

Perturbation resilience for the facility location problem

Bodo Manthey

Matthijs B. Tijkink

University of Twente
Department of Applied Mathematics
P. O. Box 217, 7500 AE Enschede, The Netherlands
`b.manthey@utwente.nl`

We analyze a simple local search heuristic for the facility location problem using the notion of perturbation resilience: an instance is γ -perturbation resilient if all costs can be perturbed by a factor of γ without changing the optimal solution.

We prove that local search for FLP succeeds in finding the optimal solution for γ -perturbation resilient instances for $\gamma \geq 3$, and we show that this is tight.

Keywords: facility location problem, perturbation resilience, local search

1 Introduction

The uncapacitated facility location problem (FLP) is a classical optimization problem: given a set of customers and a set of locations, the goal is to open facilities at a subset of the locations in order to serve all customers in the cheapest possible way. The FLP is NP-hard and even hard to approximate [7], but in practice, it seems to be much less difficult to solve than indicated by its worst-case complexity [8].

Of particular importance to solve such problems in practice are local search heuristics: they are often simple and easy to implement, yet show a remarkable performance in practice, although their worst-case behavior is usually poor. In order to get a better understanding of the performance of local search heuristics for the FLP, we apply the notion of perturbation resilience [2] to a simple local search strategy for the FLP.

1.1 Facility location problem

In this paper, we consider the uncapacitated, metric facility location problem (FLP):

Instances: a finite set F of locations, a finite set D of customers, facility costs $f_i \geq 0$ for all $i \in F$, and service costs $c_{ij} \geq 0$ for all $i \in F$ and $j \in D$. The service costs are metric, i.e., $c_{ij} \leq c_{i'j} + c_{i'j'} + c_{ij'}$ for all $i, i' \in F, j, j' \in D$.

Solutions: a solution (X, σ) for an instance (F, D, f, c) of the FLP consists of a nonempty set of open facilities $X \subseteq F$ and a customer assignment $\sigma : D \rightarrow X$ to open facilities.

Objective function and goal: The cost of a solution is $c(X, \sigma) = \sum_{i \in X} f_i + \sum_{j \in D} c_{\sigma(j)j}$. The goal is to minimize $c(X, \sigma)$.

This variant is called uncapacitated since all facilities can handle an arbitrary number of customers and it is called metric because the service costs satisfy an extension of the triangle inequality. For the discussion of other variants, we refer to two surveys [11, 10].

An optimal solution of an FLP instance is denoted as (X^*, σ^*) . The costs of a solution (X, σ) can be split into two parts: the facility costs $c_F(X) = \sum_{i \in X} f_i$ and service costs $c_S(X, \sigma) = \sum_{j \in D} c_{\sigma(j)j}$. Additionally, given an instance of the FLP and a set of open facilities X , it is easy to compute an optimal corresponding customer assignment σ : $\sigma(j) = \operatorname{argmin}_{i \in X} c_{ij}$, breaking ties arbitrarily. Thus, the customer assignment is often dropped in the cost notations, which implies that an optimal assignment is used.

The FLP cannot be approximated in polynomial time with a factor smaller than 1.463 unless $\text{NP} \subseteq \text{DTime}(n^{O(\log \log n)})$ [7]. On the other hand, there are approximation algorithms that guarantee a solution within a factor 1.5 of the optimal solution [6], and there is a randomized 1.488 approximation algorithm [9].

Local search heuristics are a popular way to solve optimization problems. They often show very poor performance in theoretical studies, but are often very powerful in practice.

We consider the following simple local search heuristic for the FLP: starting with an initial solution, we can obtain a new solution Y from our current solution X by

- removing one facility of X , or
- adding one new facility to X , or
- replacing one facility in X by a new one.

A solution is a local minimum with respect to this neighborhood if we cannot find a cheaper solution by applying any of these rules once. Since we are only concerned with the existence of non-global, local minima in this paper, a specification how the initial solution is computed or which step is chosen in case of multiple possibilities is not needed.

1.2 Perturbation resilience

Attacking the discrepancy between theoretical intractability and practical solvability has gained significant attention in the last couple of years. By either applying probabilistic input models or identifying properties that seem to make problems easier, one tries to find rigorous explanations why certain problems are easy in practice yet difficult in the worst case.

In particular in the area of clustering problems, properties have been identified that – if satisfied by the instance at hand – allow for efficient algorithms to solve or approximate the clustering problem [2, 3, 5, 4]. Given the close relationship between clustering on the one hand and FLP on the other hand, it seems natural to transfer these notions also to the FLP.

In this paper, we apply the notion of perturbation resilience to the FLP. Perturbation resilience, introduced and applied to clustering by Awasthi et al. [2], is a degeneracy condition: by how much can the costs be perturbed without changing the optimal solution? Intuitively, the more the instance can be perturbed without changing the optimal solution, the easier it should be to find the optimal solution.

Definition 1. An instance (F, D, f', c') of the non-metric FLP is a γ -perturbation of instance (F, D, f, c) of the FLP, with $\gamma \geq 1$, if and only if $f_i \leq f'_i \leq \gamma f_i$ for all $i \in F$ and $c_{ij} \leq c'_{ij} \leq \gamma c_{ij}$ for all $i \in F, j \in D$. If it is clear from the context which F and D are used, a γ -perturbed instance can also be denoted as (f', c') .

A γ -perturbed instance is any instance that is a γ -perturbation of some fixed FLP instance.

Note that γ -perturbed instances are not necessarily metric. This definition allows only to increase the costs, but this is only a matter of scaling. Using the definition of γ -perturbed instances, we can define γ -perturbation resilience for the FLP.

Definition 2. An instance (F, D, f, c) of the FLP is γ -perturbation resilient with $\gamma \geq 1$ if and only if all γ -perturbations (f', c') of (F, D, f, c) have the same unique optimal solution (X^*, σ^*) .

If costs of γ -perturbed instances are compared, the notation of the perturbed costs follow from the names given to the perturbed instances. So, for example, if (f', c') is a γ -perturbed instance, then $c'(X)$ denotes the cost of solution X in the perturbed instance.

Any γ -perturbation resilient instance of the FLP is a valid γ -perturbation of itself, so γ -perturbation resilience implies that the original instance has the same optimal solution as any of its perturbations. If an instance has multiple optimal solutions, it is not γ -perturbation resilient for any γ . Furthermore, if an instance is γ -perturbation resilient, then it is also γ' -perturbation resilient for any $\gamma' < \gamma$.

1.3 Our contribution

In this paper, we prove that local search for FLP finds the optimal solution for all 3-perturbation resilient instances. Furthermore, we show that this bound is tight. We admit that 3-perturbation resilience is a quite strong assumption. Nevertheless, this is a first step towards applying recent techniques from the analysis of algorithms to the FLP in order to better understand its practical tractability.

2 Local minima are global minima for $\gamma \geq 3$

In this section, we show that, for $\gamma \geq 3$, γ -perturbation resilient instances do not have local minima except for the unique global minimum. In the next section, we show that this is tight.

The following theorem is a result for local search algorithms on any FLP instance.

Theorem 3 (Arya et al. 2004 [1]). For a given instance of the FLP, let X be a local minimum, and let X^* be an optimal solution. Then $c_F(X) \leq c_F(X^*) + 2c_S(X^*)$ and $c_S(X) \leq c_F(X^*) + c_S(X^*)$. Combined, this yields $c(X) \leq 3c(X^*)$.

This theorem will be used to show that there are no local minima except for the optimal solution for 3-perturbation resilient FLP instances (Theorem 5). To prove this, we first the following lemma, which states the following: if there exists some γ -perturbation resilient instance with a non-global local optimum, then we can remove all facilities that belong neither to the non-global local optimum nor to the global optimum. This newly constructed instance is still γ -perturbation resilient and both the local and global optimum are maintained.

Lemma 4. Assume a γ -perturbation resilient instance (F, D, f, c) of the FLP exists with a local minimum (X, σ) not equal to the optimal solution (X^*, σ^*) .

Then an instance (F', D', f', c') of the FLP exists with the following properties:

- the instance (F', D', f', c') is γ -perturbation resilient, and
- the instance (F', D', f', c') has a local minimum (X', σ') not equal to the optimal solution (X'^*, σ'^*) , and
- $F' = X' \cup X'^*$, and
- $f'_i = 0$ for all $i \in X' \cap X'^*$, and
- for all $j \in D'$, we have $\sigma'(j) \notin X' \cap X'^*$ or $\sigma'^*(j) \notin X' \cap X'^*$.

Proof. Take any such instance (F, D, f, c) . We transform this instance to a new instance (F', D', f', c') using the following steps. We apply each step until its conditions are satisfied for all local minima $X \neq X^*$, possibly applying earlier steps again in the process if their conditions are not valid any more after applying a later step. Thus, at the beginning of every step, the conditions of all previous steps are satisfied.

Step 1: Condition to satisfy: $F = X \cup X^*$. Drop all facilities not in X or X^* , i.e., $F' = X \cup X^*$. All customer assignments of σ and σ^* remain valid and the costs $c(X)$ and $c(X^*)$ are unchanged. Thus, the resulting instance (F', D, f, c) is still γ -perturbation resilient with optimal solution X^* . Because all subsets of F' also are a subset of F , $X \neq X^*$ is still a local minimum. After doing this, the conditions of step 1 are satisfied.

Step 2: Condition to satisfy: $f_i = 0$ for all $i \in X \cap X^*$. Change the facility costs to the following:

$$\tilde{f}_i = \begin{cases} 0 & \text{if } i \in X \cap X^*, \text{ and} \\ f_i & \text{otherwise.} \end{cases}$$

Note that the resulting instance (F, D, \tilde{f}, c) still has X as a local minimum, since the cost of adding a facility to X is identical, compared to instance (F, D, f, c) , and the cost of dropping or swapping a facility from X is the equal or higher, as compared to instance (F, D, f, c) . To show that the instance (F, D, \tilde{f}, c) is γ -perturbation resilient with optimal solution X^* , consider all nonempty sets of open facilities $Y \subseteq F$ and customer assignments σ' in all perturbations of costs f'_i and c'_{ij} and equivalent perturbations of \tilde{f}_i and c_{ij} :

$$\begin{aligned} c'(X^*, \sigma^*) &= c'_F(X^*) + c'_S(X^*, \sigma^*) \\ &= \sum_{i \in X^* \setminus X} f'_i + \sum_{i \in X \cap X^*} f'_i + c'_S(X^*, \sigma^*) = \tilde{c}'(X^*, \sigma^*) + \sum_{i \in X \cap X^*} f'_i, \\ c'(Y, \sigma') &= c'_F(Y) + c'_S(Y, \sigma') \\ &= \sum_{i \in Y \setminus (X \cap X^*)} f'_i + \sum_{i \in Y \cap X \cap X^*} f'_i + c'_S(Y, \sigma') \\ &= \tilde{c}'(Y, \sigma') + \sum_{i \in Y \cap X \cap X^*} f'_i. \end{aligned}$$

This implies that

$$\begin{aligned} \tilde{c}'(X^*, \sigma^*) + \sum_{i \in X \cap X^*} f'_i &\leq \tilde{c}'(Y, \sigma') + \sum_{i \in Y \cap X \cap X^*} f'_i, \\ \tilde{c}'(X^*, \sigma^*) + \sum_{i \in (X \cap X^*) \setminus Y} f'_i &\leq \tilde{c}'(Y, \sigma'), \text{ and} \\ \tilde{c}'(X^*, \sigma^*) &\leq \tilde{c}'(Y, \sigma'). \end{aligned}$$

So the instance (F, D, \tilde{f}, c) is indeed still γ -perturbation resilient. This satisfies the conditions for step 2.

Step 3. Condition to satisfy: for all $j \in D$, $(\sigma(j) \notin X \cap X^*$ or $\sigma^*(j) \notin X \cap X^*)$ must be true. Choose an arbitrary $j \in D$ with $\sigma(j) \in X \cap X^*$ and $\sigma^*(j) \in X \cap X^*$. By the definition of γ -perturbation resilience, j cannot be assigned to any other facility in X^* in all of the γ -perturbed costs c' . Thus $c_{\sigma^*(j)j} < \gamma c_{ij}$ for all $i \in X^* \setminus \{\sigma^*(j)\}$. Also, because $F = X \cap X^*$ and X is a local minimum, $c_{\sigma(j)j} \leq c_{ij}$ for $i \in F \setminus X^* = X \setminus X^*$. Thus, the assignment $\sigma(j) = \sigma^*(j)$ is the best possible assignment in F .

Let \tilde{c} denote the costs in instance (F, D', f, c) .

The new instance (F, D', f, c) is created by removing customer j , i.e., $D' = D \setminus \{j\}$. As a result, X still is a local minimum in (F, D', f, c) :

- Dropping a facility $i \in X$ (if $|X| > 2$); If $i = \sigma(j)$, then $f_i = 0$ because $\sigma(j) \in X \cap X^*$. Thus, $\tilde{c}(X \setminus \{i\}) \geq \tilde{c}(X)$. If $i \neq \sigma(j)$, then customer j is not connected to facility i , so $\tilde{c}(X) - \tilde{c}(X \setminus \{i\}) = c(X) - c(X \setminus \{i\}) \geq 0$. Hence, dropping facility $i \in X$ does not result in a better solution.
- Adding a facility $i \in F \setminus X$; Removing customer j does not change the cost of adding a facility, since $\sigma^*(j) = \sigma(j)$ and $c_{\sigma(j)j} \leq c_{ij}$: $\tilde{c}(X) - \tilde{c}(X \cup \{i\}) = c(X) - c(X \cup \{i\}) \geq 0$. Thus, adding facility $i \in F \setminus X$ does not result in a better solution.
- Swapping an open facility $i \in X$ with closed facility $i' \in F \setminus X$; If $i \neq \sigma(j)$, the same reasoning as in adding a facility holds. If $i = \sigma(j)$, $f_i = 0$ and this swap is not better than just adding facility i' , which did not improve the cost either. Thus, swapping open facility $i \in X$ with closed facility $i' \in F \setminus X$ does not result in a better solution.

To show that the new instance is γ -perturbation resilient with optimal solution (X^*, σ^*) , consider all $Y \subseteq F$ and γ -perturbations (f', c') . Let $Y' = Y \cup (X \cap X^*)$ with optimal assignment (i.e., $\sigma'(x) = \operatorname{argmin}_{i \in Y'} c'_{ix}$) and note that $c'(Y') \leq c'(Y)$ since $f'_i = 0$ for all $i \in X \cap X^*$, even in the new instance (F, D, f', c') . Let \tilde{c}' denote the costs in instance (F, D', f', c') . Thus:

$$\begin{aligned} c'(X^*, \sigma^*) &= c'_F(X^*) + c'_S(X^*, \sigma^*) \\ &= c'_F(X^*) + c'_{\sigma^*(j)j} + \sum_{x \in D'} c'_{\sigma^*(x)x} = \tilde{c}'(X^*) + c'_{\sigma^*(j)j}, \end{aligned} \quad (1)$$

$$\begin{aligned} c'(Y', \sigma') &= c'_F(Y') + c'_S(Y', \sigma') \\ &= c'_F(Y') + c'_{\sigma'(j)j} + \sum_{x \in D'} c'_{\sigma'(x)x} = \tilde{c}'(Y', \sigma) + c'_{\sigma'(j)j}. \end{aligned} \quad (2)$$

Note that $c'_{\sigma'(j)j} \leq c'_{\sigma^*(j)j}$, because $\sigma^*(j) \in Y'$. This together with (1) and (2) imply

$$\begin{aligned} \tilde{c}'(X^*, \sigma^*) + c'_{\sigma^*(j)j} &\leq \tilde{c}'(Y', \sigma') + c'_{\sigma'(j)j} \text{ and} \\ \tilde{c}'(X^*, \sigma^*) &\leq \tilde{c}'(Y', \sigma') \leq \tilde{c}'(Y), \end{aligned}$$

so even after removal of customer j , the instance (F, D', f, c) is γ -perturbation resilient with optimal solution X^* . After doing this a couple of times, the condition for step 3 is satisfied.

After step 3, all conditions required for the lemma are satisfied. Note that all steps make the instance smaller in some way (less facilities, less facilities with nonzero cost, less customers), so this process terminates eventually. \square

This lemma can be interpreted as removing complications from the instance, except those which are necessary for either γ -perturbation resilience or the existence of the local minimum X . The following theorem uses Lemma 4 to show that any local minimum of 3-perturbation resilient FLP instances always is the global minimum.

Theorem 5. *All local minima (X, σ) of a γ -perturbation resilient instance (F, D, f, c) of the FLP with $\gamma \geq 3$ are equal to the optimal solution (X^*, σ^*) of the instance.*

Proof. We assume the contrary and use Lemma 4 to get an instance (F, D, f, c) with $F = X \cup X^*$, $f_i = 0$ for $i \in X \cap X^*$ and, for all $j \in D$, $(\sigma(j) \notin X \cap X^*$ or $\sigma^*(j) \notin X \cap X^*)$. Here (X^*, σ^*) is the optimal solution and $(X, \sigma) \neq (X^*, \sigma^*)$ the local minimum.

We perturb the costs as follows:

$$f'_i = \begin{cases} 3c_{ij} & \text{if } i \in X \cap X^*, \text{ and} \\ 3f_i & \text{if } i \in X^* \setminus X, \text{ and} \\ f_i & \text{otherwise, and} \end{cases}$$

$$c'_{ij} = \begin{cases} 3c_{ij} & \text{if } \sigma^*(j) = i, \text{ and} \\ c_{ij} & \text{otherwise.} \end{cases}$$

This is a valid γ -perturbation. Because the instance (F, D, f, c) is γ -perturbation resilient, the following holds:

$$\begin{aligned} c'(X^*, \sigma^*) &= c'_F(X^*) + c'_S(X^*, \sigma^*) \\ &= \sum_{i \in X^* \setminus X} 3f_i + \sum_{j \in D} 3c_{\sigma^*(j)j} \quad (\text{by choice of } f'_i \text{ and } c'_{ij}) \\ &= 3c(X^*, \sigma^*), \\ c'(X, \sigma) &= c_F(X) + c_S(X, \sigma) \\ &= \sum_{i \in X \setminus X^*} f_i + \sum_{j \in D} c_{\sigma(j)j}, \quad (\text{by the properties of } (F, D, f, c), f'_i \text{ and } c'_{ij}) \\ &= c(X, \sigma), \end{aligned}$$

so $3c(X^*, \sigma^*) < c(X, \sigma)$.

By Theorem 3, $c(X, \sigma) \leq 3c(X^*, \sigma^*)$, so $3c(X^*) < 3c(X^*)$ which is a contradiction. Thus, for γ -perturbation resilient instances of the FLP with $\gamma \geq 3$, no local minima exist except for the global optimum. \square

3 Non-global local minima for $\gamma < 3$

Now we show that Theorem 5 is tight: For every $\varepsilon > 0$, there is a $(3 - \varepsilon)$ -perturbation resilient instance that possesses a local optimum that is no global optimum.

Theorem 6. *There exist γ -perturbation resilient instances of the FLP for all $\gamma < 3$ with local minima $X \neq X^*$.*

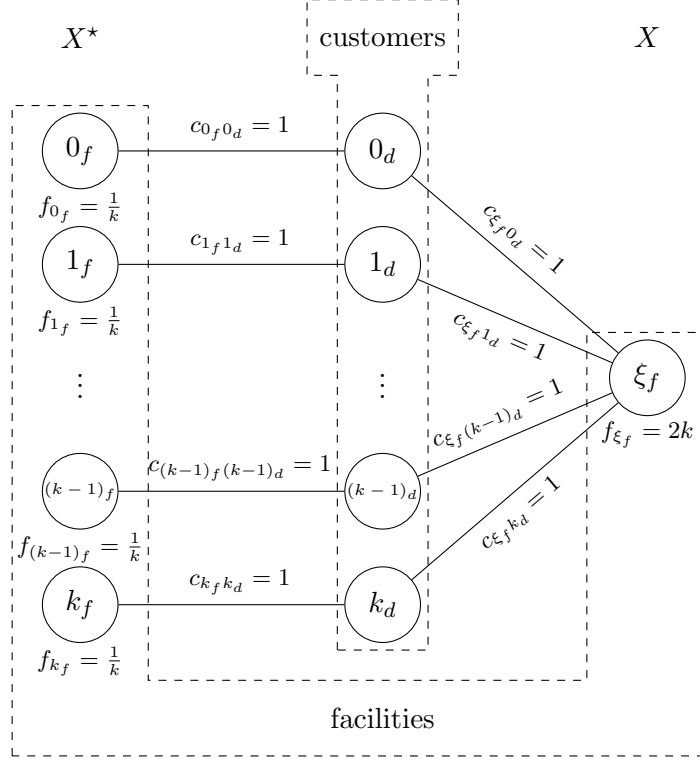


Figure 1: Example used in Theorem 6 (indirect service costs not shown)

Proof. The following example, inspired by Arya et al. [1], proves the theorem (see also Figure 1). Choose a $k \in \mathbb{N}$, whose value will be specified later on. Let

$$\begin{aligned}
 F &= \{0_f, 1_f, \dots, (k-1)_f, k_f, \xi_f\}, \\
 D &= \{0_d, \dots, k_d\}, \\
 f_i &= \begin{cases} 2k & \text{if } i = \xi_f, \text{ and} \\ \frac{1}{k} & \text{otherwise, and} \end{cases} \\
 c_{ij} &= \begin{cases} 1 & \text{if } i = \xi_f, \text{ and} \\ 1 & \text{if } i = k_f, j = k_d \text{ for some } k, \text{ and} \\ 3 & \text{otherwise.} \end{cases}
 \end{aligned}$$

In the optimal solution $X^* = \{0_f, \dots, k_f\}$, all $k+1$ facilities of costs $1/k$ are opened. Together with the service costs of 1 per client, this yields $c(X^*, \sigma^*) = (k+1)(1 + \frac{1}{k})$. This instance is $\frac{3k}{(k+1)(1+1/k)}$ -perturbation resilient, as we will show by comparing all solutions $X \neq X^*$ with X^* for all γ -perturbations (F, D, f', c') . By letting k go to infinity, the perturbation resilience gets arbitrarily close to 3.

We consider two cases: $\xi_f \in X$ and $\xi_f \notin X$. For the first case, assume $\xi_f \in X$. If this is the case, then $c(X, \sigma) \geq 3k+1$. Thus, $c'(X^*, \sigma^*) \leq 3k < 3k+1 \leq c'(X, \sigma)$, which completes the first case.

The second case is when $\xi_f \notin X$. Without loss of generality, let $X = \{0_f, \dots, (|X|-1)_f\}$. Since $c'_{i_f i_d} < 3 \leq c'_{i'_d i_d}$ for $i \in \{0, \dots, |X|-1\}, i' \in F \setminus \{i_f\}$, all customers $\{0_d, \dots, (|X|-1)_d\}$

are assigned to the same facility in both X^* and X , i.e., $\sigma^*(i_d) = \sigma(i_d) = i_f$. For the other customers j_d with $j \in \{|X|, \dots, k\}$, we look at the following quantity:

$$\begin{aligned} -f'_{j_f} + c'_{\sigma(j_d)j_d} - c'_{\sigma^*(j_d)j_d} &\geq -\frac{3}{(k+1)(1+\frac{1}{k})} + 3 - \frac{3k}{(k+1)(1+\frac{1}{k})} \\ &= \frac{3}{k+1} > 0. \end{aligned}$$

Using this quantity, we can bound the difference between $c'(X^*, \sigma^*)$ and $c'(X, \sigma)$:

$$c'(X, \sigma) - c'(X^*, \sigma^*) = \sum_{j=|X|}^k c'_{\sigma(j_d)j_d} - c'_{\sigma^*(j_d)j_d} - f'_{j_f} > 0,$$

so $c'(X^*, \sigma^*) < c'(X, \sigma)$, completing the second case.

For both cases it holds that $c'(X^*, \sigma^*) < c'(X, \sigma)$, for all solutions $X \neq X^*$, so instance (F, D, f, c) is $\frac{3k}{(k+1)(1+1/k)}$ -perturbation resilient. Also, $X = \{\xi_f\} \neq X^*$ is a local minimum, completing the proof. \square

References

- [1] Vijay Arya, Naveen Garg, Rohit Khandekar, Adam Meyerson, Kamesh Munagala, and Vinayaka Pandit. Local search heuristics for k -median and facility location problems. *SIAM Journal on Computing*, 33(3):544–562, 2004.
- [2] Pranjal Awasthi, Avrim Blum, and Or Sheffet. Center-based clustering under perturbation stability. *Information Processing Letters*, 112(1-2):49–54, 2012.
- [3] Maria-Florina Balcan, Avrim Blum, and Anupam Gupta. Clustering under approximation stability. *Journal of the ACM*, 60(2), 2013.
- [4] Yonatan Bilu, Amit Daniely, Nathan Linial, and Michael E. Saks. On the practically interesting instances of MAXCUT. In Natacha Portier and Thomas Wilke, editors, *Proc. of the 30th Int. Symp. on Theoretical Aspects of Computer Science (STACS)*, volume 20 of *LIPICs*, pages 526–537. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2013.
- [5] Yonatan Bilu and Nathan Linial. Are stable instances easy? *Combinatorics, Probability & Computing*, 21(5):643–660, 2012.
- [6] Jaroslaw Byrka and Karen Aardal. An optimal bifactor approximation algorithm for the metric uncapacitated facility location problem. *SIAM Journal on Computing*, 39(6):2212–2231, 2010.
- [7] Sudipto Guha and Samir Khuller. Greedy strikes back: Improved facility location algorithms. *Journal of Algorithms*, 31(1):228–248, 1999.
- [8] Kamal Jain, Mohammad Mahdian, Evangelos Markakis, Amin Saberi, and Vijay V. Vazirani. Greedy facility location algorithms analyzed using dual fitting with factor-revealing LP. *Journal of the ACM*, 50(6):795–824, 2003.

- [9] Shi Li. A 1.488 approximation algorithm for the uncapacitated facility location problem. *Information and Computation*, 222:45–58, 2013.
- [10] M. Teresa Melo, Stefan Nickel, and Francisco Saldanha-da Gama. Facility location and supply chain management – a review. *European Journal of Operational Research*, 196(2):401–412, 2009.
- [11] Jens Vygen. Approximation algorithms for facility location problems. Lecture Notes, Forschungsinstitut für Diskrete Mathematik, Rheinische Friedrich-Wilhelms-Universität Bonn, 2005.