Thursday, November 10, 2022

12:30–13:00	Registration
13:00–13:25	Tim Hartman *Continuous Facility Location on Graphs*
13:25–13:50	Stefan Walzer *Simple Linear Set Sketching*
13:50–14:15	Meike Neuwohner *The Limits of Local Search for Weighted k-Set Packing*
14:15–14:45	Coffee Break
14:45–15:10	Jesse van Rhijn *Improved Smoothed Analysis of 2-Opt for the Euclidean TSP*
15:10–15:35	Niklas Jost *Approximation Algorithms for Hub Location Problems*
15:35–16:00	Kelin Luo *Package Delivery Using Drones with Restricted Movement Areas*
16:00–16:15	Coffee Break

16:15–16:40	Anna Arutyunova *The Price of Hierarchical Clustering*
16:40–17:05	Abhiruk Lahiri *Reconfiguring shortest paths in graphs*
17:05–17:30	Andreas Padalkin *Reconfigurable Circuits in the Geometric Amoebot Model*
17:30–17:45	Break
17:45–18:10	Daniel Mock *Evaluating Restricted First-Order Counting Properties on Nowhere Dense Classes and Beyond*
18:10–18:35	Timon Barlag *Logical Characterizations of algebraic circuit classes over rings*
18:35	Dinner

Friday, November 11, 2022

08:00–09:00 Breakfast

09:00–10:00 Karl Bringmann
Hardness of Approximation in P: Fine-Grained Complexity of Graph Distance Approximation

10:00–10:15 Break

10:15–10:40 Martin Nägele
Congruency-Constrained Optimization

10:40–11:05 Arindam Biswas
Sublinear-Space Approximation Algorithms for Hitting Set

11:05–11:30 Ekin Ergen
Farthest Insertion Heuristic for TSP

11:30–12:00 Coffee Break

12:00–12:25 Benedikt Kolbe
On the computational geometry of hyperbolic surfaces

12:25–12:50 Felix Tschirb
Dynamic complexity of regular languages: Big changes, small work

12:50–13:15 Kevin Mann
Minimal Roman Dominating Functions: Extensions and Enumeration

Continuous Facility Location on Graphs

Tim Hartman
RWTH Aachen

We study Facility Location problems on graphs where all edges have unit length and where the facilities may also be positioned in the interior of the edges. One goal for example, is to position as many facilities as possible subject to the condition that any two facilities have at least distance δ from each other, for some positive real δ; a problem we denote as δ-Dispersion.

Beside being a suitable model for undesirable facilities in, for example, a street networks, we also draw a close connection to the well-known Independent Set problem. While the latter is NP-hard, its direct counter-part, 2-Dispersion, turns out to be in P. We resolve the complete landscape of the complexity depending on δ - there are infinity many singular values of δ for which delta-Dispersion is in P.

We also resolve the parameterized complexity depending on the solution size for every δ. Further, we investigate the facilities complexity depending on the structure of the input graph, such as treedepth, treewidth and neighborhood diversity.

Simple Linear Set Sketching

Stefan Walzer
University of Cologne

Imagine a hash table with n buckets into which m keys are inserted. However, instead of using a collision resolution strategy we combine keys that hash to the same bucket using xor. This seems like a terrible idea since there is no way to obtain x and y from $x \oplus y$.

This strategy becomes useful if each key is stored in three independent positions. If the load m/n is less than ≈ 0.81 then the set of keys can be reconstructed with high probability.

This way of obtaining a data structure $f(X)$ from a set X can be viewed as linear set sketching. "Sketching" means that the sizes m and n of X and $f(X)$ can be selected independently, which makes $f(X)$ a lossy encoding of X if $m << n$. "Linear" means that $f(X \Delta Y) = f(X) \oplus f(Y)$ where Δ denotes symmetric difference of sets and $\oplus$ denotes bucket-wise xor of the sketches.

The Limits of Local Search for Weighted k-Set Packing

Meike Neuwohner
University of Bonn

We consider the weighted k-Set Packing problem, where, given a collection $\mathcal{S}$ of sets, each of cardinality at most k, and a positive weight function $w : \mathcal{S} \to \mathbb{R}_{>0}$, the task is to find a sub-collection of $\mathcal{S}$ consisting of pairwise disjoint sets of maximum total weight. As this problem does not permit a polynomial-time $o(\frac{k}{\log k})$-approximation unless $P = NP$, most previous approaches rely on local search. For twenty years, Berman's algorithm *SquareImp*, which yields a polynomial-time $\frac{k+1}{2} + \epsilon$-approximation for any fixed $\epsilon > 0$, has remained unchallenged. Only recently, it could be improved to $\frac{k+1}{2} - \frac{1}{63{,}700{,}992} + \epsilon$ by Neuwohner. In her paper, she showed that instances for which the analysis of SquareImp is almost tight are "close to unweighted" in a certain sense. But for the unit weight variant, the best known approximation guarantee is $\frac{k+1}{3} + \epsilon$. Using this observation as a starting point, we conduct a more in-depth analysis of close-to-tight instances of SquareImp. This finally allows us to generalize techniques used in the unweighted case to the weighted setting. In doing so, we obtain approximation guarantees of $\frac{k+\epsilon_k}{2}$, where $\lim_{k\to\infty} \epsilon_k = 0$. On the other hand, we prove that this is asymptotically best possible in that searching for local improvements of logarithmically bounded size cannot produce an approximation ratio below $\frac{k}{2}$.

Improved Smoothed Analysis of 2-Opt for the Euclidean TSP

Jesse van Rhijn
University of Twente

The 2-opt heuristic is a simple local search heuristic for the Travelling Salesperson Problem (TSP). Although it usually performs well in practice, its worst-case runtime is poor. Attempts to reconcile this difference have used smoothed analysis, in which adversarial instances are perturbed probabilistically.

We are interested in the classical model of smoothed analysis for the Euclidean TSP, in which the points in an adversarial instance are perturbed by Gaussian random variables. This model was previously used by Manthey & Veenstra, who obtained smoothed complexity bounds polynomial in n, the dimension d, and the perturbation strength σ^{-1}. However, their analysis only works for $d \geq 4$. The only previous analysis for $d \geq 3$ was performed by Englert, Röglin & Vöcking, who used a different perturbation model which can be translated to Gaussian perturbations. Their model yields bounds polynomial in n and σ^{-d}, and super-exponential in d.

As it is somewhat unsatisfactory that no direct analysis existed for Gaussian perturbations that yields polynomial bounds for all d, we perform this missing analysis. Along the way, we improve all existing smoothed complexity bounds.

Approximation Algorithms for Hub Location Problems

Niklas Jost
TU Dortmund

The task in my Ph.D. thesis is to establish approximation algorithms for Hub location problems. The basis of this task is the review of different Hub location problems by Campbell. He defined four different variants of Hub location problems. In the talk, these variants will be explained. To the best of the author's knowledge, only for two of them, approximation algorithms exist. An approximation algorithm was designed for the third variant -the p-Hub center problems- which will be sketched during the talk. Furthermore, in a project with a logistics provider, another variant of Hub location problem appeared (not covered by Campbell), which was solved by a modification of an approximation algorithm for the uncapacitated Hub location problem. The talk will be finished by mentioning aspects of this project and giving an outlook for further research.

Package Delivery Using Drones with Restricted Movement Areas

Kelin Luo
University of Bonn

For the problem of delivering a package from a source node to a destination node in a graph using a set of drones, we study the setting where the movements of each drone are restricted to a certain subgraph of the given graph. We consider the objectives of minimizing the delivery time (problem DDT) and of minimizing the total energy consumption (problem DDC). For general graphs, we show a strong inapproximability result and a matching approximation algorithm for DDT as well as NP-hardness and a 2-approximation algorithm for DDC. For the special case of a path, we show that DDT is NP-hard if the drones have different speeds. For trees, we give optimal algorithms under the assumption that all drones have the same speed or the same energy consumption rate. The results for trees extend to arbitrary graphs if the subgraph of each drone is isometric.

This is a collaborative work with Thomas Erlebach and Frits C.R. Spieksma.

The Price of Hierarchical Clustering

Anna Arutjunova
University of Bonn

Hierarchical Clustering is a popular tool for understanding the hereditary properties of a data set. Such a clustering is actually a sequence of clusterings that starts with the trivial clustering in which every data point forms its own cluster and then successively merges two existing clusters until all points are in the same cluster. A hierarchical clustering achieves an approximation factor of α if the costs of each k-clustering in the hierarchy are at most α times the costs of an optimal k-clustering. We study as cost functions the maximum (discrete) radius of any cluster (k-center problem) and the maximum diameter of any cluster (k-diameter problem). In general, the optimal clusterings do not form a hierarchy and hence an approximation factor of 1 cannot be achieved. We call the smallest approximation factor that can be achieved for any instance the price of hierarchy. For the k-diameter problem we improve the upper bound on the price of hierarchy to $3+2\sqrt{2} \approx 5.83$. Moreover we significantly improve the lower bounds for k-center and k-diameter, proving a price of hierarchy of exactly 4 and $3+2\sqrt{2}$, respectively.

Reconfiguring shortest paths in graphs

Abhiruk Lahiri
Heinrich-Heine-University Düsseldorf

Reconfiguring a system (generally represented by a graph) means gradually transforming one configuration of the system to another configuration by modifying it in a slow and steady step-by-step manner. In the Shortest Path Reconfiguration (SPR) problem, we are given two shortest paths, and the goal is to transform one shortest path to the other by changing one vertex at a time so that all the intermediate configurations are also shortest paths. SPR has several real-world applications like repaving roads in a systematic way and cargo container stowage on ships.

In this talk, we will show that SPR can be solved in polynomial time for many graph classes, including circle graphs, bridged graphs, the Boolean hypercube and bounded diameter graphs. Most of our proofs are by providing a complete characterization of the shortest paths for the graph class. We will also explore a generalization of SPR known as k-SPR, and show that k-SPR is PSPACE-complete, even for some graph classes (viz. line graphs) for which SPR is known to be solvable in polynomial time.

This is joint work with Kshitij Gajjar, Agastya Vibhuti Jha and Manish Kumar.
`https://ojs.aaai.org/index.php/AAAI/article/view/21211`

Reconfigurable Circuits in the Geometric Amoebot Model

Andreas Padalkin
University of Paderborn

The (geometric) amoebot model [Derakhshandeh et al., SPAA 2014] has been proposed as a model for programmable matter consisting of tiny, robotic elements called amoebots. Since information can only travel amoebot by amoebot, many problems have a linear lower bound. Inspired by the nervous system, we propose the reconfigurable circuit extension that allows the amoebot structure to interconnect amoebots by so-called circuits [Feldmann et al., JCB 2022]. A circuit permits the instantaneous transmission of signals between the connected amoebots. We show that such an extension allows for significantly faster solutions to a variety of problems related to programmable matter [Feldmann et al., JCB 2022; Padalkin et al., DNA 2022].

Michael Feldmann, Andreas Padalkin, Christian Scheideler, Shlomi Dolev: Coordinating Amoebots via Reconfigurable Circuits. J. Comput. Biol. 29(4): 317-343 (2022)
`https://www.liebertpub.com/doi/10.1089/cmb.2021.0363`

Andreas Padalkin, Christian Scheideler, Daniel Warner: The Structural Power of Reconfigurable Circuits in the Amoebot Model. DNA 2022: 8:1-8:22
`https://drops.dagstuhl.de/opus/volltexte/2022/16793/`

Evaluating Restricted First-Order Counting Properties on Nowhere Dense Classes and Beyond

Daniel Mock
RWTH Aachen

It is known that first-order logic with some counting extensions can be efficiently evaluated on graph classes with bounded expansion, where depth-r minors have constant density. More precisely, the formulas are $\exists x_1 \dots x_k \# y\, \phi(x_1, \dots, x_k, y) > N$, where ϕ is an FO-formula. If ϕ is quantifier-free, we can extend this result to *nowhere dense* graph classes with an almost linear FPT run time. Lifting this result further to slightly more general graph classes, namely almost nowhere dense classes, where the size of depth-r clique minors is subpolynomial, is impossible unless $\mathrm{FPT} = \mathrm{W}[1]$. On the other hand, in almost nowhere dense classes we can approximate such counting formulas with a small additive error.

In particular, it follows that partial covering problems, such as partial dominating set, have fixed parameter algorithms on nowhere dense graph classes with almost linear running time.

Joint work with Jan Dreier and Peter Rossmanith.

Logical Characterizations of algebraic circuit classes over rings

Timon Barlag
Leibniz University Hannover

We are currently working on adapting the construction of algebraic circuits over rings introduced by Blum, Shub and Smale to arbitrary integral domains. For integral domains we give a theorem in the style of Immerman's Theorem and show that families of such circuits of polynomial size and constant depth decide exactly those sets of vectors of ring elements that can be defined in first-order logic on R-structures as a generalization of $\mathbb{R}$-structures in the sense of Cucker and Meer. Additionally, we talk about a generalization of the guarded predicative logic by Durand, Haak and Vollmer and we show characterizations for the AC_R and NC_R hierarchy. Furthermore we compare some of the aforementioned complexity classes with different underlying rings and give a conjecture about which inclusions of classes could be strict.

Hardness of Approximation in P: Fine-Grained Complexity of Graph Distance Approximation

Karl Bringmann
University of Saarland and Max-Planck-Institute for Informatics

Fine-grained Complexity Theory is the design of reductions that prove running time lower bounds assuming a plausible hypothesis. In the last 10 years, this approach has been very successful in proving matching upper and lower bounds for the time complexity of many problems from various problem domains. A natural question is whether fine-grained complexity not only works for exact algorithms, but can also determine optimal approximation ratios. That is, can fine-grained complexity be extended to hardness of approximation in P? This talk surveys some developments in this direction. Our focus is a recent result for distance approximation in undirected graphs, essentially showing that the classic Thorup-Zwick distance oracles attain a near-optimal tradeoff between approximation ratio, preprocessing time, and query time.

Congruency-Constrained Optimization

Martin Nägele
University of Bonn

A long-standing open question in Integer Programming is whether integer programs with constraint matrices with bounded subdeterminants are efficiently solvable. An important special case thereof are congruency-constrained integer programs of the form $min\{c^\top x : Tx \leq b, \gamma^\top x \equiv r \pmod{m}, x \in \mathbb{Z}^n\}$ with a totally unimodular constraint matrix T. Such problems have been shown to be polynomial-time solvable for $m = 2$, which led to an efficient algorithm for integer programs with bimodular constraint matrices, i.e., full-rank matrices whose $n \times n$ subdeterminants are bounded by two in absolute value. In this talk, we delve into recent progress on congruency-constrained TU problems beyond the case of $m = 2$, and explore the challenges that come with generalizing parity constraints to arbitrary congruency constraints even in very simple combinatorial problems.

Based on joint work with Ch. Nöbel, R. Santiago, B. Sudakov, and R. Zenklusen (all ETH Zürich).

Sublinear-Space Approximation Algorithms for Hitting Set

Arindam Biswas
TU Ilmenau

Over the years, much research has been conducted on ways of dealing with the apparent intractability of NP-hard problems. The two most well known frameworks utilized to this end are approximation and parameterization. Algorithms in these frameworks generally use space which is at least polynomial in the input size. However, for some problems, one can devise algorithms that use sublinear space.

In this talk, we look at the following polynomial-time sublinear-space algorithms for approximating HITTING SET.
- An $O(d\delta \log n)$-space $n^{O(d\delta)}$-time d-approximation algorithm for HITTING SET restricted to sets of size at most d and element multiplicity δ.
- A $O((d^2+(d/\epsilon)) \log n)$-space $n^{O(d^2+(d/\epsilon))}$-time $((d/\epsilon)n^{\epsilon})$-approximation scheme for HITTING SET restricted to sets of size at most d.

Farthest Insertion Heuristic for TSP

Ekin Ergen
TU Berlin

The Traveling Salesman Problem (TSP) is one of the most well-known NP-hard problems in combinatorial optimization. Given a complete graph with non-negative edge weights c, the problem inquires a minimum-weight Hamiltonian cycle. Many approximation algorithms for this problem are studied and used in practice.

A family of heuristics commonly used in practice are the so-called *insertion heuristics* that maintain a subtour S and insert one vertex v at a time between two vertices $w, y \in V(S)$ in a way that minimizes $c(\{v, w\}) + c(\{v, y\}) - c(\{w, y\})$. Such a heuristic has been proven to have an approximation ratio of $O(\log(n))$. Different strategies of choosing v give rise to different algorithms such as Nearest Insertion, Cheapest Insertion and Farthest Insertion. Nearest Insertion and Cheapest Insertion have an approximation ratio of 2, whereas an improved worst-case approximation ratio of Farthest Insertion is not known, even though this variant is observed to provide better solutions than the constant-ratio variants in practice.

In this talk, we present a simplified proof for the $O(log(n))$ approximation bound for Farthest Insertion. We also mention a slight improvement on the lower bounds of this heuristic and discuss the possibility to find instances admitting non-constant bounds. Finally, we sketch an analysis for the special case where edge weights are assumed to be 1 or 2.

On the computational geometry of hyperbolic surfaces

Benedikt Maximilian Kolbe
University of Bonn

Hyperbolic surfaces play a ubiquitous role in mathematics, with applications and relationships with areas as divers as geometric group theory, complex networks, mathematical physics, crystallography, biology, as well as image manipulation and recognition. Motivated by the varied applications, there has been a push to understand hyperbolic structures on surfaces from a computational geometry point of view. With the recent advent of robust software for different types of Euclidean Delaunay triangulations, Delaunay triangulations on spaces with a hyperbolic structure have garnered increasing interest. In this talk, we present recent findings on the structure of Delaunay triangulations of hyperbolic surfaces that have led to the only software for the robust computation of such triangulations in the case of genus 2 surfaces.

Dynamic complexity of regular languages: Big changes, small work

Felix Tschirbs
Ruhr-University Bochum

Whether a changing string is member of a certain regular language can be maintained in the DynFO framework of Patnaik and Immerman: after changing the symbol at one position of the string, a first-order update formula can express – using additionally stored information – whether the resulting string is in the regular language.

We extend this and further known results by considering changes of many positions at once. We also investigate to which degree the obtained update formulas imply work-efficient parallel dynamic algorithms.

Minimal Roman Dominating Functions: Extensions and Enumeration

Kevin Mann
University of Trier

Roman domination is one of the many variants of domination that keeps most of the complexity features of the classical domination problem. We prove that Roman domination behaves differently in two aspects: enumeration and extension. We develop non-trivial enumeration algorithms for minimal Roman domination functions with polynomial delay and polynomial space. Recall that the existence of a similar enumeration result for minimal dominating sets is open for decades. Our result is based on a polynomial-time algorithm for Extension Roman Domination: Given a graph $G = (V, E)$ and a function $f : V \to \{0, 1, 2\}$, is there a minimal Roman domination function $\tilde{f}$ with $f \leq \tilde{f}$? Here, $\leq$ lifts $0 < 1 < 2$ pointwise; minimality is understood in this order. Our enumeration algorithm is also analyzed from an input-sensitive viewpoint, leading to a run-time estimate of $O(1.9332^n)$ for graphs of order n; this is complemented by an example showing a lower bound of $\Omega(1..7441^n)$.