Advanced Approximation Algorithms **(CMU 18-854B, Spring 2008)**

Lecture 6: Facility location: greedy and local search algorithms

31-Jan-08

Lecturer: Anupam Gupta *Scribe: Viswanath Nagarajan*

1 A modified greedy algorithm

Recall the greedy algorithm for non-metric facility location that was done in Lecture 4, which was obtained by formulating the problem as set cover. In this section we present a modified greedy algorithm for the *metric* facility location problem that achieves a constant approximation ratio. Let the facility location instance consist of clients D, facilities F, metric d on $D \cup F$, and facility opening costs $f_i \in \mathbb{R}_+$ for each $i \in F$. Then the goal is to open a set $F^* \subseteq F$ of facilities that minimizes $\sum_{i \in F^*} f_i + \sum_{j \in D} d(j, F^*)$. Above for any client $j \in D$, $d(j, F^*) = \min_{i \in F^*} d(j,i)$ denotes the distance from j to its nearest facility in F^*.

As seen in Lecture 4, the facility location problem can be cast as a set covering problem where elements are the clients D, and sets correspond to 'stars' centered at some facility. Any set can be represented as (i, A) where $i \in F$ is a facility and $A \subseteq D$ is a subset of clients; the cost of this set is $f_i + \sum_{j \in A} d(i,j)$. The modified greedy algorithm is as follows.

- Set $F' \leftarrow \phi$.
- While ($D \neq \phi$) do:
  1. Pick a set (i, A) where $i \in F$ and $A \subseteq D$ that minimizes the ratio $\frac{f_i + \sum_{j \in A} d(i,j)}{|A|}$.
  2. Set $F' \leftarrow F' \cup \{i\}$, $f_i \leftarrow 0$, and $D \leftarrow D \setminus A$.
- Output F' as the set of open facilities.

Note that the only difference from the greedy algorithm of Lecture 4 is that the same facility may be picked multiple times in step 1 (however the facility cost is non-zero only the first time it is picked). In order to analyze this algorithm, we study a primal-dual procedure (different from one seen in Lecture 5) which turns out to be identical to the modified greedy algorithm!

1.1 Another primal-dual algorithm

Recall the primal and dual LPs for the facility location problem:

$$
(P)\quad
\begin{array}{lll}
\min & \sum_i f_i y_i + \sum_j \sum_i d(i,j) x_{i,j} & \\
& \sum_i x_{i,j} \geq 1 & \forall j \in D \\
& x_{i,j} \leq y_i & \forall j \in D \;\; i \in F \\
& y_i \geq 0 & \forall i \in F \\
& x_{i,j} \geq 0 & \forall j \in D \;\; i \in F
\end{array}
$$

$$
\begin{array}{lll}
\max & \sum_j \alpha_j & \\
 & \sum_j \beta_{i,j} \leq f_i & \forall i \in F \\
(D) & \alpha_j - \beta_{i,j} \leq d(i,j) & \forall j \in D \;\; i \in F \\
 & \alpha_j \geq 0 & \forall j \in D \\
 & \beta_{i,j} \geq 0 & \forall j \in D \;\; i \in F
\end{array}
$$

Consider the following procedure that maintains an integral primal solution (implicitly given by F') and an infeasible dual solution (α, β).

1. Initialize $F' \leftarrow \phi$ (open facilities), $D' \leftarrow D$ (clients with active duals), $\alpha \leftarrow 0$ and $\beta \leftarrow 0$.

2. While $D' \neq \phi$, uniformly raise duals α_j (for $j \in D'$) and $\beta_{i,j}$ (for $j \in D'$ and $i \in F \setminus F'$ such that $\alpha_j \geq d(i,j)$) until one of the following happens:

   (a) $\alpha_j = d(i,j)$ for some $j \in D'$ and $i \in F'$. In this case, set $D' \leftarrow D' \setminus \{j\}$ and $\beta_{i,j} \leftarrow 0$ for all $i \in F$. *Client j is assigned to already open facility i.*

   (b) $\sum_j \beta_{i,j} = f_i$ for some $i \in F \setminus F'$. In this case, let $A = \{j \in D' \mid \beta_{i,j} > 0\}$. Set $\beta_{i',j} \leftarrow 0$ for all $i' \in F \setminus \{i\}$ and $j \in A$. Modify $F' \leftarrow F' \cup \{i\}$ and $D' \leftarrow D' \setminus A$. *Clients A are assigned to newly opened facility i.*

Let (α, β) be the dual solution at the end of the algorithm, and F' the final set of facilities opened. It is clear that for each client $j \in D$, $\beta_{i,j} > 0$ for at most one facility $i \in F$. Also for any facility $i \in F'$, if $A_i \subseteq D$ denotes the set of clients assigned to facility i when it was added to F' then $f_i = \sum_{j \in A_i} \beta_{i,j} = \sum_{j \in A_i} (\alpha_j - d(i,j))$. Finally note that for any client $j \in D \setminus \cup_{i \in F'} A_i$ that got assigned to an already open facility i, we have $\alpha_j = d(i,j)$. Since the A_is are disjoint,

$$
\begin{aligned}
\sum_{i \in F'} f_i + \sum_{j \in D} d(j, F') &= \sum_{i \in F'} \sum_{j \in A_i} (\alpha_j - d(i,j)) + \sum_{j \in D} d(j, F') \\
&= \sum_{i \in F'} \sum_{j \in A_i} \alpha_j + \sum_{j \in (D \setminus \cup_{i \in F'} A_i)} d(j, F') \\
&= \sum_{j \in D} \alpha_j
\end{aligned}
$$

So we have an integral primal solution of objective value at most the infeasible dual solution (α, β). We next show that $(\alpha/3, \beta)$ is a feasible dual solution, which would show that that the above primal-dual algorithm achieves an approximation guarantee of 3.

Claim 1.1. *For any facility $i \in F$ and clients $j, j' \in D$ with $\alpha_j \geq d(i,j)$ and $\alpha_{j'} \geq d(i,j')$, we have $\alpha_j \leq d(i,j) + 2\alpha_{j'}$*

Proof. If $\alpha_j \leq \alpha_{j'}$ the claim is trivial; so we assume otherwise. Note that each client $l \in D$ gets assigned to some facility exactly at time α_l in the algorithm. Since $\alpha_{j'} < \alpha_j$, it follows that client j' got assigned to some facility i' at time $\alpha_{j'}$, when client j was still unassigned. Note that facility i' is open at least from time $\alpha_{j'}$, but client j is not assigned to it until (at least) time $\alpha_j > \alpha_{j'}$: so $d(i',j) \geq \alpha_j$ (otherwise j would have got assigned to i' at some time $< \alpha_j$). Using the triangle inequality, $d(i',j) \leq d(i',j') + d(j',i) + d(i,j) \leq 2\alpha_{j'} + d(i,j)$ (note that $d(j',i') \leq \alpha_{j'}$ as j' is assigned to facility i', and $d(j',i) \leq \alpha_{j'}$ by assumption). Thus we have $\alpha_j \leq d(i',j) \leq 2\alpha_{j'} + d(i,j)$. ∎

Claim 1.2. *For any facility* $i \in F$, $\sum_{j \in D} \max\{0, \frac{\alpha_j}{3} - d(i,j)\} \le f_i$.

Proof. Fix a facility i for the rest of the proof. Let $C_i = \{j \in D \mid \frac{\alpha_j}{3} \ge d(i,j)\}$, note that $\sum_{j \in D} \max\{0, \frac{\alpha_j}{3} - d(i,j)\} = \sum_{j \in C_i} (\frac{\alpha_j}{3} - d(i,j))$. Number the clients in C_i, 1 through $k = |C_i|$ such that $\alpha_1 \le \alpha_2 \le \cdots \le \alpha_k$, so $C_i = \{1, \cdots, k\}$. We now claim that $\sum_{j=1}^{k} \max\{0, \alpha_1 - d(i,j)\} \le f_i$. Suppose (for a contradiction) that this were not true: then at time α_1 in the algorithm, none of the clients in C_i have been assigned and hence at that point $\beta_{i,j} = \max\{0, \alpha_1 - d(i,j)\}$ for all $j \in C_i$; and this would imply $\sum_j \beta_{i,j} > f_i$ at time α_1, a contradiction!

Using Claim 1.1 with $j' = 1$ and any $j \in C_i = \{1, \cdots, k\}$, we have $\alpha_j \le d(i,j) + 2\alpha_1$. Now,

$$\sum_{j=1}^{k} \alpha_j \le \sum_{j=1}^{k} (2\alpha_1 + d(i,j)) = 3\sum_{j=1}^{k} d(i,j) + 2\sum_{j=1}^{k} (\alpha_1 - d(i,j)) \le 3\sum_{j=1}^{k} d(i,j) + 2f_i$$

So we have $\sum_{j=1}^{k} (\frac{\alpha_j}{3} - d(i,j)) \le f_i$, which proves the claim. $\square$

Claim 1.2 implies the feasibility of the dual solution $(\alpha/3, \beta)$ to linear program (D), and hence:

Theorem 1.3. *The primal-dual algorithm described above is a 3-approximation algorithm for metric facility location.*

In fact, a more careful analysis [3] shows that $(\alpha/\rho, \beta)$ is also a feasible dual for $\rho \approx 1.861$ (this uses an auxiliary LP to capture the worst realization of the infeasible dual α, distances d and facility costs f). So this algorithm has an approximation guarantee at most 1.861.

1.2 Relating the greedy and primal-dual algorithms

Consider any instant of time t in the primal-dual algorithm, when the set of 'active' (or unassigned) clients is D': the duals $\alpha_j = t$ for all $j \in D'$. If condition 2(a) does not currently apply, $t \le d(i,j)$ for all $i \in F'$ (already open facility) and $j \in D'$ (unassigned client). If condition 2(b) does not currently apply, $\sum_{j \in D'} \max\{0, t - d(i,j)\} \le f_i$ for all $i \in F \setminus F'$ (unopened facility). So the next event in the primal-dual algorithm corresponds occurs exactly at time:

$$\hat{t} = \min\left\{ \min_{j \in D',\, i \in F'} d(i,j), \quad \min_{i \notin F',\, S \subseteq D'} \frac{f_i + \sum_{j \in S} d(i,j)}{|S|} \right\}$$

Observe that this is precisely the criterion that is used in deciding the greedy augmentation (the facilities F' have their costs reset to 0, so for $i \in F'$ minimizing $\frac{\sum_{j \in S} d(i,j)}{|S|}$ is equivalent to minimizing $d(i,j)$). Hence the two algorithms make an identical sequence of decisions. Thus:

Theorem 1.4. *The modified greedy algorithm achieves an approximation guarantee of 3 for metric facility location.*

2 Local search

Local search is a technique that involves starting with any feasible solution to the problem at hand, and repeatedly making small changes in the solution as long as these (strictly) improve the objective. In the context of facility location, if $S \subseteq F$ denotes the current solution (set of open facilities) then a local change involves moving to a solution S' where the symmetric difference $|S \triangle S'| \leq 1$. Let $\psi(S)$ denote the cost of solution S to the facility location problem. The algorithm chooses to move to any neighbor S' of S that decreases the total cost (i.e. $\psi(S') < \psi(S)$). A solution S is said to be locally optimal if it has no neighbor that decreases total cost. The following are two important questions regarding the local search procedure.

- If S^* denotes a globally optimal solution and S denotes any locally optimal solution, how large can $\frac{\psi(S)}{\psi(S^*)}$ be?
- How fast does the algorithm converge to a local optimum?

The above mentioned local search procedure for facility location satisfies $\frac{\psi(S)}{\psi(S^*)} \leq 3$ for any locally optimal S [1]. It has also been shown [1] that for every solution S there exists a neighbor S' with $\psi(S) - \psi(S') \geq \frac{\psi(S) - 3\psi(S^*)}{n^2}$, which implies that one can obtain a $3 + o(1)$ approximate solution in polynomial time by this local search algorithm.

2.1 Local search for k-Median

In this section we study a local search algorithm for the k-median problem, which is closely related to facility location. In the k-median problem, we are given n points in a metric space (V, d) and a parameter k; the goal is to find a set $F \subseteq V$ with $|F| \leq k$ that minimizes $\sum_{j \in V} d(j, F)$. Given a set F of at most k facilities, the k-median cost is denoted

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

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 [1], and that the bound of 5 is tight for these local-search dynamics. Here we give a simpler proof of this 5-approximation.[1]

2.1.1 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 locally optimal, each of these swaps must be non-improving, which 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 $|F| = |F^*| = k$.

[1]Thanks to Kanat Tangwongsan for providing his writeup on this simpler version!

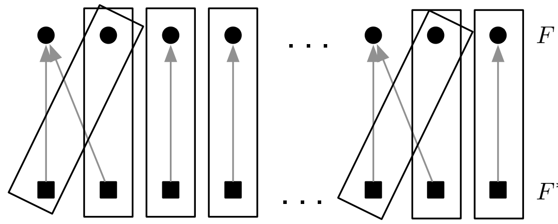

Figure 1: An example mapping $\eta\colon F^* \to F$ and a set of test swaps $\mathcal{P}$.

Define a map $\eta : F^* \to F$ 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 $\mathcal{P} = \{(r, f^*)\} \subseteq R \times F^*$ such that

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

The procedure is simple: for each $r \in R$ with in-degree 1, construct the pair $(f, \eta^{-1}(r))$—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 $(r, f^*) \in \mathcal{P}$ and $\widehat{f^*} \in F^*$ with $\widehat{f^*} \neq f^*$, $\eta(\widehat{f^*}) \neq r$.*

Intuition for the Pairing. To get some intuition for why the pairing $\mathcal{P}$ 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.

2.1.2 Bounding the Cost of a Local Optimum

In this section, we use the fact that each of the swaps in set $\mathcal{P}$ constructed in Section 2.1.1 are non-improving to show that that the local optimum has small cost.

Breaking ties arbitrarily, assume that $\varphi : V \to F$ and $\varphi^* : V \to F^*$ 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 it's 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 $(r, f^*) \in \mathcal{P}$,*

$$0 \leq \mathsf{kmed}(F + f^* - r) - \mathsf{kmed}(F) \leq \sum_{j \in N^*(f^*)} (O_j - A_j) + \sum_{j \in N(r)} 2\, O_j. \tag{1}$$

Proof. The first inequality is obvious since we are at a local optimum. 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 $\widehat{f^*} = \varphi^*(j)$ (note that $\widehat{f^*} \neq f^*$): assign j to $\widehat{f} = \eta(\widehat{f^*})$, the closest facility in F to $\widehat{f^*}$. Note that by Fact 2.1, $\widehat{f} \neq r$, 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.

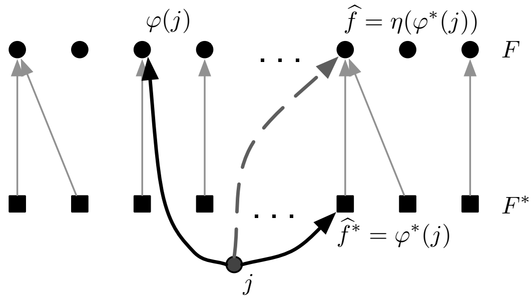

Note that for any client $j \in N^*(f^*)$, the change in cost is exactly $O_j - A_j$: summing over all these clients gives us the first term in the expression (1).

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

$$d(j, \widehat{f}) - d(j, r) \leq d(j, \widehat{f^*}) + d(\widehat{f^*}, \widehat{f}) - d(j, r) \tag{2}$$

$$\leq d(j, \widehat{f^*}) + d(\widehat{f^*}, r) - d(j, r) \tag{3}$$

$$\leq d(j, \widehat{f^*}) + d(j, \widehat{f^*}) = 2\, O_j. \tag{4}$$

with (2) and (4) following by the triangle inequality, and (3) using the fact that $\widehat{f}$ is the closest vertex in F to $\widehat{f^*}$. Summing up, the total change for all these clients is at most

$$\sum_{j \in N(r) \setminus N^*(f^*)} 2\, O_j \leq \sum_{j \in N(r)} 2\, O_j, \tag{5}$$

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

Note that summing (1) over all tuples in $\mathcal{P}$, 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:

$$0 \leq \sum_{f^* \in F^*} \sum_{j \in N^*(f^*)} (O_j - A_j) + 2 \sum_{r \in F} \sum_{j \in N(r)} 2\, O_j = 5\ \mathsf{kmed}(F^*) - \mathsf{kmed}(F)$$

Theorem 2.3. *At a local minimum F, the cost $\mathsf{kmed}(F) \leq 5 \cdot \mathsf{kmed}(F^*)$.*

3 Hardness of approximation

It is known [2] that the metric facility location problem has no polynomial time algorithm achieving an approximation guarantee better than 1.463, unless $NP \subseteq DTIME(n^{O(\log\log n)})$. This is essentially a reduction from the set-cover problem (see Guha and Khuller [2] for more details). A more direct reduction from the max-coverage problem can be used to show that the k-median problem is hard to approximate better than a $1 + 2/e$ factor.

References

[1] V. Arya, N. Garg, R. Khandekar, K. Munagala, and V. Pandit. Local search heuristic for k-median and facility location problems. In *STOC '01: Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 21–29, 2001.

[2] S. Guha and S. Khuller. Greedy strikes back: improved facility location algorithms. In *SODA '98: Proceedings of the ninth annual ACM-SIAM symposium on Discrete algorithms*, pages 649–657, 1998.

[3] K. Jain, M. Mahdian, E. Markakis, A. Saberi, and V. V. Vazirani. Greedy facility location algorithms analyzed using dual fitting with factor-revealing lp. *J. ACM*, 50(6):795–824, 2003.