

Publications of Thomas Bläsius

This document lists all peer-reviewed publications of Thomas Bläsius, Chair for Algorithm Engineering, Hasso Plattner Institute, Potsdam, Germany. This listing was automatically generated on October 29, 2024. An up-to-date version is available online at hpi.de/friedrich/docs/publist/blaesius.pdf.

Journal articles

- [1] Bläsius, T., Fischbeck, P., [On the External Validity of Average-Case Analyses of Graph Algorithms](#). In: *ACM Transactions on Algorithms* 20, 2024.

The number one criticism of average-case analysis is that we do not actually know the probability distribution of real-world inputs. Thus, analyzing an algorithm on some random model has no implications for practical performance. At its core, this criticism doubts the existence of external validity, i.e., it assumes that algorithmic behavior on the somewhat simple and clean models does not translate beyond the models to practical performance real-world input. With this paper, we provide a first step towards studying the question of external validity systematically. To this end, we evaluate the performance of six graph algorithms on a collection of 2740 sparse real-world networks depending on two properties; the heterogeneity (variance in the degree distribution) and locality (tendency of edges to connect vertices that are already close). We compare this with the performance on generated networks with varying locality and heterogeneity. We find that the performance in the idealized setting of network models translates surprisingly well to real-world networks. Moreover, heterogeneity and locality appear to be the core properties impacting the performance of many graph algorithms.

- [2] Bläsius, T., Friedrich, T., Krejca, M. S., Molitor, L., [The Impact of Geometry on Monochrome Regions in the Flip Schelling Process](#). In: *Computational Geometry (CGTA)* 108, pp. 101902, 2023.

Schelling's classical segregation model gives a coherent explanation for the wide-spread phenomenon of residential segregation. We introduce an agent-based saturated open-city variant, the Flip Schelling Process (FSP), in which agents, placed on a graph, have one out of two types and, based on the predominant type in their neighborhood, decide whether to change their types; similar to a new agent arriving as soon as another agent leaves the vertex. We investigate the probability that an edge u,v is monochrome, i.e., that both vertices u and v have the same type in the FSP, and we provide a general framework for analyzing the influence of the underlying graph topology on residential segregation. In particular, for two adjacent vertices, we show that a highly decisive common neighborhood, i.e., a common neighborhood where the absolute value of the difference between the number of vertices with different types is high, supports segregation and, moreover, that large common neighborhoods are more decisive. As an application, we study the expected behavior of the FSP on two common random graph models with and without geometry: (1) For random geometric graphs, we show that the existence of an edge u,v makes a highly decisive common neighborhood for u and v more likely. Based on this, we prove the existence of a constant $c>0$ such that the expected fraction of monochrome edges after the FSP is at least $1/2+c$. (2) For Erdős-Rényi graphs we show that large common neighborhoods are unlikely and that the expected fraction of monochrome edges after the FSP is at most $1/2+o(1)$. Our results indicate that the cluster structure of the underlying graph has a significant impact on the obtained segregation strength.

- [3] Bläsius, T., Fischbeck, P., Friedrich, T., Katzmann, M., [Solving Vertex Cover in Polynomial Time on Hyperbolic Random Graphs](#). In: *Theory of Computing Systems* 67, pp. 28–51, 2023.

The computational complexity of the VertexCover problem has been studied extensively. Most notably, it is NP-complete to find an optimal solution and typically NP-hard to find an approximation with reasonable factors. In contrast, recent experiments suggest that on many real-world networks the run time to solve VertexCover is way smaller than even the best known FPT-approaches can explain. We link these observations to two properties that are observed in many real-world networks, namely a heterogeneous degree distribution and high clustering. To formalize these properties and explain the observed behavior, we analyze how a branch-and-reduce algorithm performs on hyperbolic random graphs, which have become increasingly popular for modeling real-world networks. In fact, we are able to show that the VertexCover problem on hyperbolic random graphs can be solved in polynomial time, with high probability. The proof relies on interesting structural properties of hyperbolic random graphs. Since these predictions of the model are interesting in their own right, we conducted experiments on real-world networks showing that these properties are also observed in practice.

- [4] Bläsius, T., Freiberger, C., Friedrich, T., Katzmann, M., Montenegro-Retana, F., Thieffry, M., [Efficient Shortest Paths in Scale-Free Networks with Underlying Hyperbolic Geometry](#). In: *ACM Transactions on Algorithms*, pp. 1–32, 2022.

A standard approach to accelerating shortest path algorithms on networks is the bidirectional search, which explores the graph from the start and the destination, simultaneously. In practice this strategy performs particularly well on scale-free real-world networks. Such networks typically have a heterogeneous degree distribution (e.g., a power-law distribution) and high clustering (i.e., vertices with a common neighbor are likely to be connected themselves). These two properties can be obtained by assuming an underlying hyperbolic geometry. To explain the observed behavior of the bidirectional search, we analyze its running time on hyperbolic random graphs and prove that it is $O(n^{2-1/\alpha} + n^{1/(2\alpha)} + \delta_{\max})$ with high probability, where $\alpha \in (0.5, 1)$ controls the power-law exponent of the degree distribution, and δ_{\max} is the maximum degree. This bound is sublinear, improving the obvious worst-case linear bound. Although our analysis depends on the underlying geometry, the algorithm itself is oblivious to it.

- [5] Bläsius, T., Friedrich, T., Schirneck, M., [The Complexity of Dependency Detection and Discovery in Relational Databases](#). In: *Theoretical Computer Science* 900, pp. 79–96, 2022.

Multi-column dependencies in relational databases come associated with two different computational tasks. The detection problem is to decide whether a dependency of a certain type and size holds in a given database, the discovery problem asks to enumerate all valid dependencies of that type. We settle the complexity of both of these problems for unique column combinations (UCCs), functional dependencies (FDs), and inclusion dependencies (INDs). We show that the detection of UCCs and FDs is W[2]-complete when parameterized by the solution size. The discovery of inclusion-wise minimal UCCs is proven to be equivalent under parsimonious reductions to the transversal hypergraph problem of enumerating the minimal hitting sets of a hypergraph. The discovery of FDs is equivalent to the simultaneous enumeration of the hitting sets of multiple input hypergraphs. We further identify the detection of INDs as one of the first natural W[3]-complete problems. The discovery of maximal INDs is shown to be equivalent to enumerating the maximal satisfying assignments of antimonotone, 3-normalized Boolean formulas.

- [6] Bläsius, T., Friedrich, T., Lischeid, J., Meeks, K., Schirneck, M., [Efficiently Enumerating Hitting Sets of Hypergraphs Arising in Data Profiling](#). In: *Journal of Computer and System Sciences* 124, pp. 192–213, 2022.

The transversal hypergraph problem asks to enumerate the minimal hitting sets of a hypergraph. If the solutions have bounded size, Eiter and Gottlob [SICOMP’95] gave an algorithm running in output-polynomial time, but whose space requirement also scales with the output. We improve this to polynomial delay and space. Central to our approach is the extension problem, deciding for a set X of vertices whether it is contained in any minimal hitting set. We show that this is one of the first natural problems to be W[3]-complete. We give an algorithm for the extension problem running in time $O(m |X| + n)$ and prove a SETH-lower bound showing that this is close to optimal. We apply our enumeration method to the discovery problem of minimal unique column combinations from data profiling. Our empirical evaluation suggests that the algorithm outperforms its worst-case guarantees on hypergraphs stemming from real-world databases.

- [7] Birnick, J., Bläsius, T., Friedrich, T., Naumann, F., Papenbrock, T., Schirneck, M., [Hitting Set Enumeration with Partial Information for Unique Column Combination Discovery](#). In: *Proceedings of the VLDB Endowment* 13, pp. 2270–2283, 2020.

Unique column combinations (UCCs) are a fundamental concept in relational databases. They identify entities in the data and support various data management activities. Still, UCCs are usually not explicitly defined and need to be discovered. State-of-the-art data profiling algorithms are able to efficiently discover UCCs in moderately sized datasets, but they tend to fail on large and, in particular, on wide datasets due to run time and memory limitations. In this paper, we introduce HPIValid, a novel UCC discovery algorithm that implements a faster and more resource-saving search strategy. HPIValid models the metadata discovery as a hitting set enumeration problem in hypergraphs. In this way, it combines efficient discovery techniques from data profiling research with the most recent theoretical insights into enumeration algorithms. Our evaluation shows that HPIValid is not only orders of magnitude faster than related work, it also has a much smaller memory footprint.

- [8] Bläsius, T., Friedrich, T., Katzmann, M., Krohmer, A., [Hyperbolic Embeddings for Near-Optimal Greedy Routing](#). In: *Journal of Experimental Algorithmics (JEA)* 25, pp. 1–18, 2020.

Greedy routing computes paths between nodes in a network by successively moving to the neighbor closest to the target with respect to coordinates given by an embedding into some metric space. Its advantage is that only local information is used for routing decisions. We present different algorithms for generating graph embeddings into the hyperbolic plane that are well suited for greedy routing. In particular, our embeddings guarantee that greedy routing always succeeds in reaching the target, and we try to minimize the lengths of the resulting greedy paths. We evaluate our algorithm on multiple generated and real-world networks. For networks that are generally assumed to have a hidden underlying hyperbolic geometry, such as the Internet graph, we achieve near-optimal results (i.e., the resulting greedy paths are only slightly longer than the corresponding shortest paths). In the case of the Internet graph, they are only 6% longer when using our best algorithm, which greatly improves upon the previous best known embedding, whose creation required substantial manual intervention. In addition to measuring the stretch, we empirically evaluate our algorithms regarding the size of the coordinates of the resulting embeddings and observe how it impacts the success rate when coordinates are not represented with very high precision. Since numerical difficulties are a major issue when performing computations in the hyperbolic plane, we consider variations of our algorithm that improve the success rate when using coordinates with lower precision.

- [9] Bläsius, T., Rademacher, M., Rutter, I., [How to Draw a Planarization](#). In: *Graph Algorithms and Applications* 23, pp. 653–682, 2019.

We study the problem of computing straight-line drawings of non-planar graphs with few crossings. We assume that a crossing-minimization algorithm is applied first, yielding a planarization, i.e., a planar graph with a dummy vertex for each crossing, that fixes the topology of the resulting drawing. We present and evaluate two different approaches for drawing a planarization in such a way that the edges of the input graph are as straight as possible. The first approach is based on the planarity-preserving force-directed algorithm IMPRED [Simonetto et al. Computer Graphics Forum 2011], the second approach, which we call Geometric Planarization Drawing, iteratively moves vertices to their locally optimal positions in the given initial drawing. Our evaluation shows that both approaches significantly improve the initial drawing and that our geometric approach outperforms the force-directed approach. To the best of our knowledge, this is the first paper concerned with the generation of a straight-line drawing that respects an arbitrary planarization.

- [10] Bläsius, T., Friedrich, T., Krohmer, A., Laue, S., [Efficient Embedding of Scale-Free Graphs in the Hyperbolic Plane](#). In: *IEEE/ACM Transactions on Networking* 26, pp. 920–933, 2018.

Hyperbolic geometry appears to be intrinsic in many large real networks. We construct and implement a new maximum likelihood estimation algorithm that embeds scale-free graphs in the hyperbolic space. All previous approaches of similar embedding algorithms require at least a quadratic runtime. Our algorithm achieves quasilinear runtime, which makes it the first algorithm that can embed networks with hundreds of thousands of nodes in less than one hour. We demonstrate the performance of our algorithm on artificial and real networks. In all typical metrics, like Log-likelihood and greedy routing, our algorithm discovers embeddings that are very close to the ground truth.

- [11] Bläsius, T., Karrer, A., Rutter, I., [Simultaneous Embedding: Edge Orderings, Relative Positions, Cutvertices](#). In: *Algorithmica* 80, pp. 1214–1277, 2018.

A simultaneous embedding (with fixed edges) of two graphs G^1 and G^2 with common graph $G = G^1 \cap G^2$ is a pair of planar drawings of G^1 and G^2 that coincide on G . It is an open question whether there is a polynomial-time algorithm that decides whether two graphs admit a simultaneous embedding (problem SEFE). In this paper, we present two results. First, a set of three linear-time preprocessing algorithms that remove certain substructures from a given SEFE instance, producing a set of equivalent Sefe instances without such substructures. The structures we can remove are (1) cutvertices of the union graph $G^\cup = G^1 \cup G^2$, (2) most separating pairs of G^\cup ,

and (3) connected components of G that are biconnected but not a cycle. Second, we give an $O(n^3)$ -time algorithm solving Sefe for instances with the following restriction. Let u be a pole of a P-node μ in the SPQR-tree of a block of G^1 or G^2 . Then at most three virtual edges of μ may contain common edges incident to u . All algorithms extend to the sunflower case, i.e., to the case of more than two graphs pairwise intersecting in the same common graph.

- [12] Baum, M., Bläsius, T., Gemsa, A., Rutter, I., Wegner, F., [Scalable Exact Visualization of Isocontours in Road Networks via Minimum-Link Paths](#). In: *Journal of Computational Geometry* 9, pp. 27–73, 2018.

Isocontours in road networks represent the area that is reachable from a source within a given resource limit. We study the problem of computing accurate isocontours in realistic, large-scale networks. We propose isocontours represented by polygons with minimum number of segments that separate reachable and unreachable components of the network. Since the resulting problem is not known to be solvable in polynomial time, we introduce several heuristics that run in (almost) linear time and are simple enough to be implemented in practice. A key ingredient is a new practical linear-time algorithm for minimum-link paths in simple polygons. Experiments in a challenging realistic setting show excellent performance of our algorithms in practice, computing near-optimal solutions in a few milliseconds on average, even for long ranges.

- [13] Bläsius, T., Stumpf, P., Ueckerdt, T., [Local and Union Boxicity](#). In: *Discrete Mathematics* 341, pp. 1307–1315, 2018.

The boxicity $box(H)$ of a graph H is the smallest integer d such that H is the intersection of d interval graphs, or equivalently, that H is the intersection graph of axis-aligned boxes in R^d . These intersection representations can be interpreted as covering representations of the complement H^c of H with co-interval graphs, that is, complements of interval graphs. We follow the recent framework of global, local and folded covering numbers (Knauer and Ueckerdt, 2016) to define two new parameters: the local boxicity $box_l(H)$ and the union boxicity $box_u(H)$ of H . The union boxicity of H is the smallest d such that H^c can be covered with d vertex-disjoint unions of co-interval graphs, while the local boxicity of H is the smallest d such that H^c can be covered with co-interval graphs, at most d at every vertex. We show that for every graph H we have $box_l(H) \leq box_u(H) \leq box(H)$ and that each of these inequalities can be arbitrarily far apart. Moreover, we show that local and union boxicity are also characterized by intersection representations of appropriate axis-aligned boxes in R^d . We demonstrate with a few striking examples, that in a sense, the local boxicity is a better indication for the complexity of a graph, than the classical boxicity.

- [14] Bläsius, T., Friedrich, T., Krohmer, A., [Cliques in Hyperbolic Random Graphs](#). In: *Algorithmica* 80, pp. 2324–2344, 2018.

Most complex real world networks display scale-free features. This characteristic motivated the study of numerous random graph models with a power-law degree distribution. There is, however, no established and simple model which also has a high clustering of vertices as typically observed in real data. Hyperbolic random graphs bridge this gap. This natural model has recently been introduced by Krioukov et al. and has shown theoretically and empirically to fulfill all typical properties of real world networks, including power-law degree distribution and high clustering. We study cliques in hyperbolic random graphs G and present new results on the expected number of k -cliques $E[K_k]$ and the size of the largest clique $\omega(G)$. We observe that there is a phase transition at power-law exponent $\beta = 3$. More precisely, for $\beta \in (2, 3)$ we prove $E[K_k] = n^{k(3-\beta)/2} \Theta(k)^{-k}$ and $\omega(G) = \Theta(n^{(3-\beta)/2})$, while for $\beta \geq 3$ we prove $E[K_k] = n\Theta(k)^{-k}$ and $\omega(G) = \Theta(\log(n)/\log \log n)$. Furthermore, we show that for $\beta \geq 3$, cliques in hyperbolic random graphs can be computed in time $O(n)$. If the underlying geometry is known, cliques can be found with worst-case runtime $O(mn^{2.5})$ for all values of β .

- [15] Bläsius, T., Rutter, I., [A new perspective on clustered planarity as a combinatorial embedding problem](#). In: *Theoretical Computer Science* 609, pp. 306–315, 2016.

The clustered planarity problem (c-planarity) asks whether a hierarchically clustered graph admits a planar drawing such that the clusters can be nicely represented by regions. We introduce the cd-tree data structure and give a new characterization of c-planarity. It leads to efficient algorithms for c-planarity testing in the following cases. (i) Every cluster and every co-cluster (complement of a cluster) has at most two connected components. (ii) Every cluster has at most five outgoing edges. Moreover, the cd-tree reveals interesting connections between c-planarity and planarity with constraints on the order of edges around vertices. On one hand, this gives rise to a bunch of new open problems related to c-planarity, on the other hand it provides a new perspective on previous results.

- [16] Bläsius, T., Rutter, I., Wagner, D., [Optimal Orthogonal Graph Drawing with Convex Bend Costs](#). In: *ACM Transactions on Algorithms* 12, pp. 33, 2016.

Traditionally, the quality of orthogonal planar drawings is quantified by the total number of bends or the maximum number of bends per edge. However, this neglects that, in typical applications, edges have varying importance. We consider the problem OptimalFlexDraw that is defined as follows. Given a planar graph G on n vertices with maximum degree 4 (4-planar graph) and for each edge e a cost function $cost_e: N_0 \rightarrow R$ defining costs depending on the number of bends e has, compute a planar orthogonal drawing of G of minimum cost. In this generality OptimalFlexDraw is NP-hard. We show that it can be solved efficiently if (1) the cost function of each edge is convex and (2) the first bend on each edge does not cause any cost. Our algorithm takes time $O(n \cdot T_{flow}(n))$ and $O(n^2 \cdot T_{flow}(n))$ for biconnected and connected graphs, respectively, where $T_{flow}(n)$ denotes the time to compute a minimum-cost flow in a planar network with multiple sources and sinks. Our result is the first polynomial-time bend-optimization algorithm for general 4-planar graphs optimizing over all embeddings. Previous work considers restricted graph classes and unit costs.

- [17] Bläsius, T., Lehmann, S., Rutter, I., [Orthogonal Graph Drawing with Inflexible Edges](#). In: *Computational Geometry* 55, pp. 26–40, 2016.

We consider the problem of creating plane orthogonal drawings of 4-planar graphs (planar graphs with maximum degree 4) with constraints on the number of bends per edge. More precisely, we have a flexibility function assigning to each edge e a natural number $flex(e)$, its flexibility. The problem FlexDraw asks whether there exists an orthogonal drawing such that each edge e has at most $flex(e)$ bends. It is known that FlexDraw is NP-hard if $flex(e) = 0$ for every edge e . On the other hand, FlexDraw can be solved efficiently if $flex(e) \geq 1$ and is trivial if $flex(e) \geq 2$ for every edge e . To close the gap between the NP-hardness for $flex(e) = 0$ and the efficient algorithm for $flex(e) \geq 1$, we investigate the computational complexity of FlexDraw in case only few edges are inflexible (i.e., have flexibility 0). We show that for any $\epsilon > 0$ FlexDraw is NP-complete for instances with $O(n^\epsilon)$ inflexible edges with pairwise distance $\Omega(n^{1-\epsilon})$ (including the case where they induce a matching). On the other hand, we give an FPT-algorithm with running time $O(2^k \cdot n \cdot T_{flow}(n))$, where $T_{flow}(n)$ is the time necessary to compute a maximum flow in a planar flow network with multiple sources and sinks, and k is the number of inflexible edges having at least one endpoint of degree 4.

- [18] Bläsius, T., Rutter, I., [Simultaneous PQ-Ordering with Applications to Constrained Embedding Problems](#). In: *Transactions on Algorithms* 12, pp. 16, 2016.

In this article, we define and study the new problem of SIMULTANEOUS PQ-ORDERING. Its input consists of a set of PQ-trees, which represent sets of circular orders of their leaves, together with a set of child-parent relations between these PQ-trees, such that the leaves of the child form a subset of the leaves of the parent. SIMULTANEOUS PQ-ORDERING asks whether orders of the leaves of each of the trees can be chosen simultaneously; that is, for every child-parent relation, the order chosen for the parent is an extension of the order chosen for the child. We show that SIMULTANEOUS PQ-ORDERING is NP-complete in general, and we identify a family of instances that can be solved efficiently, the 2-fixed instances. We show that this result serves as a framework for several other problems that can be formulated as instances of SIMULTANEOUS PQ-ORDERING. In particular, we give linear-time algorithms for recognizing simultaneous interval graphs and extending partial interval representations. Moreover, we obtain a linear-time algorithm for PARTIALLY PQ-CONSTRAINED PLANARITY for biconnected graphs, which asks for a planar embedding in the presence of 16 PQ-trees that restrict the possible orderings of edges around vertices, and a quadratic-time algorithm for SIMULTANEOUS EMBEDDING WITH FIXED EDGES for biconnected graphs with a connected intersection. Both results can be extended to the case where the input graphs are not necessarily biconnected but have the property that each cutvertex is contained in at most two nontrivial blocks. This includes, for example, the case where both graphs have a maximum degree of 5.

- [19] Bläsius, T., Rutter, I., [Disconnectivity and relative positions in simultaneous embeddings](#). In: *Computational Geometry* 48, pp. 459–478, 2015.

For two planar graphs $G^1 = (V^1, E^1)$ and $G^2 = (V^2, E^2)$ sharing a common subgraph $G = G^1 \cap G^2$ the problem Simultaneous Embedding with Fixed Edges (SEFE) asks whether they admit planar drawings such that the common graph is drawn the same. Previous algorithms only work for cases where G is connected, and hence do not need to handle relative positions of connected components. We consider the problem where G , G^1 and G^2 are not necessarily connected. First, we show that a general instance of SEFE can be reduced in linear time to an equivalent instance where $V^1 = V^2$ and G^1 and G^2 are connected. Second, for the case where G consists of disjoint cycles, we introduce the CC-tree which represents all embeddings of G that extend to planar embeddings of G^1 . We show that CC-trees can be computed in linear time, and that their intersection is again a CC-tree. This yields a linear-time algorithm for SEFE if all k input graphs (possibly $k > 2$) pairwise share the same set of disjoint cycles. These results, including the CC-tree, extend to the case where G consists of arbitrary connected components, each with a fixed planar embedding on the sphere. Then the running time is $O(n^2)$.

- [20] Angelini, P., Bläsius, T., Rutter, I., [Testing Mutual duality of Planar graphs](#). In: *Computational Geometry and Applications* 24, pp. 325–346, 2014.

We introduce and study the problem Mutual Planar Duality, which asks for planar graphs G_1 and G_2 whether G_1 can be embedded such that its dual is isomorphic to G_2 . We show NP-completeness for general graphs and give a linear-time algorithm for biconnected graphs. We consider the common dual relation \sim , where $G_1 \sim G_2$ if and only if they admit embeddings that result in the same dual graph. We show that \sim is an equivalence relation on the set of biconnected graphs and devise a succinct, SPQR-tree-like representation of its equivalence classes. To solve Mutual Planar Duality for biconnected graphs, we show how to do isomorphism testing for two such representations in linear time. A special case of Mutual Planar Duality is testing whether a graph is self-dual. Our algorithm can handle the case of biconnected graphs in linear time and our NP-hardness proof extends to self-duality and also to map self-duality testing (which additionally requires to preserve the embedding).

- [21] Bläsius, T., Krug, M., Rutter, I., Wagner, D., [Orthogonal Graph Drawing with Flexibility Constraints](#). In: *Algorithmica* 68, pp. 859–885, 2014.

Traditionally, the quality of orthogonal planar drawings is quantified by either the total number of bends, or the maximum number of bends per edge. However, this neglects that in typical applications, edges have varying importance. In this work, we investigate an approach that allows to specify the maximum number of bends for each edge individually, depending on its importance. We consider a new problem called FlexDraw that is defined as follows. Given a planar graph $G = (V, E)$ on n vertices with maximum degree 4 and a function $flex: E \rightarrow \mathbb{N}_0$ that assigns a flexibility to each edge, does G admit a planar embedding on the grid such that each edge e has at most $flex(e)$ bends? Note that in our setting the combinatorial embedding of G is not fixed. FlexDraw directly extends the problem β -embeddability asking whether G can be embedded with at most β bends per edge. We give an algorithm with running-time $O(n^2)$ solving FlexDraw when the flexibility of each edge is positive. This includes 1-embeddability as a special case and thus closes the complexity gap between 0-embeddability, which is NP-hard to decide, and 2-embeddability, which is efficiently solvable since every planar graph with maximum degree 4 admits a 2-embedding except for the octahedron. In addition to the polynomial-time algorithm we show that FlexDraw is NP-hard even if the edges with flexibility 0 induce a tree or a union of disjoint stars.

Conference papers

- [22] Bläsius, T., Friedrich, T., Katzmann, M., Ruff, J., Zeif, Z., [On the Giant Component of Geometric Inhomogeneous Random Graphs](#). In: *European Symposium on Algorithms (ESA)*, pp. 20:1–20:13, 2023.

In this paper we study the threshold model of geometric inhomogeneous random graphs (GIRGs); a generative random graph model that is closely related to hyperbolic random graphs (HRGs). These models have been observed to capture complex real-world networks well with respect to the structural and algorithmic properties. Following comprehensive studies regarding their connectivity, i.e., which parts of the graphs are connected, we have a good understanding under which circumstances a giant component (containing a constant fraction of the graph) emerges. While previous results are rather technical and challenging to work with, the goal of this paper is to provide more accessible proofs. At the same time we significantly improve the previously known probabilistic guarantees, showing that GIRGs contain a giant component with probability $1 - \exp(-\Omega(n^{(3-\tau)/2}))$ for graph size n and a degree distribution with power-law exponent $\tau \in (2, 3)$. Based on that we additionally derive insights about the connectivity of certain induced subgraphs of GIRGs.

- [23] Bläsius, T., Friedrich, T., Katzmann, M., Stephan, D., [Strongly Hyperbolic Unit Disk Graphs](#). In: *Symposium Theoretical Aspects of Computer Science (STACS)*, pp. 13:1–13:17, 2023.

The class of Euclidean unit disk graphs is one of the most fundamental and well-studied graph classes with underlying geometry. In this paper, we identify this class as a special case in the broader class of hyperbolic unit disk graphs and introduce strongly hyperbolic unit disk graphs as a natural counterpart to the Euclidean variant. In contrast to the grid-like structures exhibited by Euclidean unit disk graphs, strongly hyperbolic networks feature hierarchical structures, which are also observed in complex real-world networks. We investigate basic properties of strongly hyperbolic unit disk graphs, including adjacencies and the formation of cliques, and utilize the derived insights to demonstrate that the class is useful for the development and analysis of graph algorithms. Specifically, we develop a simple greedy routing scheme and analyze its performance on strongly hyperbolic unit disk graphs in order to prove that routing can be performed more efficiently on such networks than in general.

- [24] Bläsius, T., Fischbeck, P., [On the External Validity of Average-Case Analyses of Graph Algorithms](#). In: *European Symposium on Algorithms (ESA)*, pp. 21:1–21:14, 2022.

The number one criticism of average-case analysis is that we do not actually know the probability distribution of real-world inputs. Thus, analyzing an algorithm on some random model has no implications for practical performance. At its core, this criticism doubts the existence of external validity, i.e., it assumes that algorithmic behavior on the somewhat simple and clean models does not translate beyond the models to practical performance real-world input. With this paper, we provide a first step towards studying the question of external validity systematically. To this end, we evaluate the performance of six graph algorithms on a collection of 2751 sparse real-world networks depending on two properties; the heterogeneity (variance in the degree distribution) and locality (tendency of edges to connect vertices that are already close). We compare this with the performance on generated networks with varying locality and heterogeneity. We find that the performance in the idealized setting of network models translates surprisingly well to real-world networks. Moreover, heterogeneity and locality appear to be the core properties impacting the performance of many graph algorithms.

- [25] Bläsius, T., Fischbeck, P., Gottesbüren, L., Hamann, M., Heuer, T., Spinner, J., Weyand, C., Wilhelm, M., [A Branch-and-Bound Algorithm for Cluster Editing](#). In: *Symposium on Experimental Algorithms (SEA)*, pp. 13:1–13:19, 2022.

The editing problem asks to transform a given graph into a disjoint union of cliques by inserting and deleting as few edges as possible. We describe and evaluate an exact branch-and-bound algorithm for cluster editing. For this, we introduce new reduction rules and adapt existing ones. Moreover, we generalize a known packing technique to obtain lower bounds and experimentally show that it contributes significantly to the performance of the solver. Our experiments further evaluate the effectiveness of the different reduction rules and examine the effects of structural properties of the input graph on solver performance. Our solver won the exact track of the 2021 PACE challenge.

- [26] Bläsius, T., Friedrich, T., Weyand, C., [Efficiently Computing Maximum Flows in Scale-Free Networks](#). In: *European Symposium on Algorithms (ESA)*, pp. 21:1–21:14, 2021.

We study the maximum-flow/minimum-cut problem on scale-free networks, i.e., graphs whose degree distribution follows a power-law. We propose a simple algorithm that capitalizes on the fact that often only a small fraction of such a network is relevant for the flow. At its core, our algorithm augments Dinitz’s algorithm with a balanced bidirectional search. Our experiments on a scale-free random network model indicate sublinear run time. On scale-free real-world networks, we outperform the commonly used highest-label Push-Relabel implementation by up to two orders of magnitude. Compared to Dinitz’s original algorithm, our modifications reduce the search space, e.g., by a factor of 275 on an autonomous systems graph. Beyond these good run times, our algorithm has an additional advantage compared to Push-Relabel. The latter computes a preflow, which makes the extraction of a minimum cut potentially more difficult. This is relevant, for example, for the computation of Gomory-Hu trees. On a social network with 70000 nodes, our algorithm computes the Gomory-Hu tree in 3 seconds compared to 12 minutes when using Push-Relabel.

- [27] Bläsius, T., Friedrich, T., Krejca, M. S., Molitor, L., [The Impact of Geometry on Monochrome Regions in the Flip Schelling Process](#). In: *International Symposium on Algorithms and Computation, (ISAAC) 2021*, pp. 29:1–29:17, 2021.

Schelling’s classical segregation model gives a coherent explanation for the wide-spread phenomenon of residential segregation. We consider an agent-based saturated open-city variant, the Flip-Schelling-Process (FSP), in which agents, placed on a graph, have one out of two types and, based on the predominant type in their neighborhood, decide whether to change their types; similar to a new agent arriving as soon as another agent leaves the vertex. We investigate the probability that an edge u,v is monochrome, i.e., that both vertices u and v have the same type in the FSP, and we provide a general framework for analyzing the influence of the underlying graph topology on residential segregation. In particular, for two adjacent vertices, we show that a highly decisive common neighborhood, i.e., a common neighborhood where the absolute value of the difference between the number of vertices with different types is high, supports segregation and moreover, that large common neighborhoods are more decisive. As an application, we study the expected behavior of the FSP on two common random graph models with and without geometry: (1) For random geometric graphs, we show that the existence of an edge u,v makes a highly decisive common neighborhood for u and v more likely. Based on this, we prove the existence of a constant $c > 0$ such that the expected fraction of monochrome edges after the FSP is at least $1/2 + c$. (2) For Erdős–Rényi graphs we show that large common neighborhoods are unlikely and that the expected fraction of monochrome edges after the FSP is at most $1/2 + o(1)$. Our results indicate that the cluster structure of the underlying graph has a significant impact on the obtained segregation strength.

- [28] Bläsius, T., Fischbeck, P., Gottesbüren, L., Hamann, M., Heuer, T., Spinner, J., Weyand, C., Wilhelm, M., [PACE Solver Description: The KaPoCE Exact Cluster Editing Algorithm](#). In: *International Symposium on Parameterized and Exact Computation (IPEC)*, pp. 27:1–27:3, 2021.

The cluster editing problem is to transform an input graph into a cluster graph by performing a minimum number of edge editing operations. A cluster graph is a graph where each connected component is a clique. An edit operation can be either adding a new edge or removing an existing edge. In this write-up we outline the core techniques used in the exact cluster editing algorithm of the KaPoCE framework (contains also a heuristic solver), submitted to the exact track of the 2021 PACE challenge.

- [29] Bläsius, T., Fischbeck, P., Gottesbüren, L., Hamann, M., Heuer, T., Spinner, J., Weyand, C., Wilhelm, M., [PACE Solver Description: KaPoCE: A Heuristic Cluster Editing Algorithm](#). In: *International Symposium on Parameterized and Exact Computation (IPEC)*, pp. 31:1–31:4, 2021.

The cluster editing problem is to transform an input graph into a cluster graph by performing a minimum number of edge editing operations. A cluster graph is a graph where each connected component is a clique. An edit operation can be either adding a new edge or removing an existing edge. In this write-up we outline the core techniques used in the heuristic cluster editing algorithm of the Karlsruhe and Potsdam Cluster Editing (KaPoCE) framework, submitted to the heuristic track of the 2021 PACE challenge.

- [30] Bläsius, T., Friedrich, T., Katzmann, M., [Force-Directed Embedding of Scale-Free Networks in the Hyperbolic Plane](#). In: *Symposium on Experimental Algorithms (SEA)*, pp. 22:1–22:18, 2021.

Force-directed drawing algorithms are the most commonly used approach to visualize networks. While they are usually very robust, the performance of Euclidean spring embedders decreases if the graph exhibits the high level of heterogeneity that typically occurs in scale-free real-world networks. As heterogeneity naturally emerges from hyperbolic geometry (in fact, scale-free networks are often perceived to have an underlying hyperbolic geometry), it is natural to embed them into the hyperbolic plane instead. Previous techniques that produce hyperbolic embeddings usually make assumptions about the given network, which (if not met) impairs the quality of the embedding. It is still an open problem to adapt force-directed embedding algorithms to make use of the heterogeneity of the hyperbolic plane, while also preserving their robustness. We identify fundamental differences between the behavior of spring embedders in Euclidean and hyperbolic space, and adapt the technique to take advantage of the heterogeneity of the hyperbolic plane.

- [31] Bläsius, T., Böther, M., Fischbeck, P., Friedrich, T., Gries, A., Hüffner, F., Kießig, O., Lenzner, P., Molitor, L., Schiller, L., Wells, A., Wietheger, S., [A Strategic Routing Framework and Algorithms for Computing Alternative Paths](#). In: *Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS)*, pp. 10:1–10:14, 2020.

Traditional navigation services find the fastest route for a single driver. Though always using the fastest route seems desirable for every individual, selfish behavior can have undesirable effects such as higher energy consumption and avoidable congestion, even leading to higher overall and individual travel times. In contrast, strategic routing aims at optimizing the traffic for all agents regarding a global optimization goal. We introduce a framework to formalize real-world strategic routing scenarios as algorithmic problems and study one of them, which we call Single Alternative Path (SAP), in detail. There, we are given an original route between a single origin–destination pair. The goal is to suggest an alternative route to all agents that optimizes the overall travel time under the assumption that the agents distribute among both routes according to a psychological model, for which we introduce the concept of Pareto-conformity. We show that the SAP problem is NP-complete, even for such models. Nonetheless, assuming Pareto-conformity, we give multiple algorithms for different variants of SAP, using multi-criteria shortest path algorithms as subroutines. Moreover, we prove that several natural models are in fact Pareto-conform. The implementation and evaluation of our algorithms serve as a proof of concept, showing that SAP can be solved in reasonable time even though the algorithms have exponential running time in the worst case.

- [32] Bläsius, T., Friedrich, T., Schirneck, M., [The Minimization of Random Hypergraphs](#). In: *European Symposium on Algorithms (ESA)*, pp. 21:1–21:15, 2020.

We investigate the maximum-entropy model $\mathcal{B}_{n,m,p}$ for random n -vertex, m -edge multi-hypergraphs with expected edge size pn . We show that the expected size of the minimization $\min(\mathcal{B}_{n,m,p})$, i.e., the number of inclusion-wise minimal edges of $\mathcal{B}_{n,m,p}$, undergoes a phase transition with respect to m . If m is at most $1/(1-p)^{(1-p)^n}$, then $E[|\min(\mathcal{B}_{n,m,p})|]$ is of order $\Theta(m)$, while for $m \geq 1/(1-p)^{(1-p+\varepsilon)^n}$ for any $\varepsilon > 0$, it is $\Theta(2^{H(\alpha)+(1-\alpha)\log_2 p} n / \sqrt{n})$. Here, H denotes the binary entropy function and $\alpha = -(\log_{1-p} m)/n$. The result implies that the maximum expected number of minimal edges over all m is $\Theta((1+p)^n / \sqrt{n})$. Our structural findings have algorithmic implications for minimizing an input hypergraph. This has applications in the profiling of relational databases as well as for the Orthogonal Vectors problem studied in fine-grained complexity. We make several technical contributions that are of independent interest in probability. First, we improve the Chernoff–Hoeffding theorem on the tail of the binomial distribution. In detail, we show that for a binomial variable $Y \sim \text{Bin}(n, p)$ and any $0 < x < p$, it holds that $P[Y \leq xn] = \Theta(2^{-D(x||p)n} / \sqrt{n})$, where D is the binary Kullback–Leibler divergence between Bernoulli distributions. We give explicit upper and lower bounds on the constants hidden in the big-O notation that hold for all n . Secondly, we establish the fact that the probability of a set of cardinality i being minimal after m i.i.d. maximum-entropy trials exhibits a sharp threshold behavior at $i^* = n + \log_{1-p} m$.

- [33] Bläsius, T., Fischbeck, P., Friedrich, T., Katzmann, M., [Solving Vertex Cover in Polynomial Time on Hyperbolic Random Graphs](#). In: *Symposium on the Theoretical Aspects of Computer Science (STACS)*, pp. 25:1–25:14, 2020.

The VertexCover problem is proven to be computationally hard in different ways: It is NP-complete to find an optimal solution and even NP-hard to find an approximation with reasonable factors. In contrast, recent experiments suggest that on many real-world networks the run time to solve VertexCover is way smaller than even the best known FPT-approaches can explain. Similarly, greedy algorithms deliver very good approximations to the optimal solution in practice. We link these observations to two properties that are observed in many real-world networks, namely a heterogeneous degree distribution and high clustering. To formalize these properties and explain the observed behavior, we analyze how a branch-and-reduce algorithm performs on hyperbolic random graphs, which have become increasingly popular for modeling real-world networks. In fact, we are able to show that the VertexCover problem on hyperbolic random graphs can be solved in polynomial time, with high probability. The proof relies on interesting structural properties of hyperbolic random graphs. Since these predictions of the model are interesting in their own right, we conducted experiments on real-world networks showing that these properties are also observed in practice. When utilizing the same structural properties in an adaptive greedy algorithm, further experiments suggest that, on real instances, this leads to better approximations than the standard greedy approach within reasonable time. We link these observations to two properties that are observed in many real-world networks, namely a heterogeneous degree distribution and high clustering. To formalize these properties and explain the observed behavior, we analyze how a branch-and-reduce algorithm performs on hyperbolic random graphs, which have become increasingly popular for modeling real-world networks. In fact, we are able to show that the VertexCover problem on hyperbolic random graphs can be solved in polynomial time, with high probability. The proof relies on interesting structural properties of hyperbolic random graphs. Since these predictions of the model are interesting in their own right, we conducted experiments on real-world networks showing that these properties are also observed in practice. When utilizing the same structural properties in an adaptive greedy algorithm, further experiments suggest that this leads to even better approximations than the standard greedy approach on real instances.

- [34] Bläsius, T., Friedrich, T., Lischeid, J., Meeks, K., Schirneck, M., [Efficiently Enumerating Hitting Sets of Hypergraphs Arising in Data Profiling](#). In: *Algorithm Engineering and Experiments (ALENEX)*, pp. 130–143, 2019.

We devise an enumeration method for inclusion-wise minimal hitting sets in hypergraphs. It has delay $O(m^{k^*+1} \cdot n^2)$ and uses linear space. Hereby, n is the number of vertices, m the number of hyperedges, and k^* the rank of the transversal hypergraph. In particular, on classes of hypergraphs for which the cardinality k^* of the largest minimal hitting set is bounded, the delay is polynomial. The algorithm solves the extension problem for minimal hitting sets as a subroutine. We show that the extension problem is W[3]-complete when parameterised by the cardinality of the set which is to be extended. For the subroutine, we give an algorithm that is optimal under the exponential time hypothesis. Despite these lower bounds, we provide empirical evidence showing that the enumeration outperforms

the theoretical worst-case guarantee on hypergraphs arising in the profiling of relational databases, namely, in the detection of unique column combinations.

- [35] Bläsius, T., Friedrich, T., Katzmann, M., Meyer, U., Penschuck, M., Weyand, C., [Efficiently Generating Geometric Inhomogeneous and Hyperbolic Random Graphs](#). In: *European Symposium on Algorithms (ESA)*, pp. 21:2–21:14, 2019. **EATCS Best Paper Award**.

Hyperbolic random graphs (HRG) and geometric inhomogeneous random graphs (GIRG) are two similar generative network models that were designed to resemble complex real world networks. In particular, they have a power-law degree distribution with controllable exponent β , and high clustering that can be controlled via the temperature T . We present the first implementation of an efficient GIRG generator running in expected linear time. Besides varying temperatures, it also supports underlying geometries of higher dimensions. It is capable of generating graphs with ten million edges in under a second on commodity hardware. The algorithm can be adapted to HRGs. Our resulting implementation is the fastest sequential HRG generator, despite the fact that we support non-zero temperatures. Though non-zero temperatures are crucial for many applications, most existing generators are restricted to $T = 0$. We also support parallelization, although this is not the focus of this paper. Moreover, we note that our generators draw from the correct probability distribution, i.e., they involve no approximation. Besides the generators themselves, we also provide an efficient algorithm to determine the non-trivial dependency between the average degree of the resulting graph and the input parameters of the GIRG model. This makes it possible to specify \bar{d} as input and obtain a graph with expected average degree \bar{d} . Moreover, we investigate the differences between HRGs and GIRGs, shedding new light on the nature of the relation between the two models. Although HRGs represent, in a certain sense, a special case of the GIRG model, we find that a straight-forward inclusion does not hold in practice. However, the difference is negligible for most use cases.

- [36] Bläsius, T., Friedrich, T., Sutton, A. M., [On the Empirical Time Complexity of Scale-Free 3-SAT at the Phase Transition](#). In: *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, pp. 117–134, 2019.

The hardness of formulas at the solubility phase transition of random propositional satisfiability (SAT) has been intensely studied for decades both empirically and theoretically. Solvers based on stochastic local search (SLS) appear to scale very well at the critical threshold, while complete backtracking solvers exhibit exponential scaling. On industrial SAT instances, this phenomenon is inverted: backtracking solvers can tackle large industrial problems, where SLS-based solvers appear to stall. Industrial instances exhibit sharply different structure than uniform random instances. Among many other properties, they are often heterogeneous in the sense that some variables appear in many while others appear in only few clauses. We conjecture that the heterogeneity of SAT formulas alone already contributes to the trade-off in performance between SLS solvers and complete backtracking solvers. We empirically determine how the run time of SLS vs. backtracking solvers depends on the heterogeneity of the input, which is controlled by drawing variables according to a scale-free distribution. Our experiments reveal that the efficiency of complete solvers at the phase transition is strongly related to the heterogeneity of the degree distribution. We report results that suggest the depth of satisfying assignments in complete search trees is influenced by the level of heterogeneity as measured by a power-law exponent. We also find that incomplete SLS solvers, which scale well on uniform instances, are not affected by heterogeneity. The main contribution of this paper utilizes the scale-free random 3-SAT model to isolate heterogeneity as an important factor in the scaling discrepancy between complete and SLS solvers at the uniform phase transition found in previous works.

- [37] Bläsius, T., Fischbeck, P., Friedrich, T., Schirneck, M., [Understanding the Effectiveness of Data Reduction in Public Transportation Networks](#). In: *Workshop on Algorithms and Models for the Web Graph (WAW)*, pp. 87–101, 2019.

Given a public transportation network of stations and connections, we want to find a minimum subset of stations such that each connection runs through a selected station. Although this problem is NP-hard in general, real-world instances are regularly solved almost completely by a set of simple reduction rules. To explain this behavior, we view transportation networks as hitting set instances and identify two characteristic properties, locality and heterogeneity. We then devise a randomized model to generate hitting set instances with adjustable properties. While the heterogeneity does influence the effectiveness of the reduction rules, the generated instances show that locality is the significant factor. Beyond that, we prove that the effectiveness of the reduction rules is independent of the underlying graph structure. Finally, we show that high locality is also prevalent in instances from other domains, facilitating a fast computation of minimum hitting sets.

- [38] Bläsius, T., Friedrich, T., Katzmann, M., Krohmer, A., [Hyperbolic Embeddings for Near-Optimal Greedy Routing](#). In: *Algorithm Engineering and Experiments (ALENEX)*, pp. 199–208, 2018.

Greedy routing computes paths between nodes in a network by successively moving to the neighbor closest to the target with respect to coordinates given by an embedding into some metric space. Its advantage is that only local information is used for routing decisions. We present different algorithms for generating graph embeddings into the hyperbolic plane that are well suited for greedy routing. In particular our embeddings guarantee that greedy routing always succeeds in reaching the target and we try to minimize the lengths of the resulting greedy paths. We evaluate our algorithm on multiple generated and real world networks. For networks that are generally assumed to have a hidden underlying hyperbolic geometry, such as the Internet graph, we achieve near-optimal results, i.e., the resulting greedy paths are only slightly longer than the corresponding shortest paths. In the case of the Internet graph, they are only 6% longer when using our best algorithm, which greatly improves upon the previous best known embedding, whose creation required substantial manual intervention.

- [39] Bläsius, T., Freiberger, C., Friedrich, T., Katzmann, M., Montenegro-Retana, F., Thieffry, M., [Efficient Shortest Paths in Scale-Free Networks with Underlying Hyperbolic Geometry](#). In: *International Colloquium on Automata, Languages, and Programming (ICALP)*, pp. 20:1–20:14, 2018.

A common way to accelerate shortest path algorithms on graphs is the use of a bidirectional search, which simultaneously explores the graph from the start and the destination. It has been observed recently that this strategy performs particularly well on scale-free real-world networks. Such networks typically have a heterogeneous degree distribution (e.g., a power-law distribution) and high clustering (i.e., vertices with a common neighbor are likely to be connected themselves). These two properties can be obtained by assuming an underlying hyperbolic geometry. To explain the observed behavior of the bidirectional search, we analyze its running time on hyperbolic random graphs and prove that it is $\tilde{O}(n^{2-1/\alpha} + n^{1/(2\alpha)} + \delta_{\max})$ with high probability, where $\alpha \in (0.5, 1)$ controls the power-law exponent of the degree distribution, and δ_{\max} is the maximum degree. This bound is sublinear, improving the obvious worst-case linear bound. Although our analysis depends on the underlying geometry, the algorithm itself is oblivious to it.

- [40] Bläsius, T., Eube, J., Feldtkeller, T., Friedrich, T., Krejca, M. S., Lagodzinski, J. A. G., Rothenberger, R., Severin, J., Sommer, F., Trautmann, J., [Memory-restricted Routing With Tiled Map Data](#). In: *Systems, Man, and Cybernetics (SMC)*, pp. 3347–3354, 2018.

Modern routing algorithms reduce query time by depending heavily on preprocessed data. The recently developed Navigation Data Standard (NDS) enforces a separation between algorithms and map data, rendering preprocessing inapplicable. Furthermore, map data is partitioned into tiles with respect to their geographic coordinates. With the limited memory found in portable devices, the number of tiles loaded becomes the major factor for run time. We study routing under these restrictions and present new algorithms as well as empirical evaluations. Our results show that, on average, the most efficient algorithm presented uses more than 20 times fewer tile loads than a normal A*.

- [41] Bläsius, T., Friedrich, T., Katzmann, M., Krohmer, A., Striebel, J., [Towards a Systematic Evaluation of Generative Network Models](#). In: *Workshop on Algorithms and Models for the Web Graph (WAW)*, pp. 99–114, 2018.

Generative graph models play an important role in network science. Unlike real-world networks, they are accessible for mathematical analysis and the number of available networks is not limited. The explanatory power of results on generative models, however, heavily depends on how realistic they are. We present a framework that allows for a systematic evaluation of generative network models. It is based on the question whether real-world networks can be distinguished from generated graphs with respect to certain graph parameters. As a proof of concept, we apply our framework to four popular random graph models (Erdős-Rényi, Barabási-Albert, Chung-Lu, and hyperbolic random graphs). Our experiments for example show that all four models are bad representations for Facebook’s social networks, while Chung-Lu and hyperbolic random graphs are good representations for other networks, with different strengths and weaknesses.

- [42] Bläsius, T., Radermacher, M., Rutter, I., [How to Draw a Planarization](#). In: *Current Trends in Theory and Practice of Computer Science (SOFSEM)*, pp. 295–308, 2017.

We study the problem of computing straight-line drawings of non-planar graphs with few crossings. We assume that a crossing-minimization algorithm is applied first, yielding a planarization, i.e., a planar graph with a dummy vertex for each crossing, that fixes the topology of the resulting drawing. We present and evaluate two different approaches for drawing a planarization in such a way that the edges of the input graph are as straight as possible. The first approach is based on the planarity-preserving force-directed algorithm ImPrEd, the second approach, which we call Geometric Planarization Drawing, iteratively moves vertices to their locally optimal positions in the given initial drawing.

- [43] Kovacs, R., Seufert, A., Wall, L., Chen, H.-T., Meinel, F., Müller, W., You, S., Brehm, M., Striebel, J., Kommana, Y., Popiak, A., Bläsius, T., Baudisch, P., [TrussFab: Fabricating Sturdy Large-Scale Structures on Desktop 3D Printers](#). In: *Human Factors in Computing Systems (CHI)*, pp. 2606–2616, 2017.

We present TrussFab, an integrated end-to-end system that allows users to fabricate large scale structures that are sturdy enough to carry human weight. TrussFab achieves the large scale by complementing 3D print with plastic bottles. It does not use these bottles as "bricks" though, but as beams that form structurally sound node-link structures, also known as trusses, allowing it to handle the forces resulting from scale and load. TrussFab embodies the required engineering knowledge, allowing non-engineers to design such structures and to validate their design using integrated structural analysis. We have used TrussFab to design and fabricate tables and chairs, a 2.5 m long bridge strong enough to carry a human, a functional boat that seats two, and a 5 m diameter dome.

- [44] Baum, M., Bläsius, T., Gemsa, A., Rutter, I., Wegner, F., [Scalable Exact Visualization of Isocontours in Road Networks via Minimum-Link Paths](#). In: *European Symposium on Algorithms (ESA)*, pp. 7:1–7:18, 2016.

Isocontours in road networks represent the area that is reachable from a source within a given resource limit. We study the problem of computing accurate isocontours in realistic, large-scale networks. We propose isocontours represented by polygons with minimum number of segments that separate reachable and unreachable components of the network. Since the resulting problem is not known to be solvable in polynomial time, we introduce several heuristics that run in (almost) linear time and are simple enough to be implemented in practice. A key ingredient is a new practical linear-time algorithm for minimum-link paths in simple polygons. Experiments in a challenging realistic setting show excellent performance of our algorithms in practice, computing near-optimal solutions in a few milliseconds on average, even for long ranges.

- [45] Bläsius, T., Friedrich, T., Krohmer, A., [Hyperbolic Random Graphs: Separators and Treewidth](#). In: *European Symposium on Algorithms (ESA)*, pp. 15:1–15:16, 2016.

When designing and analyzing algorithms, one can obtain better and more realistic results for practical instances by assuming a certain probability distribution on the input. The worst-case run-time is then replaced by the expected run-time or by bounds that hold with high probability (whp), i.e., with probability $1 - O(1/n)$, on the random input. Hyperbolic random graphs can be used to model complex real-world networks as they share many important properties such as a small diameter, a large clustering coefficient, and a power-law degree-distribution. Divide and conquer is an important algorithmic design principle that works particularly well if the instance admits small separators. We show that hyperbolic random graphs in fact have comparatively small separators. More precisely, we show that a hyperbolic random graph can be expected to have a balanced separator hierarchy with separators of size $O(\sqrt{n^{(3-\beta)}})$, $O(\log n)$, and $O(1)$ if $2 < \beta < 3$, $\beta = 3$ and $3 < \beta$, respectively (β is the power-law exponent). We infer that these graphs have whp a treewidth of $O(\sqrt{n^{(3-\beta)}})$, $O(\log^2 n)$, and $O(\log n)$, respectively. For $2 < \beta < 3$, this matches a known lower bound. For the more realistic (but harder to analyze) binomial model, we still prove a sublinear bound on the treewidth. To demonstrate the usefulness of our results, we apply them to obtain fast matching algorithms and an approximation scheme for Independent Set.

- [46] Bläsius, T., Friedrich, T., Krohmer, A., Laue, S., [Efficient Embedding of Scale-Free Graphs in the Hyperbolic Plane](#). In: *European Symposium on Algorithms (ESA)*, pp. 16:1–16:18, 2016. **EATCS Best Paper Award**.

Hyperbolic geometry appears to be intrinsic in many large real networks. We construct and implement a new maximum likelihood estimation algorithm that embeds scale-free graphs in the hyperbolic space. All previous approaches of similar embedding algorithms require a runtime of $\Omega(n^2)$. Our algorithm achieves quasilinear runtime, which makes it the first algorithm that can embed networks with hundreds of thousands of nodes in less than one hour. We demonstrate the performance of our algorithm on artificial and real networks. In all typical metrics like Log-likelihood and greedy routing our algorithm discovers embeddings that are very close to the ground truth.

- [47] Bläsius, T., Friedrich, T., Schirneck, M., [The Parameterized Complexity of Dependency Detection in Relational Databases](#). In: *International Symposium on Parameterized and Exact Computation (IPEC)*, pp. 6:1–6:13, 2016.

We study the parameterized complexity of classical problems that arise in the profiling of relational data. Namely, we characterize the complexity of detecting unique column combinations (candidate keys), functional dependencies, and inclusion dependencies with the solution size as parameter. While the discovery of uniques and functional dependencies, respectively, turns out to be W[2]-complete, the detection of inclusion dependencies is one of the first natural problems proven to be complete for the class W[3]. As a side effect, our reductions give insights into the complexity of enumerating all minimal unique column combinations or functional dependencies.

- [48] Alam, M. J., Bläsius, T., Rutter, I., Ueckerdt, T., Wolff, A., [Pixel and Voxel Representations of Graphs](#). In: *Graph Drawing (GD)*, pp. 472–486, 2015.

We study contact representations for graphs, which we call pixel representations in 2D and voxel representations in 3D. Our representations are based on the unit square grid whose cells we call pixels in 2D and voxels in 3D. Two pixels are adjacent if they share an edge, two voxels if they share a face. We call a connected set of pixels or voxels a blob. Given a graph, we represent its vertices by disjoint blobs such that two blobs contain adjacent pixels or voxels if and only if the corresponding vertices are adjacent. We are interested in the size of a representation, which is the number of pixels or voxels it consists of. We first show that finding minimum-size representations is NP-complete. Then, we bound representation sizes needed for certain graph classes. In 2D, we show that, for k -outerplanar graphs with n vertices, $\Theta(kn)$ pixels are always sufficient and sometimes necessary. In particular, outerplanar graphs can be represented with a linear number of pixels, whereas general planar graphs sometimes need a quadratic number. In 3D, $\Theta(n^2)$ voxels are always sufficient and sometimes necessary for any n -vertex graph. We improve this bound to $\Theta(n \cdot \tau)$ for graphs of treewidth τ and to $O((g+1)^2 n \log^2 n)$ for graphs of genus g . In particular, planar graphs admit representations with $O(n \log^2 n)$ voxels.

- [49] Bläsius, T., Lehmann, S., Rutter, I., [Orthogonal Graph Drawing with Inflexible Edges](#). In: *International Conference on Algorithms and Complexity (CIAC)*, pp. 61–73, 2015.

We consider the problem of creating plane orthogonal drawings of 4-planar graphs (planar graphs with maximum degree 4) with constraints on the number of bends per edge. More precisely, we have a flexibility function assigning to each edge e a natural number $flex(e)$, its flexibility. The problem FlexDraw asks whether there exists an orthogonal drawing such that each edge e has at most $flex(e)$ bends. It is known that FlexDraw is NP-hard if $flex(e) = 0$ for every edge e [7]. On the other hand, FlexDraw can be solved efficiently if $flex(e) \geq 1$ [2] and is trivial if $flex(e) \geq 2$ [1] for every edge e . To close the gap between the NP-hardness for $flex(e) = 0$ and the efficient algorithm for $flex(e) \geq 1$, we investigate the computational complexity of FlexDraw in case only few edges are inflexible (i.e., have flexibility 0). We show that for any $\epsilon > 0$ FlexDraw is NP-complete for instances with $O(n^\epsilon)$ inflexible edges with pairwise distance $\Omega(n^{1-\epsilon})$ (including the case where they induce a matching). On the other hand, we give an FPT-algorithm with running time $O(2^k \cdot n \cdot T_{flow}(n))$, where $T_{flow}(n)$ is the time necessary to compute a maximum flow in a planar flow network with multiple sources and sinks, and k is the number of inflexible edges having at least one endpoint of degree 4.

- [50] Bläsius, T., Brückner, G., Rutter, I., [Complexity of Higher-Degree Orthogonal Graph Embedding in the Kandinsky Model](#). In: *European Symposium on Algorithms (ESA)*, pp. 161–172, 2014.

We show that finding orthogonal grid-embeddings of plane graphs (planar with fixed combinatorial embedding) with the minimum number of bends in the so-called Kandinsky model (which allows vertices of degree >4) is NP-complete, thus solving a long-standing open problem. On the positive side, we give an efficient algorithm for several restricted variants, such as graphs of bounded branch width and a subexponential exact algorithm for general plane graphs.

- [51] Bläsius, T., Rutter, I., [A New Perspective on Clustered Planarity as a Combinatorial Embedding Problem](#). In: *Graph Drawing (GD)*, pp. 440–451, 2014.

The clustered planarity problem (c -planarity) asks whether a hierarchically clustered graph admits a planar drawing such that the clusters can be nicely represented by regions. We introduce the cd -tree data structure and give a new characterization of c -planarity. It leads to efficient algorithms for c -planarity testing in the following cases. (i) Every cluster and every co-cluster has at most two connected components. (ii) Every cluster has at most five outgoing edges. Moreover, the cd -tree reveals interesting connections between c -planarity and planarity with constraints on the order of edges around vertices. On one hand, this gives rise to a bunch of new open problems related to c -planarity, on the other hand it provides a new perspective on previous results.

- [52] Biedl, T. C., Bläsius, T., Niedermann, B., Nöllenburg, M., Prutkin, R., Rutter, I., [Using ILP/SAT to Determine Pathwidth, Visibility Representations, and other Grid-Based Graph Drawings](#). In: *Graph Drawing (GD)*, pp. 460–471, 2013.

We present a simple and versatile formulation of grid-based graph representation problems as an integer linear program (ILP) and a corresponding SAT instance. In a grid-based representation vertices and edges correspond to axis-parallel boxes on an underlying integer grid; boxes can be further constrained in their shapes and interactions by additional problem-specific constraints. We describe a general d -dimensional model for grid representation problems. This model can be used to solve a variety of NP-hard graph problems, including pathwidth, bandwidth, optimum st-orientation, area-minimal (bar- k) visibility representation, boxicity- k graphs and others. We implemented SAT-models for all of the above problems and evaluated them on the Rome graphs collection. The experiments show that our model successfully solves NP-hard problems within few minutes on small to medium-size Rome graphs.

- [53] Bläsius, T., Karrer, A., Rutter, I., [Simultaneous Embedding: Edge Orderings, Relative Positions, Cutvertices](#). In: *Graph Drawing (GD)*, pp. 220–231, 2013.

A simultaneous embedding (with fixed edges) of two graphs G^1 and G^2 with common graph $G = G^1 \cap G^2$ is a pair of planar drawings of G^1 and G^2 that coincide on G . It is an open question whether there is a polynomial-time algorithm that decides whether two graphs admit a simultaneous embedding (problem Sefe). In this paper, we present two results. First, a set of three linear-time preprocessing algorithms that remove certain substructures from a given Sefe instance, producing a set of equivalent Sefe instances without such substructures. The structures we can remove are (1) cutvertices of the union graph $G^\cup = G^1 \cup G^2$, (2) most separating pairs of G^\cup , and (3) connected components of G that are biconnected but not a cycle. Second, we give an $O(n^3)$ -time algorithm solving Sefe for instances with the following restriction. Let u be a pole of a P -node μ in the SPQR-tree of a block of G^1 or G^2 . Then at most three virtual edges of μ may contain common edges incident to u . All algorithms extend to the sunflower case, i.e., to the case of more than three graphs pairwise intersecting in the same common graph.

- [54] Bläsius, T., Rutter, I., Wagner, D., [Optimal Orthogonal Graph Drawing with Convex Bend Costs](#). In: *International Colloquium on Automata, Languages, and Programming (ICALP)*, pp. 184–195, 2013.

Traditionally, the quality of orthogonal planar drawings is quantified by the total number of bends or the maximum number of bends per edge. However, this neglects that, in typical applications, edges have varying importance. We consider the problem OptimalFlexDraw that is defined as follows. Given a planar graph G on n vertices with maximum degree 4 (4-planar graph) and for each edge e a cost function $cost_e: N_0 \rightarrow R$ defining costs depending on the number of bends e has, compute a planar orthogonal drawing of G of minimum cost. In this generality OptimalFlexDraw is NP-hard. We show that it can be solved efficiently if (1) the cost function of each edge is convex and (2) the first bend on each edge does not cause any cost. Our algorithm takes time $O(n, \cdot, T_{flow}(n))$ and $O(n^2, \cdot, T_{flow}(n))$ for biconnected and connected graphs, respectively, where $T_{flow}(n)$ denotes the time to compute a minimum-cost flow in a planar network with multiple sources and sinks. Our result is the first polynomial-time bend-optimization algorithm for general 4-planar graphs optimizing over all embeddings. Previous work considers restricted graph classes and unit costs.

- [55] Angelini, P., Bläsius, T., Rutter, I., [Testing Mutual Duality of Planar Graphs](#). In: *International Symposium on Algorithms and Computation (ISAAC)*, pp. 350–360, 2013.

We introduce and study the problem Mutual Planar Duality, which asks for planar graphs G_1 and G_2 whether G_1 can be embedded such that its dual is isomorphic to G_2 . We show NP-completeness for general graphs and give a linear-time algorithm for biconnected graphs. We consider the common dual relation \sim , where $G_1 \sim G_2$ if and only they admit embeddings that result in the same dual graph. We show that \sim is an equivalence relation on the set of biconnected graphs and devise a succinct, SPQR-tree-like representation of its equivalence classes. To solve Mutual Planar Duality for biconnected graphs, we show how to do isomorphism testing for two such representations in linear time. A special case of Mutual Planar Duality is testing whether a graph is self-dual. Our algorithm can handle the case of biconnected graphs in linear time and our NP-hardness proof extends to self-duality and also to map self-duality testing (which additionally requires to preserve the embedding).

- [56] Bläsius, T., Rutter, I., [Simultaneous PQ-Ordering with Applications to Constrained Embedding Problems](#). In: *Symposium on Discrete Algorithms (SODA)*, pp. 1030–1043, 2013.

In this article, we define and study the new problem of SIMULTANEOUS PQ-ORDERING. Its input consists of a set of PQ-trees, which represent sets of circular orders of their leaves, together with a set of child-parent relations between these PQ-trees, such that the leaves of the child form a subset of the leaves of the parent. SIMULTANEOUS PQ-ORDERING asks whether orders of the leaves of each of the trees can be chosen simultaneously; that is, for every child-parent relation, the order chosen for the parent is an extension of the order chosen for the child. We show that SIMULTANEOUS PQ-ORDERING is NP-complete in general, and we identify a family of instances that can be solved efficiently, the 2-fixed instances. We show that this result serves as a framework for several other problems that can be formulated as instances of SIMULTANEOUS PQ-ORDERING. In particular, we give linear-time algorithms for recognizing simultaneous interval graphs and extending partial interval representations. Moreover, we obtain a linear-time algorithm for PARTIALLY PQ-CONSTRAINED PLANARITY for biconnected graphs, which asks for a planar embedding in the presence of 16 PQ-trees that restrict the possible orderings of edges around vertices, and a quadratic-time algorithm for SIMULTANEOUS EMBEDDING WITH FIXED EDGES for biconnected graphs with a connected intersection. Both results can be extended to the case where the input graphs are not necessarily biconnected but have the property that each cutvertex is contained in at most two nontrivial blocks. This includes, for example, the case where both graphs have a maximum degree of 5.

- [57] Bläsius, T., Rutter, I., [Disconnectivity and Relative Positions in Simultaneous Embeddings](#). In: *Graph Drawing (GD)*, pp. 31–42, 2012.

For two planar graphs $G^1 = (V^1, E^1)$ and $G^2 = (V^2, E^2)$ sharing a common subgraph $G = G^1 \cap G^2$ the problem Simultaneous Embedding with Fixed Edges (SEFE) asks whether they admit planar drawings such that the common graph is drawn the same. Previous algorithms only work for cases where G is connected, and hence do not need to handle relative positions of connected components. We consider the problem where G, G^1 and G^2 are not necessarily connected. First, we show that a general instance of SEFE can be reduced in linear time to an equivalent instance where $V^1 = V^2$ and G^1 and G^2 are connected. Second, for the case where G consists of disjoint cycles, we introduce the CC-tree which represents all embeddings of G that extend to planar embeddings of G^1 . We show that CC-trees can be computed in linear time, and that their intersection is again a CC-tree. This yields a linear-time algorithm for SEFE if all k input graphs (possibly $k > 2$) pairwise share the same set of disjoint cycles. These results, including the CC-tree, extend to the case where G consists of arbitrary connected components, each with a fixed planar embedding on the sphere. Then the running time is $O(n^2)$.

- [58] Bläsius, T., Krug, M., Rutter, I., Wagner, D., [Orthogonal Graph Drawing with Flexibility Constraints](#). In: *Graph Drawing (GD)*, pp. 92–104, 2010.