77. Theorietag
Workshop über Algorithmen und Komplexität

Book of Abstracts

Philipps Universität Marburg

March 28 2019

Program

9:30-9:45	Welcome reception
9:45-10:10	Markus L. Schmid: Graph and String Parameters: Connections Between Pathwidth, Cutwidth and the Locality Number
10:10-10:35	Andrej Sajenko: Space-Efficient $O(1)$-Approximation Algorithm for Treewidth
10:40-11:15	Klaus Heeger: Structural Parameterizations of Stable Roommates with Ties

---

11:15-11:30	Coffee break

11:30-11:55	Sarah Morell: Diversity Maximization in Doubling Metrics
11:55-12:20	Malte Renken: Comparing Temporal Graphs with Time Warping
12:20-12:45	Hendrik Molter: Sliding Window Temporal Graph Coloring

---

12:45-14:00	Lunch

14:00-15:00	Peter Rossmanith: What One Has to Know When Attacking P vs. NP
15:00-15:15	Coffee break

15:15-15:40	Petra Wolf: On the Decidability and Complexity of Finding a Positive Instance among Infinitely Many
15:40-16:05	Jens Kosiol: Algorithmic Problems in Algebraic Graph Transformation
16:05-16:30	Rob van Stee: The Price of Clustering in Bin-Packing with Applications to Bin-Packing with Delays

---

17:30-19:00	City tour
19:15	Dinner

Graph and String Parameters: Connections Between Pathwidth, Cutwidth and the Locality Number

Markus L. Schmid, Universität Trier

The locality number is a recently introduced structural parameter for strings (with applications in pattern matching with variables). A word is k-local if there exists an order of its symbols such that, if we mark the symbols in the respective order (which is called a marking sequence), at each stage there are at most k contiguous blocks of marked symbols in the word. The locality number of a word is the smallest k for which that word is k-local.

We investigate the problem of computing the locality number (and respective marking sequences) by relating it to the two important graph-parameters cutwidth and pathwidth. It turns out that the locality number is very closely connected to the cutwidth (by natural transformations of words into graphs and vice versa), which allows us to show that computing the locality number is NP-hard but fixed parameter tractable (when the locality number or the alphabet size is treated as a parameter). Moreover, by relating the locality number to the pathwidth of graphs, we can show that it can be approximated with ratio $O(\sqrt{\log opt} \log n)$.

As a surprising by-product of independent interest, we also relate cutwidth via the locality number to pathwidth, which improves the currently best known approximation algorithm for cutwidth.

In addition to these main results, we also consider the possibility of greedy-based approximation algorithms for the locality number.

This talk is based on joint work with Katrin Casel (Hasso Plattner Institute, University of Potsdam, Germany), Joel D. Day (Loughborough University, UK), Pamela Fleischmann (Kiel University, Germany), Tomasz Kociumaka (University of Warsaw, Poland, and Bar-Ilan University, Israel), and Florin Manea (Kiel University, Germany).

Space-Efficient $O(1)$-Approximation Algorithm for Treewidth

Andrej Sajenko, THM Gießen

Given an input n-vertex graph G and an integer $k > 0$, we give a streaming algorithm that uses $O(kn)$ bits and runs in $2^{O(k)}n(\log^* n)\log\log n$ total time. The algorithm either outputs that the treewidth of G is larger than k or returns a tree decomposition of G of width at most $O(k)$.

We also show that our algorithm can be used to solve several monadic second-order problems using only $O(kn + c^k \log^2 n)$ bits for some constant $c > 1$.

Structural Parameterizations of Stable Roommates with Ties

Klaus Heeger, TU Berlin

Given a set of agents with strict preferences over all other agents, the polynomial-time solvable Stable Roommates problem asks to find a stable matching of the agents. A matching is stable if there is no pair of agents who both prefer each other over their partners in the matching.

We consider the well-known NP-complete generalization Stable Roommates with Ties and Incomplete Preferences (SRTI). In SRTI, agents can be unacceptable (that is, some pairs of agents cannot be matched) and agents are allowed to tie two or more persons in their preference list.

We study the parameterized complexity of this problem, focussing on graph-theoretic parameters. Solving an open question by Adil et al. (TCS 2018), we show that SRTI is W[1]-hard when parameterized by tree-width. In fact we even show that SRTI is W[1]-hard for feedback vertex set size as well as for tree-depth. On the positive side SRTI becomes fixed-parameter tractable when parameterized by feedback edge set or tree-cut-width. Furthermore, we show that finding a maximum-size stable matching is W[1]-hard for tree-cut-width, but there exists an algorithm that computes a factor-two approximation in FPT time for this parameter.

Diversity Maximization in Doubling Metrics

Sarah Morell, TU Berlin

Diversity maximization is an important geometric optimization problem with many applications in recommender systems, machine learning or search engines among others. A typical diversification problem is as follows: Given a finite metric space (X, d) and a positive integer k, find a subset of k elements of X that has maximum diversity. There are many functions that measure diversity. One of the most popular measures, called remote-clique, is the sum of the pairwise distances of the chosen elements. In this talk, we present novel results on three widely used diversity measures: Remote-clique, remote-star and remote-bipartition.

Our main result are polynomial time approximation schemes for these three diversification problems under the assumption that the metric space is doubling. This setting has been discussed in the recent literature. The existence of such a PTAS however was left open. Our results also hold in the setting where the distances are raised to a fixed power $q \geq 1$, giving rise to more variants of diversity functions, similar in spirit to the variations of clustering problems depending on the power applied to the distances.

This talk is based on joint work with Alfonso Cevallos (ETH Zürich, Switzerland) and Friedrich Eisenbrand (EPFL Lausanne, Switzerland).

Comparing Temporal Graphs with Time Warping

Malte Renken, TU Berlin

The connections within many real-world networks change over time, leading to the study of so-called temporal graphs. Recognizing patterns in these requires a similarity measure to compare different temporal graphs. To this end, we propose an approach using dynamic time warping (an established concept in the context of time series). The resulting measure is called the temporal graph warping distance. Finally, we talk about the hardness of computing this distance, give ways to get around it, and show some examples involving real-world data.

This talk is based on joint work with Vincent Froese (TU Berlin, Germany), Brijnesh Jain (TU Berlin, Germany), and Rolf Niedermeier (TU Berlin, Germany).

Sliding Window Temporal Graph Coloring

Hendrik Molter, TU Berlin

Graph coloring is one of the most famous computational problems with applications in a wide range of areas such as planning and scheduling, resource allocation, and pattern matching. So far coloring problems are mostly studied on static graphs, which often stand in stark contrast to practice where data is inherently dynamic and subject to discrete changes over time. A temporal graph is a graph whose edges are assigned a set of integer time labels, indicating at which discrete time steps the edge is active. In this paper we present a natural temporal extension of the classical graph coloring problem. Given a temporal graph and a natural number Δ, we ask for a coloring sequence foreach vertex such that (i) in every sliding time window of Δ consecutive time steps, in which an edge is active, this edge is properly colored (i.e. its endpoints are assigned two different colors) at least once during that time window, and (ii) the total number of different colors is minimized. This sliding window temporal coloring problem abstractly captures many realistic graph coloring scenarios in which the underlying network changes over time, such as dynamically assigning communication channels to moving agents. We present a thorough investigation of the computational complexity of this temporal coloring problem. More specifically, we prove strong computational hardness results, complemented by efficient exact and approximation algorithms. Some of our algorithms are linear-time fixed-parameter tractable with respect to appropriate parameters, while others are asymptotically almost optimal under the Exponential Time Hypothesis (ETH).

This talk is based on joint work with George B. Mertzios (Durham University, UK) and Viktor Zamaraev (Durham University, UK).

On the Decidability and Complexity of Finding a Positive Instance among Infinitely Many

Petra Wolf, Universität Trier

For a fixed problem P (e.g. Independent Set, Vertex Cover, TQBF, etc.), IntReg(P) is the problem to decide whether a given regular set R of P-instances (represented as a finite automaton) contains at least one positive instance.

For different choices of P, we investigate the decidability and complexity of IntReg(P) (also in comparison to the decidability and complexity of P) and demonstrate that IntReg(P) can be undecidable even though P is easy, and that IntReg(P) can be decidable even though P is PSPACE-hard. We also pose the question whether the concept of "IntReg-complexity" may serve as the basis for a new complexity framework that allows to gain insights into computational problems that are not revealed by common classifications (e.g. NP-hardness, fixed-parameter tractability, etc.).

Algorithmic Problems in Algebraic Graph Transformation

Jens Kosiol, Philipps-Universität Marburg

This talk shortly reviews basic concepts of algebraic graph transformation and points out some open algorithmic problems that are evoked by current research.

The Price of Clustering in Bin-Packing with Applications to Bin-Packing with Delays

Rob van Stee, Universität Siegen

One of the most significant algorithmic challenges in the "big data era" is handling instances that are too large to be processed by a single machine. The common practice in this regard is to partition the massive problem instance into smaller ones and process each one of them separately. In some cases, the solutions for the smaller instances are later on assembled into a solution for the whole instance, but in many cases this last stage cannot be pursued (e.g., because it is too costly, because of locality issues, or due to privacy considerations). Motivated by this phenomenon, we consider the following natural combinatorial question: Given a *bin-packing* instance (namely, a set of items with sizes in $(0, 1]$ that should be packed into unit capacity bins) I and a partition $\{I_i\}_i$ of I into *clusters*, how large is the ratio $\sum_i \mathtt{Opt}(I_i)/\mathtt{Opt}(I)$, where $\mathtt{Opt}(J)$ denotes the optimal number of bins into which the items in J can be packed?

In this paper, we investigate the supremum of this ratio over all instances I and partitions $\{I_i\}_i$, referred to as the bin-packing *price of clustering* (PoC). It is trivial to observe that if each cluster contains only one tiny item (and hence, $\mathtt{Opt}(I_i) = 1$), then the PoC is unbounded. On the other hand, a relatively straightforward argument shows that under the constraint that $\mathtt{Opt}(I_i) \geq 2$, the PoC is 2. Our main challenge was to determine whether the PoC drops below 2 when $\mathtt{Opt}(I_i) > 2$. In addition, one may hope that $\lim_{k\to\infty} \mathrm{PoC}(k) = 1$, where $\mathrm{PoC}(k)$ denotes the PoC under the restriction to clusters I_i with $\mathtt{Opt}(I_i) \geq k$. We resolve the former question affirmatively and the latter one negatively: Our main results are that $\mathrm{PoC}(k) \leq 1.951$ for any $k \geq 3$ and $\lim_{k\to\infty} \mathrm{PoC}(k) = 1.691\ldots$. Moreover, the former bound cannot be significantly improved as $\mathrm{PoC}(3) > 1.933$. In addition to the immediate contribution of this combinatorial result to "big data" kind of applications, it turns out that it is useful also for an interesting online problem called *bin-packing with delays*.

This talk is based on joint work with Yossi Azar (Tel-Aviv University, Israel), Yuval Emek (Technion, Israel), and Danny Vainstein (Tel-Aviv University, Israel).