The TTP takes into account wishes of each TOC (the ideal timetables) and creates global feasible timetable for the given railway network, while minimizing the changes of the TOC's wishes. The ITTP is in line with the new market structure and it can produce both: non-cyclic and cyclic timetables. The model is tested on the data provided by the Israel Railways (IR). The instance consists of a full demand OD Matrix of an average working day in Israel during 2008. The results are compared to the current timetable of IR.

On the a priori Travelling Salesman Problem

Christos Kalaitzis (EPFL)

We will discuss the a priori Travelling Salesman Problem, a stochastic version of the well-known TSP problem. The input instances consist of a standard TSP instance, along with a probabilistic distribution (which might or might not be known) which will determine which of the cities of the original instance the salesman actually has to visit. Our goal is to find one tour which minimizes the expected cost, with respect to the randomness of the input, when this tour is restricted to the cities which will actually have to be visited. We will present a 3.5-approximation algorithm for the special case of the input distribution being the product of one independent distribution per city, and we will discuss ways to achieve better guarantees in the future.

Waste collection routing with time windows and intermediate disposal trips

Iliya Markov (EPFL)

We consider a waste collection problem with time windows and intermediate trips to the disposal facilities during the vehicle tours. It is expressed as a mixed binary linear program, which is solved by a common MIP solver with the goal of minimizing the number of tours and their spatial and temporal costs. The formulation extends the so-called vehicle routing problem with intermediate facilities (VRP-IF) by adding a realistic cost-based objective function, multiple depots, a fixed heterogeneous fleet, accessibility restrictions, and a rest period that depends on when the vehicle started its tour. In addition, a relocation cost term incentivizes, rather than enforcing, the vehicle to return to its original depot. To solve general instances of the problem, we develop a local search heuristic whose performance is evaluated by the optimality gap compared to the exact solution on small instances.

Energy Optimization Driven by Perceived Value of Energy Services

Metin Caliskan (Haute École de Gestion de Genève)

Long-term energy and environmental planning models such as MARKAL, TIMES or OSeMOSYS are optimization models driven by useful demands. Objective function is the minimum of global discounted cost of all the investment, operation and maintenance costs and imported energy agents. Social-MARKAL is an extension to the original MARKAL formulation that includes customer behavior modeled as a virtual technology. Energy savings behavior or technology switch are described as they were technologies, while technical coefficients are obtained using special crafted sociological surveys. We propose to push this approach one step further. In service science, a service acquires value once the consumer perceives the benefits of it. Salient attributes are perceived during the service experience itself. While the neo-classical model premises suppose a perfect economic rationality, in our approach, we consider the perceived value of the ecological behavior. Instead of accounting only cost contributions, the perceived value of behavior is added to the system. Social-MARKAL allows the modeler to include consumer behavior into design of public policies. This new approach can help selecting the most efficient information campaigns in order to achieve behavior based energy savings.

Reverse ranks of integer polyhedra

Yuri Faenza (EPFL)

Chvátal-Gomory (CG) and Split closures are among the most powerful tools for solving integer programs. When applied to a rational polyhedron Q , they lead again to a rational polyhedron. By applying the closure operator iteratively, one eventually obtains the integer hull of Q . The number of repetitions needed is called the CG (resp. Split) rank of Q . It is known that, for each rational polyhedron Q , those ranks are finite, but can be arbitrarily high. But how bad can the rank be if we fix the integer hull? More formally, for an integer polyhedron P , let us define $r^*(P)$ as the supremum of the CG ranks of all rational polyhedra whose integer hull is P . Similarly, define $s^*(P)$ for splits. We call these parameters Reverse ranks. In this talk, we provide a geometric characterization in general dimension of polyhedra with finite r^* and of polyhedra with finite s^* .

Hybrid sampling-based metaheuristics for the Orienteering Problem with Stochastic Travel and Service Times

Vassilis Papapanagiotou (IDSIA)

Stochastic Combinatorial Optimization Problems (SCOPs) are of particular interest because they can model many problems more realistically since in many problems some quantities can change due to unpredictable factors such as traffic or weather conditions. Stochasticity increases the difficulty of constructing algorithms for finding solution compared to the deterministic versions. One of the reasons is that computing the objective function of such problems can be a hard problem or as in our case time-

consuming. In the past, Monte Carlo sampling has been used in order to speed up the computation of the objective function. In this presentation, we study different methods for computing the objective function of the Orienteering Problem with Stochastic Travel and Service Times (OPSTS), with the purpose of constructing faster objective functions than the exact or the pure Monte Carlo sampling based ones, while minimizing the error and keeping the solutions found by the metaheuristic almost the same.

LP-Based Algorithms for Capacitated Facility Location

Hyung-Chan An (EPFL)

Linear programming has played a key role in the study of algorithms for combinatorial optimization problems. In the field of approximation algorithms, this is well illustrated by the uncapacitated facility location problem. A variety of algorithmic methodologies, such as LP-rounding and primal-dual method, have been applied to and evolved from algorithms for this problem. Unfortunately, this collection of powerful algorithmic techniques had not yet been applicable to the more general capacitated facility location problem. In fact, all of the known algorithms with good performance guarantees were based on a single technique, local search, and no linear programming relaxation was known to efficiently approximate the problem. In this talk, we present a linear programming relaxation with constant integrality gap for capacitated facility location. We demonstrate that fundamental concepts from the matching theory, including alternating paths and residual networks, provide key insights that lead to the strong relaxation. Our algorithmic proof of integrality gap is obtained by finally accessing the rich toolbox of LP-based methodologies: we present a constant factor approximation algorithm based on LP-rounding. Our results resolve one of the ten open problems selected by the textbook on approximation algorithms of Williamson and Shmoys. Joint work with Mohit Singh and Ola Svensson.

Operating Room Management under Uncertainty

Antoine Legrain (University of Fribourg)

North American hospitals are renowned worldwide for the outstanding quality of the staff and dispensing of patient care. However, it is hard to deny that there is room for improvement when one of the most often critiqued shortcoming is the speed at which these services are available. While some swear that the answer lies in private management, it can be argued that it is rather in the tools private practices use to balance their workload. The operating room management problems are legion. This talk tackles the scheduling and planning of the workforce of an operating theatre. In order to achieve these results, we first solve a deterministic version that uses the constraint programming paradigm and then a stochastic one which embeds the former in a Monte-Carlo scheme. The work has been done in the context of the 5th AIMMS-MOPTA Optimization Modeling Competition.

Approximating the largest sub-determinant of a matrix

Carsten Moldenhauer (EPFL)

We consider the problem of finding the largest sub-determinant of a matrix. In fact, we consider an even more general problem: given n vectors in d dimensions, choose d such vectors that span the largest volume parallelepiped. These problems are motivated by applications in low rank matrix approximation and data simplification. The approach to tackle this problem is a greedy algorithm. I will present an analysis that yields insight into the geometry of the problem. Further, it is possible to show that the analysis is also tight. In addition, I will present a new, surprisingly simple, hardness of approximation lower bound.