Parameterized Approximation Algorithms Workshop

PAAW

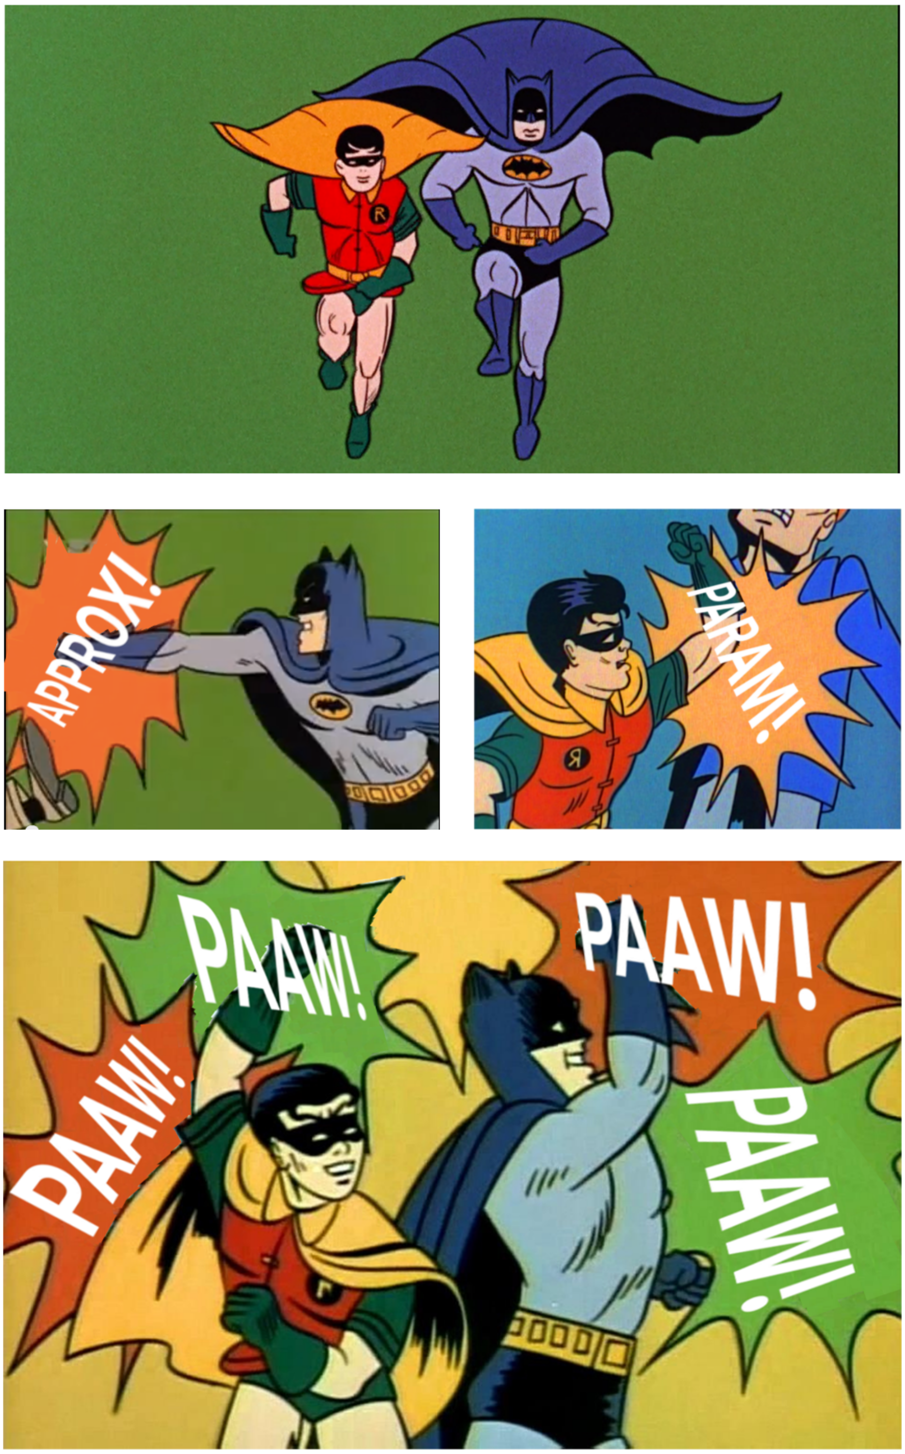

ICALP satellite workshop
Prague, Czechia
July 9th, 2018

Room: S3 (3rd floor)

Session chairs:

1. 09:00 – 10:30: Andreas Emil Feldmann
2. 11:00 – 12:30: Michael Lampis
3. 14:00 – 15:30: Michael Dinitz
4. 16:00 – 17:30: Parinya Chalermsook

Time	**Speaker**	**Title**
09:00 – 09:30	Parinya Chalermsook	From Gap-ETH to FPT-Inapproximability: Clique, Dominating Set, and More
09:30 – 10:00	Bingkai Lin	FPT-Inapproximability of Minimum Codeword Problem over Large Fields
10:00 – 10:30	Bundit Laekhanukit	On the Inapproximability of Parameterized Dominating Set
10:30 – 11:00	Coffee Break	
11:00 – 11:30	Ariel Kulik	Parameterized Approximation via Fidelity Preserving Transformations
11:30 – 12:00	Eduard Eiben	Lossy Kernels for Connected Dominating Set on Sparse Graphs
12:00 – 12:30	Krzysztof Sornat	Approximation and Parameterized Complexity of Minimax Approval Voting
12:30 – 14:00	Lunch Break	
14:00 – 14:30	Euiwoong Lee	Losing Treewidth by Separating Subsets
14:30 – 15:00	Jason Li	An FPT Algorithm Beating 2-Approximation for k-Cut
15:00 – 15:30	Michael Lampis	Parameterized (Approximate) Defective Coloring
15:30 – 16:00	Coffee Break	
16:00 – 16:30	Michael Dinitz	Characterizing Demand Graphs for (Fixed-Parameter) Shallow-Light Steiner Network
16:30 – 17:00	Daniel Vaz	Beyond Metric Embedding: Approximating Group Steiner Trees on Bounded Treewidth Graphs
17:00 – 17:30	Tomáš Masařík	Parameterized Approximation Schemes for Steiner Trees with Small Number of Steiner Vertices

From Gap-ETH to FPT-Inapproximability: Clique, Dominating Set, and More

Speaker: Parinya Chalermsook
Time: 09:00 – 9:30

We consider questions that arise from the intersection between the areas of polynomial-time approximation algorithms, subexponential-time algorithms, and fixed-parameter tractable algorithms. The questions, which have been asked several times (e.g., [Marx08, FGMS12, DF13]), are whether there is a non-trivial FPT-approximation algorithm for the Maximum Clique (Clique) and Minimum Dominating Set (DomSet) problems parameterized by the size of the optimal solution. In particular, letting OPT be the optimum and N be the size of the input, is there an algorithm that runs in $t(OPT)\text{poly}(N)$ time and outputs a solution of size $f(OPT)$, for any functions t and f that are independent of N (for Clique, we want $f(OPT) = \omega(1)$)?

In this paper, we show that both Clique and DomSet admit no non-trivial FPT-approximation algorithm, i.e., there is no $o(OPT)$-FPT-approximation algorithm for Clique and no $f(OPT)$-FPT-approximation algorithm for DomSet, for any function f (e.g., this holds even if f is the Ackermann function). In fact, our results imply something even stronger: The best way to solve Clique and DomSet, even approximately, is to essentially enumerate all possibilities. Our results hold under the Gap Exponential Time Hypothesis (Gap-ETH) [Dinur16, MR16], which states that no $2^{o(n)}$-time algorithm can distinguish between a satisfiable 3SAT formula and one which is not even (1ε)-satisfiable for some constant $\varepsilon > 0$.

Besides Clique and DomSet, we also rule out non-trivial FPT-approximation for Maximum Balanced Biclique, Maximum Subgraphs with Hereditary Properties, and Maximum Induced Matching in bipartite graphs. Additionally, we rule out $k^{o(1)}$-FPT-approximation algorithm for Densest k-Subgraph although this ratio does not yet match the trivial $O(k)$-approximation algorithm.

FPT-Inapproximability of Minimum Codeword Problem over Large Fields

Speaker: Bingkai Lin
Time: 09:30 – 10:00

Given a set C of n vectors in F_q^m and an integer k, the parameterized minimum codeword problem asks for a non-zero vector in the linear span of C with Hamming weight at most k. When the finite field F_q is binary, this problem is equivalent to the Even Set problem, whose parameterized complexity is still open. In this talk, we show that if the field size is large, i,e., $q = \Omega(n)$, then it is W[1]-hard to approximate minimum codeword within a $\log k$ ratio.

On the Inapproximability of Parameterized Dominating Set

Speaker: Bundit Laekhanukit
Time: 10:00 – 10:30

In the approximate k-Dominating Set problem we are given an integer k and a graph G on n vertices as input, and the goal is to find a dominating set of size at most $F(k)$ times k, whenever the graph G has a dominating set of size k. When such an algorithm runs in time $T(k)\text{poly}(n)$ (i.e., FPT-time) for some computable function T, it is said to be an $F(k)$-FPT approximation algorithm for k-Dominating Set. In this talk, we will show that assuming W[1] is not equal to FPT, for every computable functions F and T, there is no $F(k)$-FPT approximation algorithm for k-Dominating Set. Additionally, we will see a proof sketch of how to obtain tighter running time lower bounds under stronger time hypotheses.

Parameterized Approximation via Fidelity Preserving Transformations

Speaker: Ariel Kulik
Time: 11:00 – 11:30

We motivate and describe a parameterized approximation paradigm which studies the interaction between approximation ratio and running time for any parametrization of a given optimization problem. As a key tool, we introduce the concept of an α-shrinking transformation, for $\alpha \geq 1$. Applying such transformation to a parameterized problem instance decreases the parameter value, while preserving the approximation ratio of α (or α-fidelity).

For example, it is well-known that Vertex Cover cannot be approximated within any constant factor better than 2 (under unique games conjecture). Our parameterized α-approximation algorithm for k-Vertex Cover, parameterized by the solution size, has a running time of $1.273^{(2-\alpha)k}$, where the running time of the best FPT algorithm is 1.273^k. Our algorithms define a continuous tradeoff between running times and approximation ratios, allowing practitioners to appropriately allocate computational resources.

Moving beyond the approximation ratio, our α-shrinking transformations can be used to obtain approximative kernels which are smaller than the best known for a given problem. The smaller "α-fidelity" kernels allow us to obtain an exact solution for the reduced instance more efficiently, while obtaining an approximate solution for the original instance.

We show that such fidelity preserving transformations exist for several fundamental problems, including Vertex Cover, d-Hitting Set, Connected Vertex Cover and Steiner Tree.

Lossy Kernels for Connected Dominating Set on Sparse Graphs

Speaker: Eduard Eiben
Time: 11:30 – 12:00

For $\alpha > 1$, an α-approximate (bi-)kernel is a polynomial-time algorithm that takes as input an instance (I, k) of a problem Q and outputs an instance (I', k') (of a problem Q') of size bounded by a function of k such that, for every $c \geq 1$, a c-approximate solution for the new instance can be turned into a $(c \cdot \alpha)$-approximate solution of the original instance in polynomial time. This framework of lossy kernelization was recently introduced by Lokshtanov et al. We study Connected Dominating Set (and its distance-r variant) parameterized by solution size on sparse graph classes like biclique-free graphs, classes of bounded expansion, and nowhere dense classes. We prove that for every $\alpha > 1$, Connected Dominating Set admits a polynomial-size α-approximate (bi-)kernel on all the aforementioned classes. Our results are in sharp contrast to the kernelization complexity of Connected Dominating Set, which is known to not admit a polynomial kernel even on 2-degenerate graphs and graphs of bounded expansion, unless NP⊆coNP/poly. We complement our results by the following conditional lower bound. We show that if a class C is somewhere dense and closed under taking subgraphs, then for some value of $r \in N$ there cannot exist an α-approximate bi-kernel for the (Connected) Distance-r Dominating Set problem on C for any $\alpha > 1$ (assuming the Gap Exponential Time Hypothesis).

Approximation and Parameterized Complexity of Minimax Approval Voting

Speaker: Krzysztof Sornat
Time: 12:00 – 12:30

We present three results on the complexity of Minimax Approval Voting. First, we study Minimax Approval Voting parameterized by the Hamming distance d from the solution to the votes. We show Minimax Approval Voting admits no algorithm running in time $O^*(2^{o(d \log d)})$, unless the Exponential Time Hypothesis (ETH) fails. This means that the $O^*(d^{2d})$ algorithm of Misra et al. [AAMAS 2015] is essentially optimal. Motivated by this, we then show a parameterized approximation scheme, running in time $O^*((3/\varepsilon)^{2d})$, which is essentially tight assuming ETH. Finally, we get a new polynomial-time randomized approximation scheme for Minimax Approval Voting, which runs in time $n^{O(log(1/\varepsilon)/\varepsilon^2)} \cdot \text{poly}(m)$, almost matching the running time of the fastest known PTAS for Closest String due to Ma and Sun [SIAM J. Comp. 2009].

Losing Treewidth by Separating Subsets

Speaker: Euiwoong Lee
Time: 14:00 – 14:30

We study the problem of deleting the smallest set S of vertices (resp. edges) from a given graph G such that the induced subgraph (resp. subgraph) $G \setminus S$ belongs to some class of graphs H. When graphs in H have treewidth bounded by t, we construct a framework for both vertex and edge deletion versions where approximation algorithms for these problems can be obtained from approximation algorithms for natural graph partitioning problems called k-Subset Vertex Separator and k-Subset Edge Separator.

For the vertex deletion, our framework, combined with the previous best result for k-Subset Vertex Separator, yields improved approximation ratios for fundamental problems such as k-Treewidth Vertex Deletion and Planar-F Vertex Deletion. For the edge deletion, we give improved approximation algorithms for k-Subset Edge Separator and their applications to various problems in bounded-degree graphs.

An FPT Algorithm Beating 2-Approximation for k-Cut

Speaker: Jason Li
Time: 14:30 – 15:00

We study the classical k-Cut problem, where given a graph G and an integer k, the goal is to remove the minimum number of edges to partition G into at least k connected components. Karger's random contraction algorithm gives an exact algorithm in time $O(n^{2k})$, which has not been improved yet. In terms of FPT algorithms, Kawarabayashi and Thorup proved that an exact solution can be computed in time $f(OPT) \cdot n^{O(1)}$. In terms of approximation, Saran and Vazirani gave a 2-approximation algorithm that runs in time $\text{poly}(n, k)$, but Manurangsi recently proved that we cannot achieve $(2 - \varepsilon)$ approximation for some fixed $\varepsilon > 0$ in time $\text{poly}(n, k)$, unless the Small Set Expansion Hypothesis is false.

We prove that there is a $(2 - \varepsilon)$-approximation algorithm for some fixed $\varepsilon > 0$ for k-Cut that runs in time $f(k) \cdot n^4$. This is the first FPT algorithm that is parameterized only by k and strictly improves the 2-approximation.

Parameterized (Approximate) Defective Coloring

Speaker: Michael Lampis
Time: 15:00 – 15:30

In Defective Coloring we are given a graph $G = (V, E)$ and two integers χ_d, Δ^* and are asked if we can partition V into χ_d color classes, so that each class induces a graph of maximum degree Δ^*. We investigate the complexity of this generalization of Coloring with respect to several well-studied graph parameters, and show that the problem is W-hard parameterized by treewidth, pathwidth, tree-depth, or feedback vertex set, if $\chi_d = 2$. As expected, this hardness can be extended to larger values of χ_d for most of these parameters, with one surprising exception: we show that the problem is FPT parameterized by feedback vertex set for any χ_d different from 2, and hence 2-coloring is the only hard case for this parameter. In addition to the above, we give an ETH-based lower bound for treewidth and pathwidth, showing that no algorithm can solve the problem in $n^{o(pw)}$, essentially matching the complexity of an algorithm obtained with standard techniques.

We complement these results by considering the problems approximability and show that, with respect to Δ^*, the problem admits an algorithm which for any $\varepsilon > 0$ runs in time $(tw/\varepsilon)^{O(tw)}$ and returns a solution with exactly the desired number of colors that approximates the optimal Δ^* within $(1 + \varepsilon)$. We also give a $(tw)^{O(tw)}$ algorithm which achieves the desired Δ^* exactly while 2-approximating the minimum value of χ_d. We show that this is close to optimal, by establishing that no FPT algorithm can (under standard assumptions) achieve a better than 3/2-approximation to χ_d, even when an extra constant additive error is also allowed.

These results have appeared in STACS 2018.

Characterizing Demand Graphs for (Fixed-Parameter) Shallow-Light Steiner Network

Speaker: Michael Dinitz
Time: 16:00 – 16:30

We consider the Shallow-Light Steiner Network problem from a fixed-parameter perspective. Given a graph G, a distance bound L, and p pairs of vertices $(s_1, t_1), \ldots, (s_p, t_p)$, the objective is to find a minimum-cost subgraph G' such that s_i and t_i have distance at most L in G' (for every $i \in [p]$). Our main result is on the fixed-parameter tractability of this problem with parameter p. We exactly characterize the demand structures that make the problem "easy", and give FPT algorithms for those cases. In all other cases, we show that the problem is W[1]-hard. We also extend our results to handle general edge lengths and costs, precisely characterizing which demands allow for good FPT approximation algorithms and which demands remain W[1]-hard even to approximate.

Beyond Metric Embedding: Approximating Group Steiner Trees on Bounded Treewidth Graphs

Speaker: Daniel Vaz
Time: 16:30 – 17:00

The Group Steiner Tree (GST) problem is a classical problem in combinatorial optimization and theoretical computer science. In the Edge-Weighted Group Steiner Tree (EW-GST) problem, we are given an undirected graph $G = (V, E)$ on n vertices with edge costs, a source vertex s and a collection of k subsets of vertices, called groups, $S_1, \ldots, S_k$. The goal is to find a minimum-cost subtree H in G that connects s to some vertex from each group S_i, for all $i = 1, 2, \ldots, k$. The Node-Weighted Group Steiner Tree (NW-GST) problem has the same setting, but the costs are associated with nodes. The goal is to find a minimum-cost node set X in V such that $G[X]$ connects every group to the source.

When G is a tree, both EW-GST and NW-GST admit a polynomial-time $O(\log n \log k)$ approximation algorithm due to the seminal result of Garg et al. [SODA'98 and J. Algorithm]. The matching hardness of $\log^{2-\varepsilon} n$ is known even for tree instances of EW-GST and NW-GST [Halperin and Krauthgamer STOC'03]. In general graphs, most of polynomial-time approximation algorithms for EW-GST reduce the problem to a tree instance using the metric-tree embedding, incurring a loss of $O(\log n)$ on the approximation factor [Bartal, FOCS'96; Fakcharoenphol et al., FOCS'03 and JCSS]. This yields an approximation ratio of $O(\log^2 n \log k)$ for EW-GST. Using metric-tree embedding, this factor cannot be improved: The loss of $\Omega(\log n)$ is necessary on some input graphs (e.g., grids and expanders). There are alternative approaches that avoid metric-tree embedding, e.g., the algorithm of [Chekuri and Pal, FOCS'05], which gives a tight approximation ratio, but none of which achieves polylogarithmic approximation in polynomial-time. This state of the art shows a clear lack of understanding of GST in general graphs beyond the metric-tree embedding technique. For NW-GST (for which the metric-tree embedding does not apply), not even a polynomial-time polylogarithmic approximation algorithm is known.

In this paper, we present $O(\log n \log k)$ approximation algorithms that run in time $n^{O(tw(G)^2)}$ for both NW-GST and EW-GST, where $tw(G)$ denotes the treewidth of graph G. The key to both results is a different type of "tree-embedding" that produces a tree of much bigger size, but *does not cause any loss on the approximation factor*. Our embedding is inspired by dynamic programming, a technique which is typically not applicable to Group Steiner problems.

Parameterized Approximation Schemes for Steiner Trees with Small Number of Steiner Vertices

Speaker: Tomáš Masařík
Time: 17:00 – 17:30

We study the Steiner Tree problem, in which a set of *terminal* vertices needs to be connected in the cheapest possible way in an edge-weighted graph. This problem has been extensively studied from the viewpoint of approximation and also parametrization. In particular, on one hand Steiner Tree is known to be APX-hard, and W[2]-hard on the other, if parameterized by the number of non-terminals (*Steiner vertices*) in the optimum solution. In contrast to this we give an *efficient parameterized approximation scheme* (EPAS), which circumvents both hardness results. Moreover, our methods imply the existence of a *polynomial size approximate kernelization scheme* (PSAKS) for the considered parameter.

We further study the parameterized approximability of other variants of Steiner Tree, such as Directed Steiner Tree and Steiner Forest. For neither of these an EPAS is likely to exist for the studied parameter: for Steiner Forest an easy observation shows that the problem is APX-hard, even if the input graph contains no Steiner vertices. For Directed Steiner Tree we prove that approximating within any function of the studied parameter is W[1]-hard. Nevertheless, we show that an EPAS exists for Unweighted Directed Steiner Tree, but a PSAKS does not. We also prove that there is an EPAS and a PSAKS for Steiner Forest if in addition to the number of Steiner vertices, the number of connected components of an optimal solution is considered to be a parameter.