

# Publications of Simon Krogmann

This document lists all peer-reviewed publications of [Simon Krogmann](#), Chair for Algorithm Engineering, Hasso Plattner Institute, Potsdam, Germany. This listing was automatically generated on April 18, 2024. An up-to-date version is available online at [hpi.de/friedrich/docs/publist/krogmann.pdf](https://hpi.de/friedrich/docs/publist/krogmann.pdf).

## Journal articles

- [1] Friedrich, T., Krejca, M. S., Rothenberger, R., Arndt, T., Hafner, D., Kellermeier, T., Krogmann, S., Razmjou, A., [Routing for On-Street Parking Search using Probabilistic Data](#). In: *AI Communications* 32, pp. 113–124, 2019.

A significant percentage of urban traffic is caused by the search for parking spots. One possible approach to improve this situation is to guide drivers along routes which are likely to have free parking spots. The task of finding such a route can be modeled as a probabilistic graph problem which is NP-complete. Thus, we propose heuristic approaches for solving this problem and evaluate them experimentally. For this, we use probabilities of finding a parking spot, which are based on publicly available empirical data from TomTom International B.V. Additionally, we propose a heuristic that relies exclusively on conventional road attributes. Our experiments show that this algorithm comes close to the baseline by a factor of 1.3 in our cost measure. Last, we complement our experiments with results from a field study, comparing the success rates of our algorithms against real human drivers.

## Conference papers

- [2] Krogmann, S., Lenzner, P., Skopalik, A., Uetz, M., Vos, M. C., [Equilibria in Two-Stage Facility Location with Atomic Clients](#). In: *International Joint Conference on Artificial Intelligence (IJCAI)*, 2024.

We consider competitive facility location as a two-stage multi-agent system with two types of clients. For a given host graph with weighted clients on the vertices, first facility agents strategically select vertices for opening their facilities. Then, the clients strategically select which of the opened facilities in their neighborhood to patronize. Facilities want to attract as much client weight as possible, clients want to minimize congestion on the chosen facility. All recently studied versions of this model assume that clients can split their weight strategically. We consider clients with unsplitable weights, but allow mixed strategies. So clients may randomize over which facility to patronize. Besides modeling a natural client behavior, this subtle change yields drastic changes, e.g., for a given facility placement, qualitatively different client equilibria are possible. As our main result, we show that pure subgame perfect equilibria always exist if all client weights are identical. For this, we use a novel potential function argument, employing a hierarchical classification of the clients and sophisticated rounding in each step. In contrast, for non-identical clients, we show that deciding the existence of even approximately stable states is computationally intractable. On the positive side, we give a tight bound of 2 on the price of anarchy which implies high social welfare of equilibria, if they exist.

- [3] Bilò, D., Choudhary, K., Cohen, S., Friedrich, T., Krogmann, S., Schirneck, M., [Compact Distance Oracles with Large Sensitivity and Low Stretch](#). In: *Algorithms and Data Structures Symposium (WADS)*, pp. 149–163, 2023.

An  $f$ -edge fault-tolerant distance sensitive oracle ( $f$ -DSO) with stretch  $\sigma \geq 1$  is a data-structure that preprocesses an input graph  $G = (V, E)$ . When queried with the triple  $(s, t, F)$ , where  $s, t \in V$  and  $F \subseteq E$  contains at most  $f$  edges of  $G$ , the oracle returns an estimate  $\widehat{d}_{G-F}(s, t)$  of the distance  $d_{G-F}(s, t)$  between  $s$  and  $t$  in the graph  $G - F$  such that  $d_{G-F}(s, t) \leq \widehat{d}_{G-F}(s, t) \leq \sigma \cdot d_{G-F}(s, t)$ . For any positive integer  $k \geq 2$  and any  $0 < \alpha < 1$ , we present an  $f$ -DSO with sensitivity  $f = o(\log n / \log \log n)$ , stretch  $2k - 1$ , space  $O(n^{1+\frac{1}{k}+\alpha+o(1)})$ , and an  $\widetilde{O}(n^{1+\frac{1}{k}-\frac{\alpha}{k(F+1)}})$  query time. Prior to our work, there were only three known  $f$ -DSOs with subquadratic space. The first one by Chechik et al. [Algorithmica 2012] has a stretch of  $(8k - 2)(f + 1)$ , depending on  $f$ . Another approach is storing an  $f$ -edge fault-tolerant  $(2k - 1)$ -spanner of  $G$ . The bottleneck is the large query time due to the size of any such spanner, which is  $\Omega(n^{1+1/k})$  under the Erdős girth conjecture. Bilò et al. [STOC 2023] gave a solution with stretch  $3 + \varepsilon$ , query time  $O(n^\alpha)$  but space  $O(n^{2-\frac{\alpha}{f+1}})$ , approaching the quadratic barrier for large sensitivity. In the realm of subquadratic space, our  $f$ -DSOs are the first ones that guarantee, at the same time, large sensitivity, low stretch, and non-trivial query time. To obtain our results, we use the approximate distance oracles of Thorup and Zwick [JACM 2005], and the derandomization of the  $f$ -DSO of Weimann and Yuster [TALG 2013] that was recently given by Karthik and Parter [SODA 2021].

- [4] Bertschinger, N., Hofer, M., Krogmann, S., Lenzner, P., Schuldenzucker, S., Wilhelmi, L., [Equilibria and Convergence in Fire Sale Games](#). In: *Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 215–223, 2023.

The complex interactions between algorithmic trading agents can have a severe influence on the functioning of our economy, as witnessed by recent banking crises and trading anomalies. A common phenomenon in these situations are fire sales, a contagious process of asset sales that trigger further sales. We study the existence and structure of equilibria in a game-theoretic model of fire sales. We prove that for a wide parameter range (e.g., convex price impact functions), equilibria exist and form a complete lattice. This is contrasted with a non-existence result for concave price impact functions. Moreover, we study the convergence of best-response dynamics towards equilibria when they exist. In general, best-response dynamics may cycle. However, in many settings they are guaranteed to converge to the socially optimal equilibrium when starting from a natural initial state. Moreover, we discuss

a simplified variant of the dynamics that is less informationally demanding and converges to the same equilibria. We compare the dynamics in terms of convergence speed.

- [5] Krogmann, S., Lenzner, P., Skopalik, A., [Strategic Facility Location with Clients that Minimize Total Waiting Time](#). In: *Conference on Artificial Intelligence (AAAI)*. 5, pp. 5714–5721, 2023.

We study a non-cooperative two-sided facility location game in which facilities and clients behave strategically. This is in contrast to many other facility location games in which clients simply visit their closest facility. Facility agents select a location on a graph to open a facility to attract as much purchasing power as possible, while client agents choose which facilities to patronize by strategically distributing their purchasing power in order to minimize their total waiting time. Here, the waiting time of a facility depends on its received total purchasing power. We show that our client stage is an atomic splittable congestion game, which implies existence, uniqueness and efficient computation of a client equilibrium. Therefore, facility agents can efficiently predict client behavior and make strategic decisions accordingly. Despite that, we prove that subgame perfect equilibria do not exist in all instances of this game and that their existence is NP-hard to decide. On the positive side, we provide a simple and efficient algorithm to compute 3-approximate subgame perfect equilibria.

- [6] Bilò, D., Choudhary, K., Cohen, S., Friedrich, T., Krogmann, S., Schirneck, M., [Fault-Tolerant  \$ST\$ -Diameter Oracles](#). In: *International Colloquium on Automata, Languages and Programming (ICALP)*, pp. 24:1–24:20, 2023.

We study the problem of estimating the  $ST$ -diameter of a graph that is subject to a bounded number of edge failures. An  $f$ -edge fault-tolerant  $ST$ -diameter oracle ( $f$ -FDO- $ST$ ) is a data structure that preprocesses a given graph  $G$ , two sets of vertices  $S, T$ , and positive integer  $f$ . When queried with a set  $F$  of at most  $f$  edges, the oracle returns an estimate  $\hat{D}$  of the  $ST$ -diameter  $\text{diam}(G - F, S, T)$ , the maximum distance between vertices in  $S$  and  $T$  in  $G - F$ . The oracle has stretch  $\sigma \geq 1$  if  $\text{diam}(G - F, S, T) \leq \hat{D} \leq \sigma \text{diam}(G - F, S, T)$ . If  $S$  and  $T$  both contain all vertices, the data structure is called an  $f$ -edge fault-tolerant diameter oracle ( $f$ -FDO). An  $f$ -edge fault-tolerant distance sensitivity oracles ( $f$ -DSO) estimates the pairwise graph distances under up to  $f$  failures. We design new  $f$ -FDOs and  $f$ -FDO- $ST$ s by reducing their construction to that of all-pairs and single-source  $f$ -DSOs. We obtain several new tradeoffs between the size of the data structure, stretch guarantee, query and preprocessing times for diameter oracles by combining our black-box reductions with known results from the literature. We also provide an information-theoretic lower bound on the space requirement of approximate  $f$ -FDOs. We show that there exists a family of graphs for which any  $f$ -FDO with sensitivity  $f \geq 2$  and stretch less than  $5/3$  requires  $\Omega(n^{3/2})$  bits of space, regardless of the query time.

- [7] Gadea Harder, J., Krogmann, S., Lenzner, P., Skopalik, A., [Strategic Resource Selection with Homophilic Agents](#). In: *International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 2701–2709, 2023.

The strategic selection of resources by selfish agents is a classic research direction, with Resource Selection Games and Congestion Games as prominent examples. In these games, agents select available resources and their utility then depends on the number of agents using the same resources. This implies that there is no distinction between the agents, i.e., they are anonymous. We depart from this very general setting by proposing Resource Selection Games with heterogeneous agents that strive for joint resource usage with similar agents. So, instead of the number of other users of a given resource, our model considers agents with different types and the decisive feature is the fraction of same-type agents among the users. More precisely, similarly to Schelling Games, there is a tolerance threshold  $\tau \in [0, 1]$  which specifies the agents’ desired minimum fraction of same-type agents on a resource. Agents strive to select resources where at least a  $\tau$ -fraction of those resources’ users have the same type as themselves. For  $\tau = 1$ , our model generalizes Hedonic Diversity Games with a peak at 1. For our general model, we consider the existence and quality of equilibria and the complexity of maximizing social welfare. Additionally, we consider a bounded rationality model, where agents can only estimate the utility of a resource, since they only know the fraction of same-type agents on a given resource, but not the exact numbers. Thus, they cannot know the impact a strategy change would have on a target resource. Interestingly, we show that this type of bounded rationality yields favorable game-theoretic properties and specific equilibria closely approximate equilibria of the full knowledge setting.

- [8] Khomutovskiy, I., Dunker, R., Dierking, J., Egbert, J., Helms, C., Schöllkopf, F., Casel, K., Fischbeck, P., Friedrich, T., Isaac, D., Krogmann, S., Lenzner, P., [Applying Skeletons to Speed Up the Arc-Flags Routing Algorithm](#). In: *SIAM Symposium on Algorithm Engineering and Experiments (ALENEX)*, pp. 110–122, 2023.

The Single-Source Shortest Path problem is classically solved by applying Dijkstra’s algorithm. However, the plain version of this algorithm is far too slow for real-world applications such as routing in large road networks. To amend this, many speed-up techniques have been developed that build on the idea of computing auxiliary data in a preprocessing phase, that is used to speed up the queries. One well-known example is the Arc-Flags algorithm that is based on the idea of precomputing edge flags to make the search more goal-directed. To explain the strong practical performance of such speed-up techniques, several graph parameters have been introduced. The skeleton dimension is one such parameter that has already been used to derive runtime bounds for some speed-up techniques. Moreover, it was experimentally shown to be low in real-world road networks. We introduce a method to incorporate skeletons, the underlying structure behind the skeleton dimension, to improve routing speed-up techniques even further. As a proof of concept, we develop new algorithms called SKARF and SKARF+ that combine skeletons with Arc-Flags, and demonstrate via extensive experiments on large real-world road networks that SKARF+ yields a significant reduction of the search space and the query time of about 30% to 40% over Arc-Flags. We also prove theoretical bounds on the query time of SKARF, which is the first time an Arc-Flags variant has been analyzed in terms of skeleton dimension.

- [9] Bilò, D., Chechik, S., Choudhary, K., Cohen, S., Friedrich, T., Krogmann, S., Schirneck, M., [Approximate Distance Sensitivity Oracles in Subquadratic Space](#). In: *Symposium on Theory of Computing (STOC)*, pp. 1396–1409, 2023.

An  $f$ -edge fault-tolerant distance sensitive oracle ( $f$ -DSO) with stretch  $\sigma \geq 1$  is a data structure that preprocesses a given undirected, unweighted graph  $G$  with  $n$  vertices and  $m$  edges, and a positive integer  $f$ . When queried with a pair of vertices  $s, t$  and a set  $F$  of at most  $f$  edges, it returns a  $\sigma$ -approximation of the  $s$ - $t$ -distance in  $G - F$ . We study  $f$ -DSOs that take subquadratic space. Thorup and Zwick [JACM 2015] showed that this is only possible for  $\sigma \geq 3$ . We present, for any constant  $f \geq 1$  and  $\alpha \in (0, \frac{1}{2})$ , and any  $\varepsilon > 0$ , an  $f$ -DSO with stretch  $3 + \varepsilon$  that takes  $\tilde{O}(n^{2 - \frac{\alpha}{f+1}} / \varepsilon) \cdot O(\log n / \varepsilon)^{f+1}$  space and has an  $O(n^\alpha / \varepsilon^2)$  query time. We also give an improved construction for graphs with diameter at most  $D$ . For any constant  $k$ , we devise an  $f$ -DSO with stretch  $2k - 1$  that takes  $O(D^{f+o(1)} n^{1+1/k})$  space and has  $\tilde{O}(D^{o(1)})$  query time, with a preprocessing time of  $O(D^{f+o(1)} mn^{1/k})$ . Chechik, Cohen, Fiat, and Kaplan [SODA 2017] presented an  $f$ -DSO with stretch  $1 + \varepsilon$  and preprocessing time  $\tilde{O}_\varepsilon(n^5)$ , albeit with a super-quadratic space requirement. We show how to reduce their preprocessing time to  $O(mn^2) \cdot O(\log n / \varepsilon)^f$ .

- [10] Krogmann, S., Lenzner, P., Molitor, L., Skopalik, A., [Two-Stage Facility Location Games with Strategic Clients and Facilities](#). In: *International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 292–298, 2021.

We consider non-cooperative facility location games where both facilities and clients act strategically and heavily influence each other. This contrasts established game-theoretic facility location models with non-strategic clients that simply select the closest opened facility. In our model, every facility location has a set of attracted clients and each client has a set of shopping locations and a weight that corresponds to her spending capacity. Facility agents selfishly select a location for opening their facility to maximize the attracted total spending capacity, whereas clients strategically decide how to distribute their spending capacity among the opened facilities in their shopping range. We focus on a natural client behavior similar to classical load balancing: our selfish clients aim for a distribution that minimizes their maximum waiting times for getting serviced, where a facility’s waiting time corresponds to its total attracted client weight. We show that subgame perfect equilibria exist and give almost tight constant bounds on the Price of Anarchy and the Price of Stability, which even hold for a broader class of games with arbitrary client behavior. Since facilities and clients influence each other, it is crucial for the facilities to anticipate the selfish clients’ behavior when selecting their location. For this, we provide an efficient algorithm that also implies an efficient check for equilibrium. Finally, we show that computing a socially optimal facility placement is NP-hard and that this result holds for all feasible client weight distributions.

- [11] Arndt, T., Hafner, D., Kellermeier, T., Krogmann, S., Razmjou, A., Krejca, M. S., Rothenberger, R., Friedrich, T., [Probabilistic Routing for On-Street Parking Search](#). In: *European Symposium on Algorithms (ESA)*, pp. 6:1–6:13, 2016.

An estimated 30% of urban traffic is caused by search for parking spots. Traffic could be reduced by suggesting effective routes leading along potential parking spots. In this paper, we formalize parking search as a probabilistic problem on a road graph and show that it is NP-complete. We explore heuristics that optimize for the driving duration and the walking distance to the destination. Routes are constrained to reach a certain probability threshold of finding a spot. Empirically estimated probabilities of successful parking attempts are provided by TomTom on a per-street basis. We release these probabilities as a dataset of about 80,000 roads covering the Berlin area. This allows to evaluate parking search algorithms on a real road network with realistic probabilities for the first time. However, for many other areas, parking probabilities are not openly available. Because they are effortful to collect, we propose an algorithm that relies on conventional road attributes only. Our experiments show that this algorithm comes close to the baseline by a factor of 1.3 in our cost measure. This leads to the conclusion that conventional road attributes may be sufficient to compute reasonably good parking search routes.