

Follow-Up-Workshop to TP
“Discrete Optimization”

April 20 to April 24, 2026

organized by

Daniel Dadush (CWI & Utrecht University), Jesper Nederlof (Utrecht University), Neil Olver (LSE), Laura Sanità (Eindhoven University), László Végh (Universität Bonn)

Abstracts

Xavier Allamigeon (Inria and Ecole Polytechnique)

Tropical vs Classical Oriented Matroids

Abstract: I will report on progress on the problem of separating tropical and classical oriented matroids. This includes new results on the enumeration of coherent matching fields for small rank and size of the ground set as well as asymptotic bounds. Joint work with Yiyuan Chen, Stéphane Gaubert and Xavier Goaoc.

Kristóf Bérczi (Eötvös University)

Fixed-Parameter Tractability and Hardness for Steiner Rooted and Locally Connected Orientations

Abstract: Finding a Steiner strongly k -arc-connected orientation is particularly relevant in network design and reliability, as it guarantees robust communication between a designated set of critical nodes. We consider a rooted variant, called the Steiner Rooted Orientation problem, where one is given an undirected graph on n vertices, a root vertex, and a set of t terminals. The goal is to find an orientation of the graph such that the resulting directed graph is Steiner rooted k -arc-connected. While the maximum k for which a Steiner strongly k -arc-connected orientation exists can be determined in polynomial time via Nash-Williams’ orientation theorem, its rooted counterpart is significantly harder: the problem is NP-hard when both k and t are part of the input. In this talk, we provide a complete understanding of the problem with respect to these two parameters, thereby resolving an open problem of Király and Lau (FOCS 2006). In particular, we give an algorithm that solves the problem in time $f(k, t) \cdot n^{O(1)}$, establishing fixed-parameter tractability with respect to the number of terminals t and the target connectivity k . We further show that the problem remains NP-hard if either k or t is treated as part of the input, meaning that our algorithm is essentially optimal from a parameterized perspective.

Joint work with Florian Hörsch, András Imolay, and Tamás Schwarcz.

Jan van den Brand (Georgia Tech)

The Structural Complexity of Matrix-Vector Multiplication

Abstract: In the OMv-problem, we preprocess an $n \times n$ matrix M , and then support queries that for any vector v return the matrix-vector product Mv faster than naive matrix-vector-multiplication. From a theory perspective, this problem is hard and has become a cornerstone of fine-grained complexity theory for proving (conditional) lower bounds for data structures/dynamic algorithms. For general fields, it was proven that the problem cannot be solved faster than naive multiplication in the worst-case, and the parity version has average-case hardness. Yet, practitioners developed heuristics that are efficient in practice, despite average-case hardness of the problem. We explain this gap between theory and practice by studying the problem for *structured* matrices. We show that whenever the matrix is close to having any monotone property, then matrix-vector multiplication can be solved in subquadratic time. We show how to use this as a tool for constructing subquadratic upper bounds for various dynamic graph problems on structured graphs. Joint work with: Emile Anand and Rose McCarty

Karthik Chandrasekaran (University of Illinois)

Hedgegraph Polymatroids

Abstract: Graphs and hypergraphs combine expressive modeling power with algorithmic efficiency for a wide range of applications. Hedgegraphs generalize hypergraphs further by grouping hyperedges under a color/hedge. This allows hedgegraphs to model dependencies between hyperedges and leads to several applications. However, it poses algorithmic challenges. In particular, the cut function is not submodular, which has been a barrier to algorithms for connectivity. In this talk, I will introduce two alternative partition-based measures of connectivity in hedgegraphs and focus on their structural and algorithmic aspects. We will take a polymatroidal viewpoint of hedgegraphs. The polymatroidal lens leads to new tractability results as well as insightful generalizations of classical results on graphs and hypergraphs.

Gerard Cornuejols (Carnegie Mellon University)

Younger's Family of Minimally Non-Packing Clutters and Beyond

Abstract: Younger proposed an infinite family of counterexamples to the Edmonds-Giles conjecture. We revisit these examples. This work is joint with Ahmad Abdi and Mahsa Dalirrooyfard. This research was initiated at the workshop in Discrete Optimization when Ahmad Abdi and Gerard Cornuejols were visitors at the HIM.

Fabrizio Grandoni (SUPSI)

A Framework for k-Median and k-Means Approximation

Abstract: I will describe an algorithmic framework that led to improved approximation algorithms for k-Median and k-Means clustering.

Hung Hoang (TU Wien)

A Near-Complete Resolution of the Exponential-Time Complexity of k-Opt for TSP

Abstract: The k-opt algorithm is one of the simplest and most widely used heuristics for solving the travelling salesman problem (TSP). Starting from an arbitrary tour, the k-opt algorithm iteratively improves the current tour in each iteration by exchanging up to k edges, until no such improvement is possible. Krentel (FOCS 1989) showed that the k-opt algorithm can have exponential running time for any pivot rule. However, his proof requires $k \gg 1000$. In this talk, I will show that this property holds

for all $k \geq 3$; that is, for almost all values of k , an exponential number of iterations may be needed even if an optimal pivot rule is used. This result holds for both the general and metric TSP.

Joint work with Sophia Heimann and Stefan Hougardy.

Stefan Hougardy (Universität Bonn)

Fast Algorithms for Euclidean Perfect Matching

Abstract: We study the Euclidean minimum weight perfect matching problem for n points in the plane. Any deterministic approximation algorithm with approximation ratio depending only on n requires $\Omega(n \log n)$ time. We present an $O(n \log n)$ algorithm achieving an improved approximation ratio of $O(n^{0.206})$. Our result extends to higher fixed dimensions.

Sophie Huiberts (LIMOS)

Analyzing the Simplex Method By-the-Book

Abstract: The simplex method is an algorithm for linear programming, and this algorithm is much faster than theory is able to explain. In this talk I will describe a new methodology we introduced to address this question. We prove new strong running time guarantees, using mathematical assumptions based on software user manuals and source codes. I will discuss which features of real-world software and LP's we have managed to theoretically capture for this purpose, and what will come next.

Klaus Jansen (Christian-Albrechts-Universität zu Kiel)

A Tight Double-Exponentially Lower Bound for High-Multiplicity Bin Packing

Abstract: Consider a high-multiplicity Bin Packing instance I with d distinct item types. In 2014, Goemans and Rothvoss gave an algorithm with runtime $|I|^{2^{O(d)}}$ for this problem [SODA'14], where $|I|$ denotes the encoding length of the instance I . Although, Jansen and Klein [SODA'17] later developed an algorithm that improves upon this runtime in a special case, it has remained a major open problem by Goemans and Rothvoss [J.ACM'20] whether the doubly exponential dependency on d is necessary. We solve this open problem by showing that unless the ETH fails, there is no algorithm solving the high-multiplicity Bin Packing problem in time $|I|^{2^{o(d)}}$. To prove this, we introduce a novel reduction from 3-SAT. The core of our construction is efficiently encoding the entire information from a 3-SAT instance with n variables into an ILP with $O(\log(n))$ variables.

This result confirms that the Goemans and Rothvoss algorithm is best-possible for Bin Packing parameterized by the number d of item sizes.

This is joint work with Felix Ohnesorge and Lis Piroton.

Georg Loho (FU Berlin)

Many Rays of the Submodular Cone

Abstract: The study of the cone of submodular functions goes back to Jack Edmonds' seminal 1970 paper, which already highlighted the difficulty of characterizing its extreme rays. Since then, researchers from diverse fields have sought to characterize, enumerate, and bound the number of such rays. We introduce an inductive construction that generates new rays of the submodular cone. This allows us to establish that the n -th submodular cone has at least $2^{2^{n-2}}$ rays, which improves upon the lower bound obtained from Hien Q. Nguyen's 1986 characterization of indecomposable matroid polytopes by a factor of order $n^{3/2}$ in the exponent.

Haoyuan Ma (University of Bonn)

Trust Region Interior Point Methods: Optimal L2- and Faster Wide-Neighborhood Path Following

Abstract: We present improved running time and iteration complexities of interior point methods for linear programs parametrized by the straight line complexity, i.e., the minimum number of segments of any piecewise linear curve traversing a particular neighborhood of the central path. While the standard measure of progress is the reduction in duality gap, the straight line complexity provides a stronger instance-wise bound, reflecting the combinatorial structure of the problem.

Our first main result focuses on interior point methods that stay in the ℓ_2 -neighborhood. We give a much stronger analysis of the trust region interior point method introduced by Lan, Monteiro and Tsuchiya (SIAM J. Optim. 2009), proving that it is approximately instance optimal in this neighborhood. Namely, we show that the iteration complexity of this algorithm is a constant factor of the straight line complexity of the ℓ_2 -neighborhood. Further, each iteration can be implemented in current matrix multiplication time.

Our second main result is a wide-neighborhood interior point method whose running time is the wide-neighborhood straight line complexity times current matrix multiplication time, improving in essence a factor n over the algorithm by Allamigeon, Dadush, Loho, Natura, and Véggh (SIAM J. Comput. 2025). The algorithm can be seen as a boosted version of the robust interior point methods of Cohen, Lee and Song (JACM 2021) and van den Brand (SODA 2020) that can reduce the gap by a polynomial factor in current matrix multiplication time: our algorithm is also able to traverse any sufficiently straight segment of the central path in current matrix multiplication time, independently of the length of the segment.

A main ingredient in both methods is to solve trust region problems with ℓ_2 and ℓ_∞ -constraints, respectively. We develop fast and strongly polynomial algorithms for solving them to high accuracy. In the ℓ_2 -setting, this answers an open question by Lan, Monteiro and Tsuchiya. This is joint work with Daniel Dadush, Bento Natura, and László Véggh.

Jannik Matuschke (KU Leuven)

Stronger Hardness for Maximum Robust Flow and Randomized Network Interdiction

Abstract: We study the following fundamental network optimization problem known as Maximum Robust Flow (MRF): A planner determines a flow on s-t-paths in a given capacitated network. Then, an adversary removes k arcs from the network, interrupting all flow on paths containing a removed arc. The planner's goal is to maximize the value of the surviving flow, anticipating the adversary's response (i.e., a worst-case failure of k arcs). It has long been known that MRF can be solved in polynomial time when $k = 1$ (Aneja et al., 2001), whereas it is NP-hard when k is part of the input (Disser and Matuschke, 2020). However, the complexity of the problem for constant values of $k > 1$ has remained elusive, in part due to structure of the natural LP description preventing the use of the equivalence of optimization and separation.

This paper introduces a reduction showing that the basic version of MRF described above encapsulates the seemingly much more general variant where the adversary's choices are constrained to k -cliques in a compatibility graph on the arcs of the network. As a consequence of this reduction, we are able to prove the following results: (1) MRF is NP-hard for any constant number $k > 1$ of failing arcs. (2) When k is part of the input, MRF is $P^{NP[\log]}$ -hard. (3) The integer version of MRF is Σ_2^P -hard.

Link to paper: <https://arxiv.org/abs/2511.06505>

Nicole Megow (University of Bremen)

The Power of Two Oracles: Minimum Spanning Tree and Matroid Optimization

Abstract: The performance of learning-based algorithms benefits from high-quality predictions, while remaining robust under unreliable predictions. Querying large models such as neural networks can provide highly accurate predictions, but often at significant computational cost. At the same time, lightweight heuristics or smaller models can return approximate information much faster, albeit with reduced reliability. In this talk, we introduce a two-oracle model that formalizes this setting by giving algorithms access to both a fast but noisy oracle and a slow but accurate one, enabling them to balance efficiency and solution quality. We discuss this model in the context of matroid optimization problems, which generalize many classical problems in combinatorial optimization. Our focus is on the maximum-weight basis problem and its special case, the minimum spanning tree. We analyze scenarios involving two different kinds of information: a structural oracle (the independence oracle) and a numerical oracle (weights of elements). We design algorithms that make provably few calls to the clean oracle while remaining robust against arbitrarily poor dirty oracles, thereby approaching the performance of classic algorithms.

Matthias Mnich (Universität Hamburg)

Subexponentially Faster Quasi-Random Walks

Abstract: The rotor-router model of Propp and Spencer is a deterministic analogue of random walks. It considers a directed graph where each vertex carries some tokens; each token is deterministically routed to the outneighbor which is prescribed by a rotor and which moves to its next position in a cyclic order for the next token. This mechanism ensures that all tokens are routed highly evenly among each vertex' neighbors, similar to (and in fact better than) what a random walk would have done. A key question about such quasi-random walks is to understand the final distribution of tokens at which the routing process stabilizes. The complexity of this question is known to lie in $\text{NP} \cap \text{coNP}$, yet a polynomial-time algorithm has remained elusive. The fastest known algorithms apply either recursive methods or Tarski fixed points, both yielding subexponential-time algorithms.

We will present subexponentially faster algorithms for general quasi-random walks, and quasi-polynomial-time algorithms for sparse and dense instances.

Sarah Morell (University of Bremen)

Unsplittable Transshipments

Abstract: We introduce the Unsplittable Transshipment Problem in directed graphs with multiple sources and sinks. An unsplittable transshipment routes given supplies and demands using at most one path for each source-sink pair. Although they are a natural generalization of single source unsplittable flows, unsplittable transshipments raise interesting new challenges and require novel algorithmic techniques. As our main contribution, we give a nontrivial generalization of a seminal result of Dinitz, Garg, and Goemans by showing how to efficiently turn a given transshipment x into an unsplittable transshipment y with $y(a) < x(a) + d_{\max}$ for all arcs a , where d_{\max} is the maximum demand (or supply) value.

This is joint work with Srinwanti Debgupta and Martin Skutella.

Bento Natura (Columbia University)

Circuit Diameter of Polyhedra Is Strongly Polynomial

Abstract: We prove a strongly polynomial bound on the circuit diameter of polyhedra, resolving the circuit analogue of the polynomial Hirsch conjecture. Specifically, we show that the circuit diameter of a polyhedron $P = \{x \in \mathbb{R}^n : \mathbf{A}x = b, x \geq \mathbf{0}\}$ with $\mathbf{A} \in \mathbb{R}^{m \times n}$ is $O(m^2 \log m)$. Our construction yields monotone circuit walks, giving the same bound for the monotone circuit diameter.

The circuit diameter, introduced by Borgwardt, Finhold, and Hemmecke (SIDMA 2015), is a natural relaxation of the combinatorial diameter that allows steps along circuit directions rather than only along edges. All prior upper bounds on the circuit diameter were only weakly polynomial. Finding a circuit augmentation algorithm that matches this bound would yield a strongly polynomial time algorithm for linear programming, resolving Smale’s 9th problem.

Alantha Newman (CNRS, Laboratoire G-SCOP)

Clustering and Ranking

Abstract: The pivot algorithm was developed as a tool for ranking problems and was then applied to correlation clustering. In this talk, we go in the other direction: we show how some tools developed for clustering problems can be applied to generate rankings and discuss an application for local-to-global theorems for coloring tournaments.

Neil Olver (London School of Economics and Political Science)

Thin Trees for Structured Families

Abstract: The thin tree conjecture (in one form) states that for any k -connected graph, there exists a spanning tree that is “thin” in the sense that on every cut in the graph, only an $O(1/k)$ fraction of the edges of the cut are in the spanning tree. First conjectured in 2004 by Goddyn, this conjecture has turned out to be extremely challenging, and it remains open. It would have various interesting implications; for example, it would provide an alternative route to a constant approximation for the asymmetric traveling salesman problem.

One of the challenges of understanding thin trees is that it is a condition about every cut in a graph. Potentially easier forms of the question can be obtained by restricting our attention to some given family of cuts, and requiring the spanning tree to be “thin” only on these cuts. I will discuss some results of this type, for particular types of families, based on combinatorial approaches. In particular, I will discuss the setting where the given family is laminar; and the setting where the given family consists of near-minimum cuts in the graph.

(Includes joint work with Nathan Klein and Zi Song Yeoh.)

Thomas Rothvoss (University of Washington)

A Parameterized Linear Formulation of the Integer Hull

Abstract: Let $A \in \mathbb{Z}^{m \times n}$ be an integer matrix with components bounded by Δ in absolute value. Cook et al. (1986) have shown that there exists a universal matrix $B \in \mathbb{Z}^{m' \times n}$ with the following property: For each $b \in \mathbb{Z}^m$, there exists $t \in \mathbb{Z}^{m'}$ such that the integer hull of the polyhedron $P = \{x \in \mathbb{R}^n : Ax \leq b\}$ is described by $P_I = \{x \in \mathbb{R}^n : Bx \leq t\}$. Our main result is that t is an affine function of b as long as b is from a fixed equivalence class of the lattice $D \cdot \mathbb{Z}^m$. Here $D \in \mathbb{N}$ is a number that depends on n and Δ only. Furthermore, D as well as the matrix B can be computed in time depending on Δ and n only. An application of this result is the solution of an open problem posed by Cslovjcek et al. (SODA 2024) concerning the complexity of 2-stage-stochastic integer programming problems. The main tool of our proof is the classical theory of Chvátal-Gomory cutting planes and the elementary closure of rational polyhedra.

This is joint work with Friedrich Eisenbrand.

Lisa Sauer mann (University of Bonn)

Slicing All Edges of the Hypercube by Hyperplanes

Abstract: How many hyperplanes in \mathbb{R}^n are needed in order to slice every edge of the n -dimensional hypercube? Here, we say that a hyperplane $H \subseteq \mathbb{R}^n$ slices an edge of the hypercube if it contains exactly one interior point of the edge. The problem of determining the minimum possible size of a collection of hyperplanes in \mathbb{R}^n , such that every edge of the hypercube with vertex set $\{1, -1\}^n$ is sliced by at least one of these hyperplanes, is more than 50 years old and has been studied by many researchers. We prove a lower bound of roughly $n^{13/19}$, improving upon the best previous lower bound $\Omega(n^{2/3} \log^{-4/3} n)$ of roughly $n^{2/3}$ due to Klein. Joint work with Zixuan Xu.

Niklas Schlomberg (University of Bonn)

Improved Erdős-Pósa Inequalities for Odd Cycles in Planar Graphs

Abstract: In an undirected graph, the odd cycle packing number is the maximum number of pairwise vertex-disjoint odd cycles. The odd cycle transversal number is the minimum number of vertices that hit every odd cycle. The maximum ratio between transversal and packing number is called Erdős-Pósa ratio. We show that in planar graphs, this ratio does not exceed 4. This improves on the previously best known bound of 6 by Král', Sereni and Stacho. This is joint work with Luise Puhlmann.

András Sebő (CNRS, Laboratoire G-SCOP)

Some Paths with HIM and You over the Years

Abstract: I would like to discuss a selection of problems studied at HIM, mostly related to paths in graphs.

Some had been encountered before and were further investigated there (TSP paths), while others were raised at HIM (non-shortest paths). Some were posed there and later solved by colleagues (k-trails), while others were solved there or since then (on intersection graphs of squares). Finally, some problems were merely attempted and only partially solved at HIM and they still resist solution, leaving work in progress.

For each problem I present, I plan to highlight key ideas through one, simple illustrative example, and mention some related still open direction.

Vera Traub (ETH Zürich)

Steiner Forest: A Simplified Better-Than-2 Approximation

Abstract: In the Steiner Forest problem, we are given a graph with edge lengths, and a collection of demand pairs; the goal is to find a subgraph of least total length such that each demand pair is connected in this subgraph. For over twenty years, the best approximation ratio known for the problem was a 2-approximation due to Agrawal, Klein, and Ravi (STOC 1991), despite many attempts to surpass this bound. Finally, in a recent breakthrough, Ahmadi, Gholami, Hajiaghayi, Jabbarzade, and Mahdavi (FOCS 2025) gave a $(2 - \varepsilon)$ -approximation, where $\varepsilon \approx 10^{-11}$.

In this work, we show how to simplify and extend the work of Ahmadi et al. to obtain an improved 1.994-approximation. We combine some ideas from their work (e.g., an extended run of the moat-growing primal-dual algorithm, and identifying autarkic pairs) with other ideas—submodular maximization to find components to contract, as in the relative greedy algorithms for Steiner tree, and the use of autarkic triples. We hope that our cleaner abstraction will open the way for further improvements.

This is joint work with Anupam Gupta.

Jens Vygen (University of Bonn)

Better Approximation Guarantee for Asymmetric TSP

Abstract: We improve the approximation ratio for the Asymmetric TSP to less than 15. We also obtain improved ratios for the special case of unweighted digraphs and the generalization where we ask for a minimum-cost tour with given (distinct) endpoints. Moreover, we prove better upper bounds on the integrality ratios of the natural LP relaxations.

Karol Węgrzycki (Max Planck Institute)

A Subexponential Time Algorithm for Makespan Scheduling of Unit Jobs with Precedence Constraints

Abstract: In a classical scheduling problem, we are given a set of n jobs of unit length with precedence constraints, and the goal is to find a schedule of these jobs on m identical machines that minimizes the makespan. In standard 3-field notation, it is denoted as $Pm|prec, p_j = 1|C_{\max}$.

For $m = 2$ machines, the problem can be solved in polynomial time. Settling the complexity for any constant $m \geq 3$ is a longstanding open question in the field, asked by Lenstra and Rinnooy Kan [OR 1978] in the late 70s and prominently featured in the textbook of Garey and Johnson. Since then, the problem has been thoroughly investigated, but nontrivial solutions had been found only in special cases or relaxed settings. For example, despite the possibility of the problem being polynomially solvable in the exact setting, just the existence of an approximation-scheme is widely regarded as a major open problem (see the survey of Bansal [MAPS 2017]), but so far, only superpolynomial approximations are known.

In this paper, we make the first progress on the exact complexity of $Pm|prec, p_j = 1|C_{\max}$. We present an algorithm that runs in $2^{O(\sqrt{n} \log n)}$ time for $m = O(1)$. Before our work, only a $2^{O(n)}$ time exact algorithm was known by Held and Karp [ACM 1961].

We introduce a notion of *proper decomposition* of the schedule. With this decomposition in hand, our approach can be presented in three parts. In the *branch* part, we branch on possible subschedules generated by proper decomposition. In the *reconstruct* part, we compute the schedule based on the partial answer. Interestingly, despite the fact that the depth of the branching may be linear, by *memoization* we can show that the total number of different subproblems is bounded by $2^{O(\sqrt{n} \log n)}$.

Join work with Jesper Nederlof and Céline M. F. Swennenhuis

Sorrachai Yingchareonthawornchai (ETH Zürich)

A Little Clairvoyance Is All You Need

Abstract: In the classical problem of minimizing the total flow time of jobs on a single machine in the online setting where jobs arrive over time, it has long been known that the Shortest Remaining Processing Time (SRPT) algorithm is optimal (i.e., 1-competitive) when the job sizes are known upfront [Schrage, 1968]. But in the non-clairvoyant setting where job sizes are revealed only when the job finishes, no algorithm can be constant-competitive [Motwani, Phillips, and Torng, 1994].

We consider the ε -clairvoyant setting, where $\varepsilon \in [0, 1]$ is a fixed constant, and each job's processing time becomes known once its remaining processing time equals an ε fraction of its processing time. This captures settings where the system user uses the initial $(1 - \varepsilon)$ fraction of a job's processing time to learn its true length, which it can then reveal to the algorithm. The model was proposed by Yingchareonthawornchai and Torng (2017), and it smoothly interpolates between the clairvoyant setting (when $\varepsilon = 1$) and the non-clairvoyant setting (when $\varepsilon = 0$). In a concrete sense, we are asking: how much knowledge is required to circumvent the hardness of this problem?

We show that a little knowledge is enough, and that a constant competitive algorithm exists for every constant $\varepsilon > 0$. More precisely, for all $\varepsilon \in (0, 1)$, we present a deterministic $\lceil 1/\varepsilon \rceil$ -competitive algorithm, which is optimal for deterministic algorithms. Our algorithm to achieve this bound is remarkably simple and applies the “optimism in the face of uncertainty” principle. For each job, we

form an optimistic estimate of its length, based on the information revealed thus far and run SRPT on these optimistic estimates.

Joint work with Anupam Gupta, Haim Kaplan, Alexander Lindermayr, and Jens Schloter
