GOETHE UNIVERSITÄT FRANKFURT AM MAIN

GESELLSCHAFT FÜR INFORMATIK

82. Theorietag

08.06 - 09.06.2022

Frankfurt am Main

GI-Fachgruppen Komplexität und Algorithmen

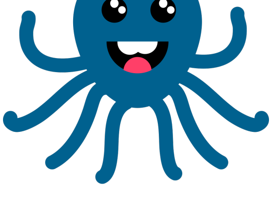

Pascal Schweitzer (TU Darmstadt): The Theory of Practical Graph Isomorphism Solving

The graph isomorphism problem is one of the few remaining prominent problems in NP for which we do not know whether it is in P and also do not know whether it is NP-complete. It is of theoretical interest due to cross connections to other theoretical questions arising in quantum computing, computational group theory, and complexity theory itself.

However, it is also of practical interest. Fast algorithmic solutions for graph isomorphism have major applications in practice. For example, symmetries of combinatorial objects can be found with graph isomorphism solvers and once found they can be used to prune search trees. In this talk I will discuss the theory behind current practical graph isomorphism solvers.

The talk is based on various results obtained in joint works with several authors (Markus Anders, Jendrik Brachter, Daniel Neuen, and Florian Wetzels).

Marcus Wilhelm (Karlsruhe Institute of Technology): Deterministic Performance Guarantees for Bidirectional BFS on Real-World Networks

A common technique to speed up shortest path queries in graphs is to use a bidirectional search, i.e., performing a forward search from the start and a backwards search from the destination until a common vertex on a shortest path is found. In practice, this has a tremendous impact on the performance on some real-world networks, while it only seems to save a constant factor on other types of networks. Even though finding shortest paths is a ubiquitous problem, there are only few studies attempting to understand the apparently asymptotic speedups on some networks, using average case analysis on certain models for real-world networks.

In this paper we give a new perspective on this, by analyzing deterministic properties that permit theoretical analysis and that can easily be checked on any particular instance. We prove that these parameters imply sublinear running time for the bidirectional BFS in several regimes, some of which are tight. Moreover, we perform experiments on a large set of real-world networks showing that our parameters capture the concept of practical running time well.

Nils Morawietz (Philipps-Universität Marburg): Parameterized Local Search for Max c-Cut

In the NP-hard Max c-Cut problem, one is given an undirected graph G and wants to color the vertices of G with c colors such that the number of edges with distinctly colored endpoints is maximal. The case with $c = 2$ is the famous Max Cut problem. To deal with the NP-hardness of this problem, we study parameterized local search algorithms. More precisely, we study LS-Max c-Cut where we are additionally given a vertex coloring f and an integer k and the task is to find a better coloring f' that differs from f in at most k entries, if such a coloring exists; otherwise, f is k-optimal. We show that LS-Max c-Cut presumably cannot be solved in $g(k) \cdot n^{O(1)}$ time even on bipartite graphs, for all $c \geq 2$. We then show an algorithm for LS-Max c-Cut with running time $O((3e\Delta)^k \cdot c \cdot k^2\Delta \cdot n)$, where Δ is the maximum degree of the input graph, and evaluate its practical performance for moderate values of k in combination with a state-of-the-art heuristic algorithm for Max c-Cut.

This is joined work with Jaroslav Garvardt, Niels Grüttemeier, and Christian Komusiewicz (all from Philipps-Universität Marburg).

Johannes Meintrup (TU Mittelhessen): Space-Efficient Graph Coarsening with Applications to Succinct Planar Encodings

We present a novel space-efficient graph coarsening technique for n-vertex separable graphs G, in particular for planar graphs, called *cloud partition*, which partitions the vertices $V(G)$ into disjoint sets C of size $O(\log n)$ such that each C induces a connected subgraph of G. Using this partition $\mathcal{P}$ we construct a so-called *structure-maintaining minor* F of G via specific contractions within the disjoint sets such that F has $O(n/\log n)$ vertices. The combination of $(F, \mathcal{P})$ is referred to as a *cloud decomposition*.

We call a graph $G = (V, E)$ separable if it admits to an $O(n^c)$-separator theorem for some constant $c < 1$ meaning there exists a separator $S \subset V$ that partitions V into $\{A, S, B\}$ such that no vertices of A and B are adjacent in G and neither A nor B contain more than $c'n$ vertices for a fixed constant $c' < 1$. Due to the last property such separators are called balanced. This famously includes planar graphs, which admit an $O(\sqrt{n})$-separator theorem. For planar graphs we show that a cloud decomposition can be constructed in $O(n)$ time and using $O(n)$ bits. Given a cloud decomposition $(F, \mathcal{P})$ constructed for a planar graph G we are able to find a balanced separator of G in $O(n/\log n)$ time. Contrary to related publications, we do not make use of an embedding of the input graph. This allows us to construct the succinct encoding scheme for planar graphs due to Blelloch and Farzan (CPM 2010) in $O(n)$ time and $O(n)$ bits improving both runtime and space by a factor of $\Theta(\log n)$.

As an additional application of our cloud decomposition we show that a tree decomposition for planar graphs of width $O(n^{1/2+\epsilon})$ for any $\epsilon > 0$ can be constructed in $O(n)$ bits and a time linear in the size of the tree decomposition. A similar result by Izumi and Otachi (ICALP 2020) constructs a tree decomposition of width $O(k\sqrt{n}\log n)$ for graphs of treewidth $k \leq \sqrt{n}$ in sublinear space and polynomial time.

Finally, we generalize our cloud decomposition from planar graphs to arbitrary separable graphs.

Tomohiro Koana (TU Berlin): The Complexity of Finding Fair Many-to-One Matchings

We analyze the (parameterized) computational complexity of "fair" variants of bipartite many-to-one matching, where each vertex from the "left" side is matched to exactly one vertex and each vertex from the "right" side may be matched to multiple vertices. We want to find a "fair" matching, in which each vertex from the right side is matched to a "fair" set of vertices. Assuming that each vertex from the left side has one color modeling its attribute, we study two fairness criteria. In one of them, we deem a vertex set fair if for any two colors, the difference between the numbers of their occurrences does not exceed a given threshold. Fairness is relevant when finding many-to-one matchings between students and colleges, voters and constituencies, and applicants and firms. Here colors may model sociodemographic attributes, party memberships, and qualifications, respectively. We show that finding a fair many-to-one matching is NP-hard even for three colors and maximum degree five. Our main contribution is the design of fixed-parameter tractable algorithms with respect to the number of vertices on the right side. Our algorithms make use of a variety of techniques including color coding. At the core lie integer linear programs encoding Hall like conditions. To establish the correctness of our integer programs, we prove a new separation result, inspired by Frank's separation theorem [Frank, Discrete Math. 1982], which may also be of independent interest. We further obtain complete complexity dichotomies regarding the number of colors and the maximum degree of each side.

This is joint work with Niclas Boehmer.

Augusto Modanese (Karlsruhe Institute of Technology): Sublinear-Time Probabilistic Cellular Automata

As every model of computation, cellular automata have been widely studied as language acceptors. Nevertheless, these efforts seem to have been almost exclusively devoted to the linear- or real-time case—to the detriment of the *sublinear-time* one. This is particularly unfortunate since, as it was recently shown, the study of sublinear-time variants of cellular automata might help better direct efforts in resolving outstanding problems in computational complexity theory.

In this talk, I propose and present a *probabilistic* model of sublinear-time cellular automata. The main motivation is to analyze to what extent (if at all) the addition of randomness to the model is able to make up for its inherent limitations. The talk is based on ongoing work as well as a recent preprint, see `https://arxiv.org/abs/2203.14614`.

The model. The most widely studied acceptance condition for sublinear-time cellular automata is that of all cells simultaneously accepting, which yields the model of *ACA* (where the first "A" in the acronym indicates that *all* cells must accept). I present a probabilistic version of the ACA model called *probabilistic ACA* (PACA). The underlying automaton is similar to a model previously proposed by Arrighi, Schabanel, and

Theyssier: At every step, each cell tosses a fair coin $c \in \{0,1\}$ and then changes its state based on the outcome of c.

One may consider both one- and two-sided error versions of the model (as natural counterparts to their Turing machine analogues). In the talk, we will see full characterization results in the constant-time setting as well as complexity-theoretical consequences of derandomizing PACA in the general sublinear-time setting.

A connection to sliding-window algorithms. A *sliding-window algorithm* is an algorithm whose current operation depends solely on the last few symbols read. A recent paper by Pacut et al. has pointed out a connection between distributed and sliding-window algorithms. Since sublinear-time cellular automata are essentially a "bare-bones" model of distributed computation, this is a promising connection for obtaining derandomization results for PACA. In the talk, I will also address the general idea for this approach.

Markus Schmid (Humboldt-Universität zu Berlin): Subsequences With Gap Constraints: Complexity Bounds for Matching and Analysis Problems

We consider subsequences with gap constraints, i.e., length-k subsequences p that can be embedded into a string w such that the induced gaps (i.e., the factors of w between the positions to which p is mapped to) satisfy given gap constraints $gc = (C_1, C_2, \ldots, C_{k-1})$; we call p a gc-subsequence of w. In the case where the gap constraints gc are defined by lower and upper length bounds $C_i = (L_i^-, L_i^+) \in \mathbb{N}^2$ and/or regular languages $C_i \in$ REG, we prove tight (conditional on the orthogonal vectors (OV) hypothesis) complexity bounds for checking whether a given p is a gc-subsequence of a string w. We also consider the whole set of all gc-subsequences of a string, and investigate the complexity of the universality, equivalence and containment problems for these sets of gc-subsequences.

Armin Weiß (Universität Stuttgart): Equation satisfiability for finite solvable groups

Over twenty years ago, Goldmann and Russell initiated the study of the complexity of the equation satisfiability problem (PolSat) and the NUDFA program satisfiability problem (ProgSat) in finite groups. They showed that these problems are decidable in polynomial time for nilpotent groups while they are NP-complete for non-solvable groups. However, for a long time the case of solvable but non-nilpotent groups remained wide open – in a long sequence of papers only the case of p-by-abelian groups could be shown to be in polynomial time.

Only in 2020 Idziak, Kawałek, Krzaczkowski and myself succeeded to show that in groups of Fitting length at least three, PolSat is not in P under the condition that the exponential time hypothesis (ETH) holds. In my talk I will review this result and also provide some details on very recent work considering PolSat for groups of Fitting length two (ie groups which have a nilpotent normal subgroup with a nilpotent quotient). Moreover, I will explain the related problems ProgSat and ListPolSat for which, under ETH and the so-called constant degree hypothesis, we can get a complete classification in which cases they are in P.

The talk is based on `https://drops.dagstuhl.de/opus/volltexte/2020/12509/`, `https://arxiv.org/abs/2010.11788` as well as research to appear soon on arXiv and in the ICALP 2022 Proceedings.

Jana Holznigenkemper (Philipps-Universität Marburg): On Computing Exact Means of Time Series Using the Move-Split-Merge Metric

Computing an accurate mean of a set of time series is a critical task in applications like classification and clustering of time series. While there are many distance functions for time series, the most popular distance function used for the computation of time series means is the non-metric dynamic time warping (DTW) distance. A recent algorithm for the exact computation of a DTW-Mean has a running time of $\mathcal{O}(n^{2k+1}2^k k)$, where k denotes the number of time series and n their maximum length. We study the mean problem for the move-split-merge (MSM) metric and give a dynamic program computing the exact MSM-Mean problem of time series. The running time of our algorithm is $\mathcal{O}(n^{k+3}2^k k^3)$, and thus better than the previous DTW-based algorithm. Moreover, we give new essential characteristics of the MSM metric to show that the MSM-Mean only consists of data points that are present in at least one time series of the input set.

Dominik Scheder (Shanghai Jiaotong University): PPSZ is better than you think

PPSZ, for long time the fastest known algorithm for k-SAT, works by going through the variables of the input formula in random order; each variable is then set randomly to 0 or 1, unless the correct value can be inferred by an efficiently implementable rule (like small-width resolution; or being implied by a small set of clauses). We show that PPSZ performs exponentially better than previously known, for all $k \geq 3$. In their 2019 STOC paper, Hansen, Kaplan, Zamir, and Zwick gave an improved version of PPSZ for all k. Whereas they present an *improved algorithm*, our improvements are achieved without changing the original PPSZ algorithm, i.e., solely through an *improved analysis*. The core idea is to pretend that PPSZ does not process the variables in uniformly random order, but according to a carefully designed distribution. We write “pretend” since this can be done without any actual change to the algorithm.

Yasir Mahmood (Leibniz Universität Hannover): Counter Claim Augmented Argumentation Frameworks

Dung's framework is a popular analytical tool for modeling, evaluating, and comparing arguments in abstract argumentation. Despite its popularity, various works point out that abstract argumentation frameworks lack certain features which are common in almost every form of argumentation in practice.

In this work, we augment Dung's framework with counter claims (CC), where every argument is connected with a specific ASP program. Intuitively, a counter claim is a reason not to choose its argument in an extension unless other arguments are selected such that the counter claims are invalidated altogether. This allows for a compact and elegant representation of AFs and merges benefits of answer set programming into AFs. We analyze the computational complexity of problems for counter claim frameworks, also when considering treewidth.

Sabrina Gaube (Leibniz Universität Hannover): Logical Charactarizations of algebraic circuit classes over rings

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 ringelements 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. Furthermore we show some equalities of some $\mathrm{AC}^0_R = \mathrm{AC}^0_{R'}$ classes with different underlying rings R and R' and give a conjecture stating which inclusions of classes could be strict inclusions. Lastly, we talk about a generalization of the guarded predicative logic by Durand, Haak and Vollmer and we show $\mathrm{NC}^1_R = \mathrm{FO}_R(\mathrm{GPR}_{\mathrm{bound}})$.

Tim Koglin (Goethe-Universität Frankfurt): Public Signals in Network Congestion Games

It is a well-known fact that selfish behavior degrades the performance of traffic networks. We consider a largely untapped potential of network improvement rooted in the inherent uncertainty of travel times. Travel times are subject to stochastic uncertainty resulting from various parameters such as weather condition, occurrences of road works, or traffic accidents. Large mobility services have an informational advantage over single network users as they are able to learn traffic conditions from data. A benevolent mobility service may use this informational advantage in order to steer the traffic equilibrium into a favorable direction. The resulting optimization problem is a task commonly referred to as *signaling* or *Bayesian persuasion*.

Previous work has shown that the underlying signaling problem can be NP-hard to approximate within any non-trivial bounds [Bhaskar et al., EC'16], even for affine cost functions with stochastic offsets. In contrast, we show that in this case, the signaling problem is easy for many networks. We tightly characterize the class of single-commodity networks, in which *full information revelation* is always an optimal signaling strategy.

Moreover, we construct a reduction from optimal signaling to computing an optimal collection of support vectors for the Wardrop equilibrium. For two states, this insight can be used to compute an optimal signaling scheme. The algorithm runs in polynomial time whenever the number of different supports resulting from any signal distribution is bounded to a polynomial in the input size. Using a cell decomposition technique, we extend the approach to a polynomial-time algorithm for multi-commodity parallel link networks with a constant number of commodities, even when we have a constant number of different states of nature.

This is joint work with Svenja M. Griesbach, Martin Hoefer, and Max Klimm, see `https://arxiv.org/abs/2205.09823`.