Wuppertal–Coimbra Optimization Workshop

May 3rd - 4th, 2023

University of Coimbra

Participate by zoom: https://dei.uc.pt/vc/workshops

Local Organization:
Carlos Fonseca, Kathrin Klamroth, Luís Paquete

Venue: Department of Informatics Engineering
Pólo II - Pinhal de Marrocos, 3030-290 Coimbra
Room A.5.1

From Ibis Hotel to Pólo II:

Pólo II:

Program for Wednesday, May 3rd, 2023

09:45 – 10:00	A.5.1	Welcome
10:00 – 10:30	A.5.1	Daniela Santos *Algorithms for the maximum quasi-clique problem and variants*
10:30 – 11:00	A.5.1	Lara Löhken *A multi-objective perspective on the Cable-Trench problem*
11:00 – 11:30	Cafeteria	Coffee Break
11:30 – 12:00	A.5.1	Onur Tanil Doganay *Non-convex shape optimization by dissipative Hamiltonian flows*
12:00 – 13:00	Student Restaurant	Lunch Break
13:00 – 13:30	A.5.1	Gonçalo Lopes *Greedy hypervolume representation in higher dimensions*
13:30 – 14:00	A.5.1	Hannah Borgmann *Algorithms and approximation results for variants of the Rectangular Knapsack Problem*
14:00 – 14:30	Cafeteria	Coffee Break
14:30 – 15:00	A.5.1	Julius Bauß *Augmenting multi-objective Branch and Bound by scalarization-based information*
15:00 – 15:30	A.5.1	Felipe Mota *Ranking strategies for the biobjective minimum spanning tree*
15:30 – 16:00	Cafeteria	Coffee Break & Group Photo
16:00 – 17:30	A.5.1	Group Discussion
19:30 –	Pizzeria Itália Parque de Coimbra 3000	Workshop Dinner

Program for Thursday, May 4th, 2023

10:00 – 10:30	A.5.1	João Almeida *A multi-objective mixed integer linear programming model for thesis defence scheduling*
10:30 – 11:00	A.5.1	Julia Sudhoff *Interrelation between multiobjective shortest paths and safest paths*
11:00 – 12:00	Cafeteria	Coffee Break and Group Discussion
12:00 – 13:00	Student Restaurant	Lunch Break
13:00 – 13:30	A.5.1 via zoom	François Clément *Subset selection to build low-discrepancy point sets: an MILP approach*
13:30 – 14:00	A.5.1	Daniela Gaul *Event-based MILP models for ridepooling applications*
14:00 – 14:30	Cafeteria	Coffee Break
14:30 – 15:00	A.5.1	Michael Stiglmayr *Efficient dominance filtering for unions and Minkowski sums of non-dominated sets*
15:00 – 15:30	A.5.1	Britta Schulze *Sensitivity analysis of the cost coefficients in multi-objective integer linear optimization*
15:30 –	A.5.1	Closing Discussion

Abstracts

A multi-objective mixed integer linear programming model for thesis defence scheduling

João Almeida
IST Lisbon

In this paper, we address the thesis defence scheduling problem, a critical academic scheduling management process, which has been overshadowed in the literature by its counterparts, course timetabling and exam scheduling. Specifically, we address the single defence assignment type of thesis defence scheduling problems, where each committee is assigned to a single defence, scheduled for a specific day, hour and room. We formulate a multi-objective mixed-integer linear programming model, which aims to be applicable to a broader set of cases than other single defence assignment models present in the literature, which have a focus on the characteristics of their universities. For such a purpose, we introduce a different decision variable, propose constraint formulations that are not regulation and policy specific, and cover and offer new takes on the more common objectives seen in the literature. We also include new objective functions based on our experience with the problem at our university and by applying knowledge from other academic scheduling problems. We also propose a two-stage solution approach. The first stage is employed to find the number of schedulable defences, enabling the optimisation of instances with unschedulable defences. The second stage is an implementation of the augmented e-constraint method, which allows for the search of a set of different and non-dominated solutions while skipping redundant iterations. The methodology is tested for case-studies from our university, significantly outperforming the solutions found by human schedulers. A novel instance generator for thesis scheduling problems is presented. Its main benefit is the generation of the availability of committee members and rooms in availability and unavailability blocks, resembling their real-world counterparts. A set of 96 randomly generated instances of varying sizes is solved and analysed regarding their relative computational performance, the number of schedulable defences and the distribution of the considered types of iterations. The proposed method can find the optimal number of schedulable defences and present non-dominated solutions within the set time limits for every tested instance.

Augmenting multi-objective Branch and Bound by scalarization-based information

Julius Bauß
University of Wuppertal

Although Branch and Bound based methods are the standard approach for solving (mixed-)integer optimization problems, in the multi-objective case Branch and Bound is rarely used in comparison to the predominant objective space methods. This is mainly due to two reasons: objective space methods benefit from optimized single-objective solvers, while bounding is considerably weaker in multiple objectives.

We present improvements to increase the performance of multi-objective Branch and Bound by using objective space information to limit the impact of these shortcomings in multiple objectives. We introduce a dynamic branching strategy based on the hypervolume indicator as a measurement for the gap between lower and upper bound set. By adaptively solving scalarizations in the root node to integer optimality, lower and upper bound set can be improved. The obtained information may be integrated into the lower bound set of all remaining active nodes, while the determined effcient solutions are integrated into the upper bound set. This results in a decomposition of the upper bound set into the incumbent points and the already proven non-dominated points. The improved bound sets lead to an increased fathoming rate and therefore to a decrease of the created nodes. We present preliminary numerical results showing that these improvements also have a significant impact on the total computation time.

Algorithms and approximation results for variants of the Rectangular Knapsack Problem

Hannah Borgmann
RPTU Kaiserslautern-Landau

We look at variants of the Cardinality Constrained Rectangular Knapsack Problem, where the cardinality constraint $\sum_{i=1}^{n} x_i \leq \kappa$ is replaced by constraints of the form $\sum_{v \in V(F)} x_v \leq \kappa_F$ for some connected subgraphs F of a given graph G. We explore several types of graphs G and subgraphs F. In particular, we consider the problem where G is a path and the constraints bound the number of nodes that can be chosen from some subpaths of G, or the problem where at most one node can be chosen from each edge in G, i.e. any feasible solution has to be an independent set. Although some of these problems are polynomially solvable when the objective function is linear, we show that the they become NP-hard when a rectangular objective function is considered. Lastly, we investigate the possibility of developing pseudo-polynomial algorithms and FPTASes for the problems that are solvable in polynomial time in the linear version.

Subset selection to build low-discrepancy point sets: an MILP approach

François Clément
Sorbonne University, Paris

Discrepancy measures are metrics designed to characterise how regularly a point set is distributed in a given space. Constructing low-discrepancy point sets has been the focus of mathematicians' work since the 1950's, but these sets are often designed for asymptotic properties, where the number of points considered tends to infinity. Tailoring low-discrepancy sets to practical applications with a finite number of points is an open problem from moderate dimensions onwards (roughly greater than 6). In this talk, I will introduce the Star Discrepancy Subset Selection, which consists in choosing from a set P of n points the subset P_k of size k smaller than n such that the discrepancy of P_k is minimized. In particular, we focus here on the L_∞-star discrepancy, arguably the most important measure. While this problem is NP-hard, we are able to formulate an MILP that can solve it. However, our experiments suggest current solvers can only provide solutions in a reasonable time in dimension 2 for up to $n = 140$. A second approach based on a feasibility version of the previous MILP will also be presented.

Non-convex shape optimization by dissipative Hamiltonian flows

Onur Tanil Doganay
University of Wuppertal

Shape optimization with constraints given by partial differential equations (PDE) is a highly developed field of optimization theory. The elegant adjoint formalism allows to compute shape gradients at the computational cost of a further PDE solve. Thus, gradient descent methods can be applied to shape optimization problems. However, gradient descent methods that can be understood as approximation to gradient flows get stuck in local minima, if the optimization problem is non-convex. In machine learning, the optimization in high dimensional non-convex energy landscapes has been successfully tackled by momentum methods, which can be understood as passing from gradient flow to dissipative Hamiltonian flows. In this article, we adopt this strategy for non-convex shape optimization. In particular, we provide a mechanical shape optimization problem that is motivated by optimal reliability considering also material cost and the necessity to avoid certain obstructions in installation space. We then show how this problem can be solved effectively by port Hamiltonian shape flows.

Event-based MILP models for ridepooling applications

Daniela Gaul
University of Wuppertal

Ridepooling services require efficient optimization algorithms to simultaneously plan routes and pool users in shared rides. We consider a static dial-a-ride problem (DARP) where a series of origin destination requests have to be assigned to routes of a fleet of vehicles. Thereby, all requests have associated time windows for pick-up and delivery, and may be denied if they can not be serviced in reasonable time or at reasonable cost. Rather than using a spatial representation of the transportation network we suggest an event-based formulation of the problem: nodes in the event-based graph represent feasible allocations of users to vehicles and arcs describe the transition from one user allocation to another. The event-based formulation results in significantly improved computational times. While the corresponding MILP requires more variables than standard models, it has the advantage that capacity, pairing and precedence constraints are handled implicitly. The approach is tested and validated using a standard IP-solver on benchmark data from the literature.

A multi-objective perspective on the Cable-Trench Problem

Lara Löhken
University of Wuppertal

The Cable-Trench Problem is defined as a combination of the Shortest Path and the Minimum Spanning Tree Problem. In particular, the goal is to find a spanning tree that simultaneously minimizes its total length and the total path length from a pre-defined starting vertex to all other vertices. Both, the Minimum Spanning Tree and the Shortest Path Problem are known to be efficiently solvable. However, a linear combination of these two objectives results in a highly complex problem.

While in the original publication by Vasko et al. (2002) the Cable-Trench Problem is introduced as a single-objective problem, namely a linear combination of Minimum Spanning Tree and the Shortest Path objective, we consider it as a bi-objective problem separating the two cost functions. In this sense the original form of the Cable-Trench Problem corresponds to a weighted sum scalarization of the bi-objective formulation. We show that in general the bi-objective approach may imply additional compromise solutions that cannot be found by solving the Cable-Trench Problem in its originally formulation. In order to determine the set of non-dominated points and efficient solutions, we use ϵ-constraint scalarizations in combination with problem specific cutting planes which can be adapted from the fixed-charge network flow problem. We show preliminary numerical results.

Greedy hypervolume representation in higher dimensions

Gonçalo Lopes
University of Coimbra

In multiobjective combinatorial optimization, the goal is to solve an optimization problem with several objective functions and find the nondominated set. However, the nondominated set is usually quite large and can take some time to retrieve, particularly when the number of objectives increases. Therefore, it is necessary to rapidly obtain a representation of the nondominated set that satisfies a given metric of interest.

This work aims to compute a representation of the nondominated set that maximizes the dominated hypervolume with respect to a given reference point. We mimic a greedy strategy that iteratively uses the hypervolume indicator as a scalarizing function for m-objective combinatorial problems.

In particular, we describe a generic solution approach that determines a subset of the nondominated set of an m-objective optimization problem by solving a sequence of hypervolume scalarizations with appropriate choices of local bounds. Moreover, this solution approach can be modified to use at most k local bounds to approximate the hypervolume scalarizations. We illustrate these concepts for the m-objective knapsack problem.

Ranking strategies for the biobjective minimum spanning tree

Felipe Mota
University of Coimbra

Multiobjective Combinatorial Optimization is concerned with solving optimization problems with more than one conflicting objective. In these problems, a feasible solution is efficient if there exists no other feasible solution that is better or equal in all objectives. Since there might be more than one efficient solution, the goal is to find the set of all the efficient solutions, called the efficient set, and/or its image in the objective space, called the non-dominated set. One common approach to these problems is to divide the search into two phases. The first phase takes into account that an optimal solution for a weighted sum is also an efficient solution for the multiobjective problem, which allows obtaining a subset of the efficient set by solving a sequence of weighted sum problems. The second phase consists of finding the remaining solutions by some enumeration method in the search region delimited by thesolutions found in the first phase. In the second phase, it is possible to use ranking methods, which enumerate the k-best solutions for a given weighted sum. However, this enumeration is repeated very often in the second phase, leading to large computation times due to the overlapping of solutions that are found. Our work aims to reduce redundancy by grouping search regions in the second phase and running fewer iterations of the ranking algorithm. We use the Biobjective Minimum Spanning Tree Problem to compare the performance of several grouping strategies.

Algorithms for the maximum quasi-clique problem and variants

Daniela Santos
University of Coimbra

Given an undirected graph G, a quasi-clique is a subgraph of G whose density is at least γ ($0 < \gamma \leq 1$). Two optimization problems can be defined for quasi-cliques: the Maximum quasi-clique problem (MQC), which asks for a quasi-clique with maximum vertex cardinality, and the Densest $k-$subgraph problem (DKS), which seeks the densest subgraph given a cardinality. Many exact and heuristic approaches have been proposed to address these problems. However, most of them neglect the connectedness property of the solutions, leading to quasi-cliques with isolated components that are useless for many real-life applications. Experiments on sparse graphs confirm that existing mixed integer linear programming formulations may produce disconnected subgraphs for both problems. To address this issue, two flow-based connectedness constraints are proposed and integrated into these formulations. The first ensures connectedness by using a spanning tree characterization, whereas the second uses a classic flow-based approach. The performance of the formulations enhanced with the proposed connectedness constraints is compared in terms of running time to existing exact methods, namely the Connected$_\mu$ algorithm [1] and Marinelli's connectedness constraints [2]. Results for the MQC problem indicate that our first constraints allow the formulation to solve the problem with a connectedness guarantee faster than the compared methods, whereas, for the DKS problem, the second one shows superior performance.

[1] Komusiewicz, C., Sorge, M., Stahl, K.: Finding connected subgraphs of fixed minimum density: Implementation and experiments. In: E. Bampis (ed.) Experimental Algorithms, pp. 82-93. Springer International Publishing, Cham (2015)
[2] Marinelli, F., Pizzuti, A., Rossi, F.: LP-based dual bounds for the maximum quasi-clique problem. Discrete Applied Mathematics 296, 118-140 (2021)

Sensitivity analysis of the cost coefficients in multi-objective integer linear optimization

Britta Schulze
University of Wuppertal

We consider sensitivity analysis of the cost coefficients in multi-objective integer linear programming problems and define the sensitivity region as the set of simultaneous changes to the objective function coefficients for which the efficient set and its structure remain the same. In particular, we require that the component-wise relation between efficient solutions is preserved and that inefficient solutions remain inefficient. In this talk we concentrate on changes to a single objective function coefficient and show that the sensitivity region is a convex set, i.e., an interval.

Efficient dominance filtering for unions and Minkowski sums of non-dominated sets

Michael Stiglmayr
University of Wuppertal

The repeated filtering of vectors for non-dominance is an important component in many multi-objective programming approaches, like e. g. decomposition approaches, dynamic programming or meta heuristics. Often the set of vectors to be filtered is given as the union $A \cup B$ or Minkowski sum $A+B$ of Pareto (or stable) sets, *i. e.* within both sets A and B the vectors are pairwise non-dominated. We propose several algorithms for both problems and compare them to a well-known static divide-and-conquer non-dominance filtering algorithm.

Interrelation between multiobjective shortest paths and safest paths

Julia Sudhoff
University of Wuppertal

The safety of a street for a cyclist cannot be measured by numerical values but it can be associated with ordered categories, like, e.g., safe (separate bike lane), medium safe (street with a bike lane) and unsafe (street without a bike lane). Such a problem with ordered categories is also called a problem with ordinal costs. In this talk we investigate the ordinal ordering cone and its interrelation with the Pareto cone. We show that the safest path problem can be transformed into a multiobjective shortest path problem with Pareto optimality, binary costs and a specific structure. Moreover, we analyze how the ordering cone changes, if we assume for example that each category is at least twice as good as the next-lower category.