Publications of Maximilian Katzmann

This document lists all peer-reviewed publications of Maximilian Katzmann, Chair for Algorithm Engineering, Hasso Plattner Institute, Potsdam, Germany. This listing was automatically generated on September 19, 2019. An up-to-date version is available online at hpi.de/friedrich/docs/publist/katzmann.pdf.

Journal articles

[1] Friedrich, T., Katzmann, M., Krohmer, A., Unbounded Discrepancy of Deterministic Random Walks on Grids. In: *SIAM Journal on Discrete Mathematics* 32, pp. 2441–2452, 2018.

Random walks are frequently used in randomized algorithms. We study a derandomized variant of a random walk on graphs, called rotor-router model. In this model, instead of distributing tokens randomly, each vertex serves its neighbors in a fixed deterministic order. For most setups, both processes behave remarkably similar: Starting with the same initial configuration, the number of tokens in the rotor-router model deviates only slightly from the expected number of tokens on the corresponding vertex in the random walk model. The maximal difference over all vertices and all times is called single vertex discrepancy. Cooper and Spencer (2006) showed that on $\mathbb{Z}^d$ the single vertex discrepancy is only a constant c_d. Other authors also determined the precise value of c_d for $d = 1, 2$. All these results, however, assume that initially all tokens are only placed on one partition of the bipartite graph $\mathbb{Z}^d$. We show that this assumption is crucial by proving that otherwise the single vertex discrepancy can become arbitrarily large. For all dimensions $d \geq 1$ and arbitrary discrepancies $\ell \geq 0$, we construct configurations that reach a discrepancy of at least ℓ.

Conference papers

[2] 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.

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.

[3] 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 wold 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.

[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: *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 $\delta_{\max}$ 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., 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.

[6] Katzmann, M., Komusiewicz, C., Systematic Exploration of Larger Local Search Neighborhoods for the Minimum Vertex Cover Problem. In: *Conference on Artificial Intelligence (AAAI)*, pp. 846–852, 2017.

We investigate the potential of exhaustively exploring larger neighborhoods in local search algorithms for Minimum Vertex Cover. More precisely, we study whether, for moderate values of k, it is feasible and worthwhile to determine, given a graph G with vertex cover C, if there is a k-swap S such that $(C \setminus S) \cup (S \setminus C)$ is a smaller vertex cover of G. First, we describe an algorithm running in $\Delta^{O(k)} \cdot n$ time for searching the k-swap neighborhood on n-vertex graphs with maximum degree Δ. Then, we demonstrate that, by devising additional pruning rules that decrease the size of the search space, this algorithm can be implemented so that it solves the problem quickly for $k \approx 20$. Finally, we show that it is worthwhile to consider moderately-sized k-swap neighborhoods. For our benchmark data set, we show that when combining our algorithm with a hill-climbing approach, the solution quality improves quickly with the radius k of the local search neighborhood and that in most cases optimal solutions can be found by setting $k = 21$.

[7] Friedrich, T., Katzmann, M., Krohmer, A., Unbounded Discrepancy of Deterministic Random Walks on Grids. In: *International Symposium on Algorithms and Computation (ISAAC)*, pp. 212–222, 2015.

Random walks are frequently used in randomized algorithms. We study a derandomized variant of a random walk on graphs, called rotor-router model. In this model, instead of distributing tokens randomly, each vertex serves its neighbors in a fixed deterministic order. For most setups, both processes behave remarkably similar: Starting with the same initial configuration, the number of tokens in the rotor-router model deviates only slightly from the expected number of tokens on the corresponding vertex in the random walk model. The maximal difference over all vertices and all times is called single vertex discrepancy. Cooper and Spencer (2006) showed that on $\mathbb{Z}^d$ the single vertex discrepancy is only a constant c_d. Other authors also determined the precise value of c_d for $d = 1, 2$. All these results, however, assume that initially all tokens are only placed on one partition of the bipartite graph $\mathbb{Z}^d$. We show that this assumption is crucial by proving that otherwise the single vertex discrepancy can become arbitrarily large. For all dimensions $d \geq 1$ and arbitrary discrepancies $\ell \geq 0$, we construct configurations that reach a discrepancy of at least ℓ.