

Simpler Analyses of Local Search Algorithms for Facility Location*

Anupam Gupta

Kanat Tangwongsan

Abstract

We study local search algorithms for metric instances of facility location problems: the uncapacitated facility location problem (UFL), as well as uncapacitated versions of the k -median, k -center and k -means problems. All these problems admit natural local search heuristics: for example, in the UFL problem the natural moves are to open a new facility, close an existing facility, and to swap a closed facility for an open one; in k -medians, we are allowed only swap moves. The local-search algorithm for k -median was analyzed by Arya et al. (*SIAM J. Comput.* 33(3):544-562, 2004), who used a clever “coupling” argument to show that local optima had cost at most constant times the global optimum. They also used this argument to show that the local search algorithm for UFL was 3-approximation; their techniques have since been applied to other facility location problems.

In this paper, we give a proof of the k -median result which avoids this coupling argument. These arguments can be used in other settings where the Arya et al. arguments have been used. We also show that for the problem of opening k facilities F to minimize the objective function $\Phi_p(F) = (\sum_{j \in V} d(j, F)^p)^{1/p}$, the natural swap-based local-search algorithm is a $\Theta(p)$ -approximation. This implies constant-factor approximations for k -medians (when $p = 1$), and k -means (when $p = 2$), and an $O(\log n)$ -approximation algorithm for the k -center problem (which is essentially $p = \log n$).

1 Introduction

Facility location problems have been central objects of study in the operations research and computer science community, not only for their intrinsic fundamental nature and broad applicability, but also as problems whose solutions have led to the development of new ideas and techniques. Indeed, techniques such as rounding linear relaxations, using primal-dual techniques and Lagrangean relaxations, greedy algorithms (with and without the dual-fitting approach), and local search algorithms have all been honed when applied to facility location problems.

Local search has been a popular algorithm design paradigm, and has had many successes in the design of approximation algorithms for hard combinatorial optimization problems. The focus of this paper is on local search algorithms for facility location problems on metric spaces. In fact, the best approximation algorithm known for the k -median problem is a $(3 + \varepsilon)$ -approximation via local search [AGK⁺04]. However, the analysis of this simple local search algorithm is fairly subtle, and requires a careful “coupling” argument. The same coupling argument was used by [AGK⁺04] to also analyze a local search algorithm for uncapacitated facility location, and subsequently by [DGK⁺05, Pan04] for some other problems.

In this paper, we present the following:

*Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213. Research was partly supported by the NSF awards CCF-0448095 and CCF-0729022, and an Alfred P. Sloan Fellowship.

- We give somewhat simpler analyses for the natural local search algorithms of [AGK⁺04, DGK⁺05, Pan04]: while the approximation guarantees remain the same, the proofs are arguably more intuitive than existing proofs.
- We show that the problem of opening k facilities F to minimize the objective function

$$\Phi_p(F) = \left(\sum_{j \in V} d(j, F)^p \right)^{1/p},$$

the natural swap-based local-search algorithm is a $\Theta(p)$ -approximation¹: this immediately implies constant factor approximations for k -medians (which corresponds to the case $k = 1$), k -means (the case $p = 2$), an $O(\log n)$ -approximation local search algorithm for the k -center problem (which is essentially the case $p = \log n$). To the best of our knowledge, the results for $p \neq 1$ give the first analyses of local search heuristics for these problems.

Technical Ideas. The main contribution of this work is the simplification of the proofs. As with most local search proofs, the previous papers considered a local optimum, and showed that since a carefully chosen set of local moves were non-improving, we could infer some relationship between our cost and the optimal cost. However, in the previous papers, this set of local moves have to be carefully defined by looking at how the clients served by each optimal facility were split up between facilities in the local optimum. Even bounding the change in cost due to each of these potential local moves is somewhat non-trivial.

In our paper, we define a set of local moves based only on distance information about which of the optimal facilities F^* are close to which of our facilities F . The intuition is simple: consider each of the optimal facilities in F^* , and look at the closest facility to it in F . If some facility $f \in F$ is the closest to only one facility $f^* \in F^*$, then we should try swapping f with f^* . However, if there is some facility $f \in F$ that is the closest facility to many facilities in F^* , then swapping f might be bad for our solution, and hence we do not want to close this facility in any potential move. Formalizing this natural intuition gives us the claimed simpler proofs for many problems.

1.1 A Note on the Rate of Convergence.

In this paper, we will only focus on the quality of the local optima, and not explicitly deal with the rate of convergence. This, however, is only for brevity: since all our arguments are based on averaging arguments, we can show that if at any point in the local search procedure, the current solution is very far from every local optimum, then there is a step that reduces the objective function by a large amount. Moreover, since we are dealing with approximation algorithms, we can stop when each local step improves the objective function by at most a factor of $(1 + \varepsilon)$, which would give us only a slightly worse approximation guarantee but would keep the running time polynomial in n and ε^{-1} . These details are fairly standard; e.g., see the discussion in [AGK⁺04].

1.2 Other Related Work

Facility location problems have had a long history; here, we mention only some of the results for these problems. We focus on metric instances of these problems: the non-metric cases are usually much harder [Hoc82].

¹When we say that our local search algorithm is a ρ -approximation, we mean that the cost of every local optimum is at most ρ times the optimum cost. In this paper, we also implicitly mean that one can find a solution of cost $(\rho + \varepsilon)OPT$ in time $\text{poly}(n, \varepsilon)$ —see Section 1.2 for details.

k -Median. The k -median problem seeks to find facilities F with $|F| = k$ to minimize $\sum_{j \in V} d(j, F)^2$. The first constant factor approximation for the k -median problem was given by Charikar et al. [CGTS02], which was subsequently improved by [CG05] and [AGK⁺04] to the current best factor of $3 + \varepsilon$. It is known that the natural LP relaxation for the problem has an integrality gap of 3, but the currently-known algorithm that achieves this does not run in polynomial time [ARS03]. The extension of k -median to the case when one can open at most k facilities, but also has to pay their facility opening cost was studied by [DGK⁺05], who gave a 5-approximation.

k -Means. The k -means problem minimizes $\sum_{j \in V} d(j, F)^2$, and is widely used for clustering in machine learning, especially when the point set is in the Euclidean space. For Euclidean instances, one can obtain $(1 + \varepsilon)$ -approximations in linear time, if one imagines k and ε to be constants: see [KSS04] and the references therein. The most commonly used algorithm in practice is Lloyd’s algorithm, which is a local-search procedure different from ours, and which is a special case of the EM algorithm [Llo82]. While there is no explicit mention of an approximation algorithm with provable guarantees for k -means (to the best of our knowledge), many of the constant-factor approximations for k -median can be extended to the k -means problem as well. The paper of Kanungo et al. [KMN⁺04] is closely related to ours: it analyzes the same local search algorithm we consider, and uses properties of k -means in Euclidean spaces to obtain a 9-approximation. Our results hold for general metrics, and can essentially be viewed as extensions of their results.

k -Center. Tight bounds for the k -center problem are known: there is a 2-approximation algorithm due to [Gon85, HS86], and this is tight unless $\mathbf{P} = \mathbf{NP}$.

UFL. For the uncapacitated metric facility location (UFL) problem, the first constant factor approximation was given by Shmoys et al. [STA97]; subsequent approximation algorithms and hardness results have been given by [STA97, Chu98, Svi02, CS03, Byr07, JV01, CG05, PT03, MMSV01, JMS02, MYZ02, KPR00, CG05, AGK⁺04, GK99]. It remains a tantalizing problem to close the gap between the best known approximation factor of 1.5 [Byr07], and the hardness result of 1.463 [GK99].

1.3 Organization of the Paper

The paper is organized as follows. In Section 2, we consider the metric k -Median problem and present an analysis showing a 5 approximation. Following that, we analyze a version of the local-search algorithm for k -Median that allows t simultaneous swaps, where we show a $(3 + 2/t)$ approximation. In Section 3, we consider a common generalization of the k -median, k -means, and k -center problems, called the ℓ_p -norm k -facility location problem, where we give an $O(p)$ approximation algorithm, and present an instance which is asymptotically tight. Finally, in Section 4, we present simpler proofs of two other known results: a 3 approximation for the uncapacitated facility location problem, and a 5 approximation for the k -uncapacitated facility location problem.

2 A Simpler Analysis of k -Median Local Search

In this section, we study the local search algorithm for the k -median problem, and show that any local optimum has a cost which is at most 5 times the cost of the global optimum. Recall that given a set F of at most k facilities, the k -median cost is

$$\text{kmed}(F) = \sum_{j \in V} d(j, F).$$

The local search algorithm we consider in this section is the simplest one: we start with any set of k

facilities. At each point in time, we try to find some facility in our current set of facilities, and swap it with some currently unopened facility, so that the cost of the resulting solution decreases. It is known that any local minimum is a 5-approximation to the global minimum [AGK⁺04], and that the bound of 5 is tight for these local-search dynamics. Here we give a simpler proof of this 5-approximation.

2.1 A Set of Test Swaps

To show that a local optimum is a good approximation, the standard approach is to consider a carefully chosen subset of potential swaps: if we are at a local optimum, each of these swaps must be non-improving. This gives us some information about the cost of the local optimum. To this end, consider the set F^* of facilities chosen by an optimum solution, and let F be the facilities at the local optimum. Without loss of generality, assume that $|F| = |F^*| = k$.




Figure 2.1: An example mapping $\eta: F^* \rightarrow F$ and a set of test swaps S .

Define a map $\eta: F^* \rightarrow F$ that maps each optimal facility f^* to a closest facility $\eta(f^*) \in F$: that is, $d(f^*, \eta(f^*)) \leq d(f^*, f)$ for all $f \in F$. Now define $R \subseteq F$ to be all the facilities that have *at most* 1 facility in F^* mapped to it by the map η . (In other words, if we create a directed bipartite graph by drawing an arc from f^* to $\eta(f^*)$, $R \subseteq F$ are those facilities whose in-degree is at most 1).

Finally, we define a set of k pairs $S = \{(r, f^*)\} \subseteq R \times F^*$ such that

- Each $f^* \in F^*$ appears in exactly one pair (r, f^*) .
- If $\eta^{-1}(r) = \{f^*\}$ then r appears only once in S as the tuple (r, f^*) .
- If $\eta^{-1}(r) = \emptyset$ then r appears at most in two tuples in S .

The procedure is simple: for each $r \in R$ with in-degree 1, construct the pair $(r, \eta^{-1}(r))$ —let the optimal facilities that are already matched off be denoted by F_1^* . The other facilities in R have in-degree 0: denote them by R_0 . A simple averaging argument shows that the unmatched optimal facilities $|F^* \setminus F_1^*| \leq 2|R_0|$. Now, arbitrarily create pairs by matching each node in R_0 to at most two pairs in $F^* \setminus F_1^*$ so that the above conditions are satisfied.

The following fact is immediate from the construction:

Fact 2.1 *For any tuple $(r, f^*) \in S$ and $\hat{f}^* \in F$ with $\hat{f}^* \neq f^*$, $\eta(\hat{f}^*) \neq r$.*

Intuition for the Pairing. To get some intuition for why the pairing S was chosen, consider the case when each facility in F is the closest to a unique facility in F^* , and far away from all other facilities in F^* —in this case, opening facility $f^* \in F^*$ and closing the matched facility in $f \in F$ can be handled by letting all clients attached to f be handled by f^* (or by other facilities in F). A problem case would be when a facility $f \in F$ is the closest to several facilities in F^* , since closing f and opening only one of these facilities in F^* might still cause us to pay too much—hence we never consider the gains due to closing such “popular” facilities, and instead only consider the swaps that involve facilities from the set of relatively “unpopular” facilities R .

2.2 Bounding the Cost of a Local Optimum


In this section, we use the fact that each of the swaps in set S constructed in Section 2.1 are non-improving to show that the local optimum has small cost.

Breaking ties arbitrarily, assume that $\varphi : V \rightarrow F$ and $\varphi^* : V \rightarrow F^*$ are functions mapping each client to some closest facility. For any client j , let $O_j = d(j, F^*) = d(j, \varphi^*(j))$ be the client j 's cost in the optimal solution, and $A_j = d(j, F) = d(j, \varphi(j))$ be its cost in the local optimum. Let $N^*(f^*) = \{j \mid \varphi^*(j) = f^*\}$ be the set of clients assigned to f^* in the optimal solution, and $N(f) = \{j \mid \varphi(j) = f\}$ be those assigned to f in the local optimum.

Lemma 2.2 *For each swap $(r, f^*) \in S$,*

$$\text{kmed}(F + f^* - r) - \text{kmed}(F) \leq \sum_{j \in N^*(f^*)} (O_j - A_j) + \sum_{j \in N(r)} 2O_j. \quad (2.1)$$

Proof: Consider the following candidate assignment of clients (which gives us an upper bound on the cost increase): map each client in $N^*(f^*)$ to f^* . For each client $j \in N(r) \setminus N^*(f^*)$, consider the figure below. Let the facility $\hat{f}^* = \varphi^*(j)$: assign j to $\hat{f} = \eta(\hat{f}^*)$, the closest facility in F to \hat{f}^* . Note that by Fact 2.1, $\hat{f} \neq r$, and this is a valid new assignment. All other clients in $V \setminus (N(r) \cup N^*(f^*))$ stay assigned as they were in φ .



Note that for any client $j \in N^*(f^*)$, the change in cost is exactly $O_j - A_j$: summing over all these clients gives us the first term in the expression (2.1).

For any client $j \in N(r) \setminus N^*(f^*)$, the change in cost is

$$d(j, \hat{f}) - d(j, r) \leq d(j, \hat{f}^*) + d(\hat{f}^*, \hat{f}) - d(j, r) \quad (2.2)$$

$$\leq d(j, \hat{f}^*) + d(\hat{f}^*, r) - d(j, r) \quad (2.3)$$

$$\leq d(j, \hat{f}^*) + d(j, \hat{f}^*) = 2O_j. \quad (2.4)$$

with (2.2) and (2.4) following by the triangle inequality, and (2.3) using the fact that \hat{f} is the closest vertex in F to \hat{f}^* . Summing up, the total change for all these clients is at most

$$\sum_{j \in N(r) \setminus N^*(f^*)} 2O_j \leq \sum_{j \in N(r)} 2O_j, \quad (2.5)$$

the inequality holding since we are adding in non-negative terms. This proves Lemma 2.2. \blacksquare

Note that summing (2.1) over all tuples in S , along with the fact that each $f^* \in F^*$ appears exactly once and each $r \in R \subseteq F$ appears at most twice gives us the simple proof of the following theorem.

Theorem 2.3 ([AGK⁺04]) *At a local minimum F , the cost $\text{kmed}(F) \leq 5 \cdot \text{kmed}(F^*)$.*

2.3 A Projection Lemma

Like the proof for k -Median, most proofs in the remaining of the paper also rely on a candidate assignment of clients to give an upper bound on the cost increase. In particular, these proofs often reassign a client j to $\eta(\varphi^*(j))$, the image of the facility $\varphi^*(j)$ that serves j in the optimal solution. The following lemma bounds the distance between client j and the facility $\eta(\varphi^*(j))$; its proof is exactly the same as in (2.2)-(2.4).

Lemma 2.4 (Projection Lemma) *For any client $j \in C$,*

$$d(j, \eta(\varphi^*(j))) \leq 2O_j + A_j.$$

2.4 The k -Median Local Search with Multiswaps

The paper of Arya et al. [AGK⁺04] also showed that local minima under t -swaps for the k -median algorithm have cost at most $(3 + \frac{2}{t})$ times the optimum, where a t -swap involves shutting down t of the current set of k facilities and opening t new facilities in their stead. (Note that the best t -swap can be found in time $n^{O(t)}$, and hence is polynomial-time for constant t .) Using the notation from the previous section, we will show that $\text{kmed}(F) \leq (3 + \frac{2}{t})\text{kmed}(F^*)$, where F denotes the set of facilities at the local optimum, and F^* the set of facilities in an optimal solution.

2.4.1 A Set of Test t -Swaps

In order to extend the proof of the previous section to the t -swap algorithm, we generalize the set of test swaps as follows. The set of test t -swaps relies on the same intuition as before, but it is more sophisticated for technical reasons. For each element r of F , define the degree of r to be $\text{deg}(r) = |\eta^{-1}(r)|$. The pairing is constructed by the following procedure.

Algorithm 1 Pairing construction

$i = 0$

while there exists $r \in F$ such that $\text{deg}(r) > 0$ **do**

(1) $R_i = \{r\} \cup \{\text{any } \text{deg}(r) - 1 \text{ elements of } F \text{ with degree } 0\}$.

(2) $F_i^* = \eta^{-1}(R_i)$.

(3) $F = F \setminus R_i, F^* = F^* \setminus F_i^*, i = i + 1$.

end while

$R_i = F, F_i^* = F^*, r = i$.

When the algorithm terminates, the sequences $\{R_i\}_{i=1}^r$ and $\{F_i^*\}_{i=1}^r$ form partitions of F and F^* , respectively. It is straightforward to show that **(a)** in all iterations, step (1) will be able to find $\text{deg}(r) - 1$ elements of degree 0, **(b)** $|R_i| = |F_i^*|$ for all $i = 1, 2, \dots, r$, and **(c)** all elements of R_r have degree 0. Moreover, the following holds.

Fact 2.5 *If $j \in R_i$ and $\varphi^*(j) \notin F_i^*$, then $\eta(\varphi^*(j)) \notin R_i$.*

2.4.2 The Cost of a Local Optimum: the t -Swap Case

Based on the potential t -swaps, we now give upper bounds on several possible moves. The following proofs are largely similar to the proof of Lemma 2.2.

Lemma 2.6 *If $|R_i| = |F_i^*| \leq t$, then*

$$\text{kmed}([F \setminus R_i] \cup F_i^*) - \text{kmed}(F) \leq \sum_{j \in N^*(F_i^*)} (O_j - A_j) + \sum_{j \in N(R_i)} 2O_j \quad (2.6)$$

Proof: Consider the following assignment of clients after removing R_i and adding F_i^* : map each client $j \in N^*(F_i^*)$ to $\varphi^*(j)$ and map each client $j \in N(R_i) \setminus N^*(F_i^*)$ to $\eta(\varphi^*(j))$. All other clients stay where they were. We know from Fact 2.5 that this assignment is legal, because we do not assign any clients to the facilities being removed. The assignment gives an upper bound on the cost of $\text{kmed}([F \setminus R_i] \cup F_i^*)$. Applying the projection lemma (Lemma 2.4), we know that each $j \in N(R_i)$ has an increase at most $2O_j$. The rest of the proof is similar to the proof of Lemma 2.2. ■

Lemma 2.7 *If $|R_i| = |F_i^*| = s > t$, then*

$$\frac{1}{s-1} \sum_{(f^*, r) \in F_i^* \times \widehat{R}_i} [\text{kmed}(F + f^* - r) - \text{kmed}(F)] \leq \sum_{j \in N^*(F_i^*)} (O_j - A_j) + \sum_{j \in N^*(R_i)} 2 \left(1 + \frac{1}{t}\right) O_j, \quad (2.7)$$

where \widehat{R}_i is a set of $s-1$ degree-0 elements of R_i .

Proof: Consider each individual swap in $(f^*, r) \in F_i^* \times \widehat{R}_i$ that removes r and adds f^* . By Fact 2.5, we know that the assignment in Lemma 2.2 is legal. Thus the same analysis as in Lemma 2.2 gives that $\text{kmed}(F + f^* - r) - \text{kmed}(F) \leq \sum_{j \in N^*(f^*)} (O_j - A_j) + \sum_{j \in N(r)} 2O_j$.

Now consider each $j \in N^*(F_i^*)$. We know that f^* appears in exactly $s-1$ pairs of $F_i^* \times \widehat{R}_i$, so summing over these, we have $\sum_{j \in N^*(F_i^*)} (O_j - A_j)$. Likewise, each $j \in N^*(R_i)$ appears in exactly s pairs of $F_i^* \times \widehat{R}_i$. But, $\frac{s}{s-1} \leq 1 + \frac{1}{t}$, so summing over these, we have $\sum_{j \in N^*(R_i)} 2(1 + \frac{1}{t}) O_j$. Together, we have the lemma. ■

We now combine the two Lemmas. Note that for $|R_i| = |F_i^*| \leq t$, the optimality condition implies that $\text{kmed}([F \setminus R_i] \cup F_i^*) - \text{kmed}(F) \geq 0$, and for $|R_i| = |F_i^*| > t$, the optimality condition gives that $\text{kmed}(F + f^* - r) - \text{kmed}(F) \geq 0$ for all $f^* \in F_i^*$ and $r \in R_i$. Therefore, by noting that $\{R_i\}_{i=1}^r$ and $\{F_i^*\}_{i=1}^r$ are partitions of F and F^* , respectively, we establish

$$\begin{aligned} 0 &\leq \sum_{i: |R_i| \leq t} \left(\text{kmed}([F \setminus R_i] \cup F_i^*) - \text{kmed}(F) \right) + \sum_{i: |R_i| > t} \frac{1}{|R_i| - 1} \sum_{(f^*, r) \in F_i^* \times \widehat{R}_i} \text{kmed}(F + f^* - r) - \text{kmed}(F) \\ &\leq \sum_{i=1}^r \left[\sum_{j \in N^*(F_i^*)} (O_j - A_j) + \sum_{j \in N(R_i)} 2(1 + 1/t) O_j \right] = \text{kmed}(F^*) - \text{kmed}(F) + 2(1 + 1/t) \text{kmed}(F^*), \end{aligned}$$

which proves the following theorem.

Theorem 2.8 ([AGK⁺04]) *At a local minimum F for the t -swap version of the k -median local search algorithm, $\text{kmed}(F) \leq (3 + 2/t) \text{kmed}(F^*)$.*

3 The ℓ_p -Norm k -Facility Location Problem

In this section, we consider the following common generalization of the k -median, k -means, and k -center problems: given a metric space with a point set V and distances $d(\cdot, \cdot)$, and a value $p \geq 1$, the ℓ_p -norm

k -facility location problem is to find a set F of k facilities to minimize the objective function

$$\Phi_p(F) = \left(\sum_{j \in V} d(j, F)^p \right)^{1/p}. \quad (3.8)$$

Like the k -median problem, this setting has a natural local-search algorithm: Starting with k facilities, the algorithm tries to find $i \in F$ and $j \notin F$ such that the “swap” $F' = (F \setminus \{i\}) \cup \{j\}$ improves the objective value. The algorithm stops when no improving move exists. We will show that this algorithm is a $5p$ approximation, which can be improved to a $(3 + \frac{2}{t})p$ approximation by allowing t -swaps. This algorithm is *asymptotically tight* as we will show in Section 3.3. Even for the single-swap case, this immediately implies a 5-approximation for k -median, a 10-approximation for the k -means problem (which can be improved to 9), and an $O(\log n)$ -approximation for the k -center problem. The implication for k -center follows from the fact that for vectors of length n , the norm $\|\cdot\|_\infty$ is within a constant factor of the norm $\|\cdot\|_{\log n}$, and hence $\max_{j \in V} d(j, F)$ is within a constant factor of $\Phi_{\log n}(F)$. To the best of our knowledge, the results for $p \neq 1$ give the first analyses of local-search heuristics for these problems on general metric spaces.

3.1 Analyzing the Single-swap Case

We begin with the single-swap case and analyze the t -swap case in the next section. Let $\Delta(p, q) = d^p(p, q)$, and assume that V is indexed by $\{1, 2, \dots, n\}$. We borrow notations and definitions from our analysis of k -Median in Section 2. To argue about the quality of our solution, we will use the set S of test swaps defined in Section 2.1. Using an argument similar to the proof of the k -Median problem, we establish:

$$\begin{aligned} 0 &\leq \sum_{(r, f^*) \in S} \Phi_p^p(F + f^* - r) - \Phi_p^p(F) \leq \sum_{(r, f^*) \in S} \left(\sum_{j \in N^*(f^*)} (O_j^p - A_j^p) + \sum_{j \in N(r)} \Delta(j, \eta(\varphi^*(j))) - A_j^p \right) \\ &\leq \sum_{j \in V} O_j^p - 3 \sum_{j \in V} A_j^p + 2 \sum_{j \in V} \Delta(j, \eta(\varphi^*(j))) = \Phi_p^p(F^*) - 3\Phi_p^p(F) + 2 \sum_{j \in V} \Delta(j, \eta(\varphi^*(j))) \end{aligned} \quad (3.9)$$

We proceed to derive an upper-bound on the sum $\sum_{j \in V} \Delta(j, \eta(\varphi^*(j)))$ in terms of $\Phi_p(F^*)$ and $\Phi_p(F)$ as follows.

Claim 3.1

$$\sum_{j \in V} \Delta(j, \eta(\varphi^*(j))) \leq (2\Phi_p(F^*) + \Phi_p(F))^p$$

Proof: Define $\mathbf{x} = \langle d(j, \eta(\varphi^*(j))) \rangle_{j=1}^n$, and $\mathbf{y} = \langle 2d(j, \varphi^*(j)) + d(j, \varphi(j)) \rangle_{j=1}^n$. The Projection Lemma (Lemma 2.4) gives that $\mathbf{x}_j \leq \mathbf{y}_j$ for all $j \in V$, and thus $\|\mathbf{x}\|_{\ell_p} \leq \|\mathbf{y}\|_{\ell_p}$. Now note that $\|\langle d(j, \varphi^*(j)) \rangle_{j=1}^n\|_{\ell_p} = \Phi_p(F^*)$ and $\|\langle d(j, \varphi(j)) \rangle_{j=1}^n\|_{\ell_p} = \Phi_p(F)$. Therefore, applying the triangle inequality on the (\mathbb{R}^n, ℓ_p) space, we have $\sum_{j \in V} \Delta(j, \eta(\varphi^*(j))) = \|\mathbf{x}\|_{\ell_p}^p \leq \|\mathbf{y}\|_{\ell_p}^p \leq (2\Phi_p(F^*) + \Phi_p(F))^p$. ■

Together with (3.9), this claim implies

$$0 \leq \Phi_p^p(F^*) - 3\Phi_p^p(F) + 2(2\Phi_p(F^*) + \Phi_p(F))^p \quad (3.10)$$

To complete the proof, we let α be the smallest positive number such that $\Phi_p^p(F) \leq \alpha^p \cdot \Phi_p^p(F^*)$ and show that $\alpha \leq 5p$. Suppose for a contradiction that $\alpha > 5p$. Consider the right-hand side of inequality

(3.10), which we rewrite as follows.

$$\Phi_p^p(F^*) - 3\Phi_p^p(F) + 2(2\Phi_p(F^*) + \Phi_p(F))^p = \alpha^p \Phi_p^p(F^*) \left(\frac{1}{\alpha^p} - 3 + 2(1 + 2/\alpha)^p \right) \quad (3.11)$$

Let $f(p, \alpha) = \frac{1}{\alpha^p} + 2(1 + 2/\alpha)^p - 3$. Since $\alpha^p \Phi_p^p(F^*)$ is always non-negative, if we know that, for $\alpha > 5p$, $f(\alpha) < 0$, we will have a contradiction to (3.10). For $p = 1$ and 2 , we can explicitly solve for α . For $p \geq 2$, we know that if $\alpha > 5p$, then $\frac{1}{\alpha^p} \leq 1/100$ and $(1 + 2/\alpha)^p \leq e^{2/5}$. Thus, we have $\frac{1}{\alpha^p} + 2(1 + 2/\alpha)^p \leq \frac{1}{100} + 2 \cdot e^{2/5} - 3 < 0$. The results can be summarized in the following theorem:

Theorem 3.2 *The natural local-search algorithm for the ℓ_p -facility location problem gives a $5p$ approximation guarantee. Additionally, for $p = 2$, this algorithm gives a 9 approximation guarantee on general metric spaces.*

3.2 Analyzing the t -Swap Case

We will now analyze the t -swaps case. The analysis here will be largely similar to the analysis of the single-swap case, except for few extra ingredients which we will now develop. Our analysis will use the set of test swaps from Section 2.4. We will borrow notations and definitions from our analysis of the t -swap case of k -Median in Section 2.4.

Let us recall that F and F^* denote the facilities in our solution and the facilities in the optimal solution, respectively; the sequences $\{R_i\}_{i=1}^r$ and $\{F_i^*\}_{i=1}^r$ denote partitions of F and F^* , respectively. We derive the following upper bounds.

Lemma 3.3 *If $|R_i| = |F_i^*| \leq t$, then*

$$\Phi_p^p([F \setminus R_i] \cup F_i^*) - \Phi_p^p(F) \leq \sum_{j \in N^*(F_i^*)} (O_j^p - A_j^p) + \sum_{j \in N(R_i)} (\Delta(j, \eta(\varphi^*(j))) - A_j^p) \quad (3.12)$$

Proof: Consider the following assignment of clients after removing R_i and adding F_i^* : map each client $j \in N^*(F_i^*)$ to $\varphi^*(j)$ and map each client $j \in N(R_i) \setminus N^*(F_i^*)$ to $\eta(\varphi^*(j))$. All other clients stay where they were. We know from Fact 2.5 that this assignment is legal, because we do not assign any clients to the facilities being removed. The assignment gives an upper bound on the cost of $\Phi_p^p([F \setminus R_i] \cup F_i^*)$. This immediately gives $\Phi_p^p([F \setminus R_i] \cup F_i^*) - \Phi_p^p(F) \leq \sum_{j \in N^*(F_i^*)} (O_j^p - A_j^p) + \sum_{j \in N(R_i) \setminus N^*(F_i^*)} (\Delta(j, \eta(\varphi^*(j))) - A_j^p) \leq \Phi_p^p([F \setminus R_i] \cup F_i^*) - \Phi_p^p(F) \leq \sum_{j \in N^*(F_i^*)} (O_j^p - A_j^p) + \sum_{j \in N(R_i)} (\Delta(j, \eta(\varphi^*(j))) - A_j^p)$, which proves the lemma. \blacksquare

Lemma 3.4 *If $|R_i| = |F_i^*| = s > t$, then*

$$\begin{aligned} \frac{1}{s-1} \sum_{(f^*, r) \in F_i^* \times \widehat{R}_i} [\Phi_p^p(F + f^* - r) - \Phi_p^p(F)] &\leq \sum_{j \in N^*(F_i^*)} (O_j^p - A_j^p) \\ &+ \sum_{j \in N^*(R_i)} \left(1 + \frac{1}{t}\right) (\Delta(j, \eta(\varphi^*(j))) - A_j^p), \end{aligned}$$

where \widehat{R}_i is a set of $s - 1$ degree-0 elements of R_i .

Proof: Consider each individual swap in $(f^*, r) \in F_i^* \times \widehat{R}_i$ that removes r and adds f^* . Again, we will reassign every $j \in N^*(f^*)$ to f^* and every $j \in N(r) \setminus N^*(f^*)$ to r , and keep all other clients where

they were. By Fact 2.5, we know that this assignment is legal, and thus $\Phi_p^p(F + f^* - r) - \Phi_p^p(F) \leq \sum_{j \in N^*(f^*)} (O_j^p - A_j^p) + \sum_{j \in N(r)} (\Delta(j, \eta(\varphi^*(j))) - A_j^p)$.

Now consider each $j \in N^*(F_i^*)$. We know that f^* appears in exactly $s - 1$ pairs of $F_i^* \times \widehat{R}_i$, so summing over these, we have $\sum_{j \in N^*(F_i^*)} (O_j^p - A_j^p)$. Likewise, each $j \in N^*(R_i)$ appears in exactly s pairs of $F_i^* \times \widehat{R}_i$. But, $\frac{s}{s-1} \leq 1 + \frac{1}{t}$, so summing over these, we have $\sum_{j \in N^*(R_i)} (1 + \frac{1}{t}) (\Delta(j, \eta(\varphi^*(j))) - A_j^p)$. Together, we have the lemma. \blacksquare

Using these two lemmas and the fact that $\{R_i\}_{i=1}^r$ and $\{F_i^*\}_{i=1}^r$ are partitions of F and F^* , we establish the following bound:

$$0 \leq \sum_{i=1}^r \left(\sum_{j \in N^*(F_i^*)} (O_j^p - A_j^p) + \sum_{j \in N(R_i)} (1 + \frac{1}{t}) (\Delta(j, \eta(\varphi^*(j))) - A_j^p) \right) \quad (3.13)$$

$$= \Phi_p^p(F^*) - (2 + \frac{1}{t}) \Phi_p^p(F) + (1 + \frac{1}{t}) \sum_{j \in V} \Delta(j, \eta(\varphi^*(j))). \quad (3.14)$$

As shown in Claim 3.1, the sum $\sum_{j \in V} \Delta(j, \eta(\varphi^*(j)))$ is upper-bounded by $(2\Phi_p(F^*) + \Phi_p(F))^p$; therefore, we have

$$\Phi_p^p(F^*) - (2 + \frac{1}{t}) \Phi_p^p(F) + (1 + \frac{1}{t}) (2\Phi_p(F^*) + \Phi_p(F))^p \geq 0 \quad (3.15)$$

We are now ready to prove the following theorem:

Theorem 3.5 *For any value of $t \in \mathbb{Z}_+$, the natural t -swap local-search algorithm for the ℓ_p -facility location problem yields the following guarantees:*

1. For $p = 1$, it is a $(3 + \frac{2}{t})$ approximation.
2. For $p = 2$, it is a $(5 + \frac{4}{t})$ approximation.
3. For all reals $p \geq 2$, it is a $(3 + \frac{2}{t})p$ approximation.

Proof: Let α be the smallest positive number such that $\Phi_p^p(F) \leq \alpha^p \cdot \Phi_p^p(F^*)$ in inequality (3.15). For $p = 1$ and $p = 2$, we can directly solve the inequality and obtain the desired results. For $p \geq 2$, we assume for a contradiction that $\alpha > (3 + \frac{2}{t})p$ and proceed as follows. With the assumption, we have

$$\begin{aligned} & \Phi_p^p(F^*) - (2 + \frac{1}{t}) \Phi_p^p(F) + (1 + \frac{1}{t}) (2\Phi_p(F^*) + \Phi_p(F))^p \\ & \leq \alpha^p \Phi_p^p(F^*) \left(\frac{1}{\alpha^p} - (2 + \frac{1}{t}) + (1 + \frac{1}{t}) (1 + \frac{2}{\alpha})^p \right) \\ & = \alpha^p \Phi_p^p(F^*) \left(\frac{1}{\alpha^p} + (1 + \frac{2}{\alpha})^p - 2 + \frac{1}{t} ((1 + \frac{2}{\alpha})^p - 1) \right) \end{aligned}$$

Now define $f(p) = (\frac{1}{(3+2/t)p})^p + (1 + \frac{2}{(3+2/t)p})^p$ and $g(p) = (1 + \frac{2}{(3+2/t)p})^p$. It is easy to see that for $p \geq 2$, $f(p)$ and $g(p)$ are non-decreasing functions, and thus for all $p \geq 2$, $f(p) \leq \lim_{p \rightarrow \infty} f(p) = \exp\{\frac{2t}{2+3t}\}$ and

$g(p) \leq \lim_{p \rightarrow \infty} g(p) = \exp\left\{\frac{2t}{2+3t}\right\}$. Note also that $\exp\left\{\frac{2t}{2+3t}\right\} = \exp\left\{1 - \frac{2+t}{2+3t}\right\} \leq e \cdot \left(1 - \frac{2+t}{2+3t}\right) = e \cdot \frac{2t}{2+3t}$.

$$\begin{aligned} & \Phi_p^p(F^*) - \left(2 + \frac{1}{t}\right)\Phi_p^p(F) + \left(1 + \frac{1}{t}\right)(2\Phi_p(F^*) + \Phi_p(F))^p \\ & \leq \alpha^p \Phi_p^p(F^*) \left(e^{2t/(2+3t)} - 2 + \frac{1}{t}(e^{2t/(2+3t)} - 1) \right) \\ & \leq \alpha^p \Phi_p^p(F^*) \left(\frac{1+t}{t} \cdot e \cdot \frac{2t}{2+3t} - \frac{1}{t} - 2 \right) \\ & \leq \alpha^p \Phi_p^p(F^*) \left(\frac{2e(1+t)}{2+3t} - \frac{1}{t} - 2 \right) \end{aligned}$$

Simple algebra shows that $\frac{2e(1+t)}{2+3t} - \frac{1}{t} \leq 1.9$; therefore, $\Phi_p^p(F^*) - \left(2 + \frac{1}{t}\right)\Phi_p^p(F) + \left(1 + \frac{1}{t}\right)(2\Phi_p(F^*) + \Phi_p(F))^p \leq \alpha^p \Phi_p^p(F^*) \left(\frac{2e(1+t)}{2+3t} - \frac{1}{t} - 2\right) < 0$, which gives a contradiction. \blacksquare

3.3 An Asymptotically Tight Example

Inspired by a lower bound given by Kanungo et al. [KMN⁺04], we present an example where the local-search algorithm produces an $\Omega(p)$ -approximate solution for the the ℓ_p -facility location problem. The example presented below is designed for the single-swap case, but it can be generalized to the t -swap case.

Theorem 3.6 *For every p , there are instances of the ℓ_p -norm k -facility location problem where the cost of some local minima (under the standard local moves) is at least $2p$ times the optimal cost.*

Proof: Consider a 2-dimensional torus with lattice points $\{0, 1, \dots, N-1\}^2$ for some large integer N . The lattice points are labeled even or odd according to the parity of the sum of their coordinates. For a value x to be fixed later, we define a gadget $D(x)$ to be a set of 4 points at $(\pm x, 0), (0, \pm x)$.

We overlay a graph on the torus as follows. Every lattice point is a facility node. Centered at each even lattice point is a copy of $D(x)$. These gadget points make up our client nodes. To set up a distance metric, we introduce the following edges and define the distance between any two nodes as the shortest path between them. As shown in Figure 3.3, an even lattice point is at distance x from any of its surrounding gadget points, and an odd lattice point is at distance $1-x$ from its (physical) neighboring gadget points. Under this distance metric, the distance between a and b , for example, is $1+x$.




Figure 3.2: A lower-bound instance

Let $x = 1/(2p+1)$ and $k = N^2/2$. For this choice of k and x , an optimal solution, denoted by F^* , opens facilities at all even lattice points, yielding $\Phi_p(F^*) = (4k)^{1/p}x$. To get a $2p$ approximation, consider a solution F which opens k facilities at all odd lattice points. It is easy to see that $\Phi_p(F) = (4k)^{1/p}(1-x)$, and so $\frac{\Phi_p(F)}{\Phi_p(F^*)} = 2p$. It remains to show that F is a local optimum. Consider that, for any $f^* \in F$ and $r \in F$, if we shut down a facility r and open a facility f^* , the change in cost is given by

$$\Phi_p^p(F - r + f^*) - \Phi_p^p(F) \geq 4\left(x^p - (1-x)^p\right) + 3\left((1+x)^p - (1-x)^p\right). \quad (3.16)$$

With our choice of x , it is straightforward to show that $\Phi_p^p(F - r + f^*) - \Phi_p^p(F) \geq 0$. Thus, the solution F is a local optimum as desired. \blacksquare

4 Simpler Proofs for Other Previously Known Results

4.1 Uncapacitated Facility Location

In the metric uncapacitated facility location problem, instead of a hard upper bound k on the number of facilities, we are given an opening cost f_i for each location $i \in V$, and the goal is to minimize the objective function

$$\text{UFL}(F) = \sum_{i \in F} f_i + \sum_{j \in V} d(j, F).$$

This problem has been extremely widely studied and many constant-factor approximation algorithms are known: see Section 1.2 for many references.

The Local Search Moves. Since we do not have a hard bound on the number of facilities, we can have a richer set of local moves—**(a)** opening a facility, **(b)** closing a currently open facility, and **(c)** swapping facilities as above. Again, we let F^* be the optimal set of facilities, and F be the algorithm’s set of facilities at a local minimum.

4.1.1 Bounding the Local Optimum Cost for UFL

Since we have been using f to denote a generic facility, let us use $\text{fac}(f)$ to denote the facility opening cost for facility f , and $\text{fac}(F')$ to denote the cost $\sum_{f \in F'} \text{fac}(f)$ of a set F' of facilities. Again, for a client $j \in C$, let O_j and A_j be the connection cost in the optimal and local-optimal solutions, and let φ^* and φ denote the maps assigning clients to facilities. Hence $\text{opt} = \text{fac}(F^*) + \sum_{j \in C} O_j$, whereas $\text{alg} = \text{fac}(F) + \sum_{j \in C} A_j$. Moreover, for a facility $f \in F$, let $N(f)$ denote the clients assigned to it; similarly, define $N^*(f^*)$ to be the clients assigned to it in the optimal solution. The first lemma below is easy: try opening each facility in F^* , note the change in cost, and add things up.

Lemma 4.1 (Connection Cost [KPR00]) *At a local optimum, the fact that “open new facility” moves are non-improving implies the connection cost $\sum_{j \in C} A_j \leq \text{fac}(F^*) + \sum_{j \in C} O_j$.*

Lemma 4.2 (Facility Cost [AGK⁺04]) *The facility cost $\text{fac}(F) \leq \text{fac}(F^*) + 2 \sum_{j \in C} O_j$ at a local optimum.*

Proof: Recall the notation of Section 2: given F^* and F , let $\eta : F^* \rightarrow F$ map each optimal facility to a closest facility in F . Following [AGK⁺04], call a facility $f \in F$ “good” if $\eta^{-1}(f) = \emptyset$, and “bad” otherwise.

If a facility f is good, we can consider closing the facility and assigning any clients $j \in N(f)$ to $\widehat{f} = \eta(\varphi^*(j))$: note that since f is good, we know that $\widehat{f} \neq f$, and hence this reassignment is valid. By the projection lemma (Lemma 2.4), the total increase in the assignment cost is $d(j, \widehat{f}) - d(j, f) \leq 2O_j$, and hence from local optimality, we get that for any good $f \in F$,

$$-\text{fac}(f) + \sum_{j \in N(f)} 2O_j \geq 0. \quad (4.17)$$

For a bad facility f , let P_f^* be the set $\eta^{-1}(f) = \{f_0^*, f_1^*, \dots, f_t^*\}$ (with $t \geq 0$), and let f_0^* be the closest one to f . We then consider the t possible moves of opening facility f_i^* in $P_f^* \setminus \{f_0^*\}$, and assigning any client $j \in N^*(f_i^*) \cap N(f)$ to f_i^* . The local optimality ensures that

$$\text{fac}(f_i^*) + \sum_{j \in N^*(f_i^*) \cap N(f)} (O_j - A_j) \geq 0. \quad (4.18)$$

Moreover, consider the move of opening f_0^* and closing f :

- Any client $j \in N(f)$ with $\varphi^*(j) \notin P_f^*$ is assigned to the facility $\eta(\varphi^*(j)) \neq f$: the projection Lemma 2.4 implies that the increase in connection cost for such j is at most $2O_j$.
- Any client $j \in N(f)$ with $\varphi^*(j) = f_i^* \in P_f^*$ (for some $i \in \{0, 1, \dots, t\}$) is assigned to f_0^* . The change in the connection cost is $d(j, f_0^*) - d(j, f)$.

Hence, local optimality shows that

$$\text{fac}(f_0^*) - \text{fac}(f) + \sum_{j \in N(f) \wedge \varphi^*(j) \notin P_f^*} 2O_j + \sum_{i=0}^t \sum_{j \in N^*(f_i^*) \cap N(f)} (d(j, f_0^*) - A_j) \geq 0. \quad (4.19)$$

Adding (4.19) with the t inequalities (4.18) (one for each $i \in \{1, \dots, t\}$) gives us

$$\text{fac}(P_f^*) - \text{fac}(f) + \sum_{j \in N(f) \wedge \varphi^*(j) \notin P_f^*} 2O_j + \sum_{i=0}^t \sum_{j \in N^*(f_i^*) \cap N(f)} (d(j, f_0^*) + O_j - 2A_j) \geq 0. \quad (4.20)$$

Consider the rightmost sum in (4.20): for $i = 0$, the summand is $2(O_j - A_j) \leq 2O_j$. For $i \neq 0$,

$$d(j, f_0^*) + d(j, f_i^*) - 2d(j, f) \leq (d(j, f) + d(f, f_0^*)) + d(j, f_i^*) - 2d(j, f) \quad (4.21)$$

$$\leq d(f, f_i^*) + d(j, f_i^*) - d(j, f) \quad (4.22)$$

$$\leq 2d(j, f_i^*) = 2O_j, \quad (4.23)$$

where we used the fact that $d(f, f_0^*) \leq d(f, f_i^*)$ in (4.22), and the triangle inequality at other places. Now the expression (4.20) can be simplified to say

$$\text{fac}(P_f^*) - \text{fac}(f) + \sum_{j \in N(f)} 2O_j \geq 0. \quad (4.24)$$

Summing (4.24) over all bad f , and (4.17) over all the good f , we get

$$\text{fac}(F^*) - \text{fac}(F) + \sum_{j \in C} 2O_j \geq 0. \quad (4.25)$$

which proves the claimed bound $\text{fac}(F) \leq \text{fac}(F^*) + \sum_j 2O_j$. ■

Combining the facility cost and connection cost lemmas above results in the following theorem.

Theorem 4.3 *At a local optimum, $\text{UFL}(F) \leq 2\text{fac}(F^*) + 3 \sum_j O_j \leq 3 \text{UFL}(F^*)$.*

4.2 k -Uncapacitated Facility Location

Building on the techniques developed in the previous sections, we can now give proofs for the *metric k -uncapacitated facility location problem* (k -UFL) problem. This is a common generalization of the k -median and UFL problems: not only do we have an opening cost f_i for each location (like UFL), but we also have a limit k on the number of facilities (like k -median). The goal is still to minimize the cost

$$\text{kUFL}(F) = \text{fac}(F) + \sum_{j \in V} d(j, F),$$

where $\text{fac}(F) = \sum_{i \in F} f_i$. This problem was defined by [DGK⁺05], whose main result—showing that local search was a 5-approximation—is reproved in this section.

The Local Search Moves. We start with any set F of at most k facilities, and allow the following modes: the algorithm can open a new facility if $|F| < k$, it can close a facility in F , or swap an open facility in F with a currently closed facility; as usual, the algorithm performs a move only if the total cost decreases.

Since we want to argue about some local optimum F , we can assume that $|F| = k$: indeed, if $|F| < k$, then it is a local optimum with respect to all moves—opening, closing, or swapping facilities, and then the result of Section 4.1 shows that F is a 3-approximation. However, if $|F| = k$, then we cannot open facilities even if we want to, and hence have to work harder for the proof.

4.2.1 Pairing for k -UFL

As with previous proofs, the central ingredient of the proof is an appropriate pairing that allows us to bound the cost of the local-search solution by applying the local optimality condition. To generate the pairing, we proceed as follows.

- Pair degree-1 facilities $f \in F$ with $\eta^{-1}(f)$. The degree-1 group is called “single.”
- Pair higher degree facilities in the following manner. If $\deg(f) \geq 2$, let $P_{f^*} = \eta^{-1}(f) = \{f_0^*, f_1^*, \dots, f_t^*\}$, where f_0^* is the facility closest to f , and the facilities f_1^*, \dots, f_t^* are ordered arbitrarily. Additionally, let $f_0 = f$ and f_1, \dots, f_t be any t distinct degree-0 facilities in F . Since $|F| \leq |F^*|$, we will always be able to find enough degree-0 facilities. Let $P_f = \{f_0, f_1, \dots, f_t\}$. The pairs (P_f, P_{f^*}) are called “heavy strips.”
- At this point, some degree-0 facilities are still left unmentioned; we call them “excess” facilities.

4.3 Bounding the Cost of a Local Optimum

We apply the local optimality condition to the pairs as follows. For each single pair (f, f^*) , we could swap f^* for f , assigning all $j \in N^*(f^*)$ to f^* and $j \in N(f) \setminus N^*(f^*)$ to $\eta(\varphi^*(j))$. The same reasoning as before shows that this is a valid reassignment. The local optimality condition, together with the projection Lemma 2.4, gives

$$\text{fac}(f^*) - \text{fac}(f) + \sum_{j \in N^*(f^*)} (O_j - A_j) + \sum_{j \in N(f) \setminus N^*(f^*)} 2O_j \geq 0 \quad (4.26)$$

Now consider a heavy strip (P_f, P_{f^*}) . Suppose $P_f = \{f_0, \dots, f_t\}$ and $P_{f^*} = \{f_0^*, \dots, f_t^*\}$. First, we could swap f_0^* for f_0 , assigning all $j \in N^*(f_0^*)$ to f_0^* , all $j \in N^*(f_i^*) \cap N(f_0)$ (for $i = 1, 2, \dots, t$) to f_0^* ,

and all $j \in N(f_0) \setminus (\cup_{i=1}^t N^*(f_i^*))$ to $\eta(\varphi^*(j))$. This is a legal assignment as can be easily checked. By the local optimality condition, we have

$$\begin{aligned} \text{fac}(f_0^*) - \text{fac}(f_0) + \sum_{j \in N^*(f_0^*)} (O_j - A_j) \\ + \sum_{i=1}^t \sum_{j \in N^*(f_i^*) \cap N(f_0)} (d(j, f_0^*) - A_j) + \sum_{j \in N(f_0) \setminus (\cup_{i=1}^t N^*(f_i^*))} 2O_j \geq 0 \end{aligned} \quad (4.27)$$

Within the same strip (P_f, P_{f^*}) , we can also exchange f_i for f_i^* (for each $i = 1, 2, \dots, t$):

- If we assign all $j \in N^*(f_i^*) \cap (N(f_0) \cup N(f_i))$ to f_i^* and all $j \in N(f_i) \setminus N^*(f_i^*)$ to $\eta(\varphi^*(j))$, then the local optimality condition yields

$$\text{fac}(f_i^*) - \text{fac}(f_i) + \sum_{j \in N^*(f_i^*) \cap (N(f_0) \cup N(f_i))} (O_j - A_j) + \sum_{j \in N(f_i) \setminus N^*(f_i^*)} 2O_j \geq 0. \quad (4.28)$$

- If we assign all $j \in N^*(f_i^*)$ to f_i^* and all remaining $j \in N(f_i)$ to $\eta(\varphi^*(j))$, then the local optimality condition yields

$$\text{fac}(f_i^*) - \text{fac}(f_i) + \sum_{j \in N^*(f_i^*)} (O_j - A_j) + \sum_{j \in N(f_i) \setminus N^*(f_i^*)} 2O_j \geq 0. \quad (4.29)$$

Finally, consider deleting the “excess” facilities. For each excess facility f , we could delete it and assign all facilities $j \in N(f)$ to $\eta(\varphi^*(j))$ —recall f has degree 0. By the local optimality, we have

$$-\text{fac}(f) + \sum_{j \in N(f)} 2O_j \geq 0. \quad (4.30)$$

Adding up (4.26)-(4.30) across all strips and all groups yields the following claim.

Claim 4.4 *For any $j \in C$, the increase in connection cost is upper bounded by $5O_j - A_j$.*

Proof: Consider a client j . Note that j is uniquely assigned to a facility $f^* = \varphi^*(j)$ in the optimal solution. If f^* is in the degree-1 group, the increase in connection cost is clearly upper bounded by $5O_j - A_j$. Otherwise, f^* belongs to a heavy strip $(\{f_0, \dots, f_t\}, \{f_0^*, \dots, f_t^*\})$, in which case we consider the following possibilities:

- If $j \in N^*(f_0)$, then the increase in connection cost is upper bounded by $2O_j + 2O_j + (O_j - A_j) \leq 5O_j - A_j$.
- If $j \in N^*(f_i^*) \cap N(f_0)$ for some $i \in [t]$, then the increase in connection cost is upper bounded by $d(j, f_0^*) - A_j + (O_j - A_j) + (O_j - A_j) \leq 3O_j - A_j$. The inequality follows from the fact that $d(j, f_0^*) \leq d(j, \varphi(j)) + d(\varphi(j), f_0^*) \leq d(j, \varphi(j)) + d(\varphi(j), f_i^*) \leq d(j, \varphi(j)) + d(\varphi(j), j) + d(j, f_i^*) = 2A_j + O_j$.
- If $j \in N^*(f_i^*) \cap N(f_i)$ for some $i \in [t]$, then the increase in connection cost is clearly upper bounded by $2(O_j - A_j) \leq 5O_j - A_j$.
- Otherwise j must belong to $N^*(f_i^*) \setminus (N(f_0) \cup N(f_i))$. In this case, the increase in connection cost is upper bounded by $2O_j + 2O_j + O_j - A_j = 5O_j - A_j$.

We conclude that for all $j \in C$, the increase in connection cost is at most $5O_j - A_j$. ■

It follows that $2\text{fac}(F^*) - \text{fac}(F) + \sum_{j \in C} (5O_j - A_j) \geq 0$, resulting in the following theorem.

Theorem 4.5 *At a local minimum, $\text{kUFL}(F) \leq 2\text{fac}(F^*) + 5 \sum_j d(j, F^*) \leq 5 \text{kUFL}(F^*)$.*

References

- [AGK⁺04] Vijay Arya, Naveen Garg, Rohit Khandekar, Adam Meyerson, Kamesh Munagala, and Vinayaka Pandit. Local search heuristics for k -median and facility location problems. *SIAM J. Comput.*, 33(3):544–562, 2004.
- [ARS03] Aaron Archer, Ranjithkumar Rajagopalan, and David B. Shmoys. Lagrangian relaxation for the k -median problem: new insights and continuity properties. In *Algorithms—ESA 2003*, volume 2832 of *Lecture Notes in Comput. Sci.*, pages 31–42. Springer, Berlin, 2003.
- [Byr07] Jaroslav Byrka. An optimal bifactor approximation algorithm for the metric uncapacitated facility location problem. In *Proceedings of the 10th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, 2007. 29–43.
- [CG05] Moses Charikar and Sudipto Guha. Improved combinatorial algorithms for facility location problems. *SIAM J. Comput.*, 34(4):803–824, 2005.
- [CGTS02] Moses Charikar, Sudipto Guha, Éva Tardos, and David B. Shmoys. A constant-factor approximation algorithm for the k -median problem. *J. Comput. System Sci.*, 65(1):129–149, 2002. Special issue on STOC, 1999 (Atlanta, GA).
- [Chu98] Fabián A. Chudak. Improved approximation algorithms for uncapacitated facility location. In *Integer programming and combinatorial optimization (Houston, TX, 1998)*, volume 1412 of *Lecture Notes in Comput. Sci.*, pages 180–194. Springer, Berlin, 1998.
- [CS03] Fabián A. Chudak and David B. Shmoys. Improved approximation algorithms for the uncapacitated facility location problem. *SIAM J. Comput.*, 33(1):1–25, 2003.
- [DGK⁺05] Nikhil Devanur, Naveen Garg, Rohit Khandekar, Vinayaka Pandit, Amin Saberi, and Vijay Vazirani. Price of anarchy, locality gap, and a network service provider game. In *Proceedings of the First Workshop on Internet and Network Economics*, pages 1046–1055, 2005.
- [GK99] Sudipto Guha and Samir Khuller. Greedy strikes back: improved facility location algorithms. *J. Algorithms*, 31(1):228–248, 1999.
- [Gon85] Teofilo F. Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoret. Comput. Sci.*, 38(2-3):293–306, 1985.
- [Hoc82] Dorit S. Hochbaum. Heuristics for the fixed cost median problem. *Math. Programming*, 22(2):148–162, 1982.
- [HS86] Dorit S. Hochbaum and David B. Shmoys. A unified approach to approximation algorithms for bottleneck problems. *J. Assoc. Comput. Mach.*, 33(3):533–550, 1986.
- [JMS02] Kamal Jain, Mohammad Mahdian, and Amin Saberi. A new greedy approach for facility location problems. In *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing*, pages 731–740, New York, 2002. ACM.
- [JV01] Kamal Jain and Vijay Vazirani. Approximation algorithms for metric facility location and k -median problems using the primal-dual schema and Lagrangean relaxation. *Journal of the ACM*, 48(2):274–296, 2001. (Preliminary version in *40th FOCS*, pages 2–13, 1999).
- [KMN⁺04] Tapas Kanungo, David M. Mount, Nathan S. Netanyahu, Christine D. Piatko, Ruth Silverman, and Angela Y. Wu. A local search approximation algorithm for k -means clustering. *Comput. Geom.*, 28(2-3):89–112, 2004.
- [KPR00] Madhukar R. Korupolu, C. Greg Plaxton, and Rajmohan Rajaraman. Analysis of a local search heuristic for facility location problems. *J. Algorithms*, 37(1):146–188, 2000. (Preliminary version in 9th SODA, 1998).
- [KSS04] Amit Kumar, Yogish Sabharwal, and Sandeep Sen. A simple linear time $(1 + \epsilon)$ -approximation algorithm for k -means clustering in any dimensions. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS'04)*, pages 454–462, Washington, DC, USA, 2004. IEEE Computer Society.
- [Llo82] S.P. Lloyd. Least squares quantization in PCM. *IEEE Trans. Inform. Theory*, 28(2):129–137, 1982.

- [MMSV01] Mohammad Mahdian, Evangelos Markakis, Amin Saberi, and Vijay Vazirani. A greedy facility location algorithm analyzed using dual fitting. In *Approximation, randomization, and combinatorial optimization (Berkeley, CA, 2001)*, volume 2129 of *Lecture Notes in Comput. Sci.*, pages 127–137. Springer, Berlin, 2001.
- [MYZ02] Mohammad Mahdian, Yinyu Ye, and Jiawei Zhang. Improved approximation algorithms for metric facility location problems. In *Approximation algorithms for combinatorial optimization*, volume 2462 of *Lecture Notes in Comput. Sci.*, pages 229–242. Springer, Berlin, 2002.
- [Pan04] Vinayaka Pandit. *Local Search Heuristics for Facility Location Problems*. PhD thesis, Indian Institute of Technology, New Delhi, 2004.
- [PT03] Martin Pál and Éva Tardos. Group strategyproof mechanisms via primal-dual algorithms. In *Proceedings of the 44th Symposium on the Foundations of Computer Science (FOCS)*, pages 584–593, 2003.
- [STA97] David B. Shmoys, Éva Tardos, and Karen I. Aardal. Approximation algorithms for facility location problems. In *Proceedings of the 29th ACM Symposium on the Theory of Computing (STOC)*, pages 265–274, 1997.
- [Svi02] Maxim Sviridenko. An improved approximation algorithm for the metric uncapacitated facility location problem. In *Integer programming and combinatorial optimization*, volume 2337 of *Lecture Notes in Comput. Sci.*, pages 240–257. Springer, Berlin, 2002.