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

Lecture 4: Uncapacitated Facility Location

Jan 24, 2008

Lecturer: Anupam Gupta *Scribe: S. Harsha Vardhan*

In this lecture, as well as the next two lectures, we will study the *uncapacitated facility location problem*, using it as a vehicle to illustrate several different techniques that are commonly used to devise approximation algorithms for NP-hard problems.

1 The general Uncapacitated facility location problem

Definition 1.1. (Uncapacitated facility location)
Input:

- set D of clients
- set F of potential facility locations
- a distance function $d : D \times F \to \mathbb{R}_+$ (for $i \in F$, let $f_i \triangleq f(i)$)
- a cost function $f : F \to \mathbb{R}_+$ for the set F (for $i \in F, j \in D$, let $d_{ij} \triangleq d(i,j)$)

Output: $S \subseteq F$ that minimizes $\sum_{i \in S} f_i + \sum_{j \in D}(\min_{i \in S} d_{ij})$

A Simple logarithmic approximation This problem can be formulated as a set cover problem. Given an instance (D, F, f, d) of the facility location problem, it can be transformed into an instance of a set cover problem $(X, \mathcal{S}, cost)$ by

- setting $X = D$, and
- setting $\mathcal{S} = \{(i, A) : i \in F, A \subseteq D\}$. (i, A) covers elements of the set A. The cost of the set (i, A) is $f_i + \sum_{j \in A} d_{ij}$.

We can then use the standard greedy algorithm for set cover: at every iteration, the algorithm chooses the set (i, A) which minimizes

$$\frac{Cost(i, A)}{|A \cap Uncovered|}.$$

To select such a set in polynomial time, the algorithm iterates over all possible facility locations $i \in F$ and $s \in \mathbb{Z}_{\geq 0}$, the set size of $A \cap Unconverted$. For a given i and a set size s, the minimum cost set can be found by including the s uncovered clients closest to facility i.

This algorithm gives a $\log |D|$ approximation guarantee. One can show that the (non-metric) facility location problem is at least as hard as the set cover problem, and hence this $\log |D|$ factor approximation is the best one can hope to give for.

2 The metric uncapacitated facility location problem

In this version of the problem, the clients and facilities are embedded in a metric space and thus satisfy the triangle inequality, $d_{ik} \leq d_{ij} + d_{jk}$, where $i, j, k \in D \cup F$.

2.1 The above-mentioned greedy algorithm

As we shall see, there are constant factor approximation algorithms for this problem. However, there exist instances of the problem where the greedy algorithm described in the previous section has a $\Omega(\frac{\log n}{\log \log n})$ gap [5]. Consider the instance where there are k facilities each of cost p^k at the same location. There are $k-1$ sets of clients $S_1, S_2, \ldots, S_{k-1}$. Each client S_i contains p^{k-i+1} clients all located at equal distances from the facilities, the sum of the distances being $\sum_{j=1}^{i} p^{j-1}$. The greedy algorithm chooses all the facilities incurring a cost of at least kp^k. The optimal solution chooses exactly one facility, and incurs a cost of $p^k + \Omega(p^{k-1})$. The ratio of the costs is $\Omega(k)$ and $n = \Omega(p^k)$. If we set $p = \log n$, then k, the gap between the optimal cost and the cost between the solution returned by the greedy algorithm is $\Omega(\log_p n) = \Omega(\frac{\log n}{\log p}) = \Omega(\frac{\log n}{\log \log n})$.

On the other hand, if we refine the cost of set (i, A) that the greedy algorithm uses as follows:

$$cost(i, A) = \begin{cases} f_i + \sum_{j \in A} d_{ij}, & \text{if facility } i \text{ has been selected} \\ \sum_{j \in A} d_{ij}, & \text{otherwise,} \end{cases} \tag{1}$$

the algorithm guarantees a constant factor approximation guarantee. (This will be proved in a subsequent lecture.)

2.2 Results for the metric facility location problem

The following is a brief list of results related to this problem:

- Shmoys, Tardos and Aardal [10] give a $O(1)$-approximation algorithm using LP rounding
- Jain and Vazirani [7] give a 3-approximation algorithm using Primal-Dual methods
- Arya et. al. [1] give a $O(1)$-approximation algorithm using local search
- Jain, Mahdian and Saberi [6] give a 1.61-approximation algorithm using a greedy algorithm. Mahdian, Ye and Zhang [9] proved that this algorithm is a $(1.11, 1.78)$ bi-criteria approximation. A (γ_f, γ_c) bicriteria approximation algorithm yields a solution with cost less than $\gamma_f F^* + \gamma_c C^*$, where F^* is the cost of facilities in the optimal solution and C^* is the sum of distances from clients to facilities in the optimal solution. [6] also proves the $(\gamma_f, 1 + 2e^{-\gamma_f})$ approximatability limit curve.
- Byrka [3] gives a $(1.6774, 1.3738)$ bicriteria approximation algorithm. Using this along with the $(1.11, 1.78)$ bicriteria approximation of [6], it is possible to obtain an algorithm which

outputs solution with expected cost not greater than 1.5 times the optimal (by simply choosing to use the algorithm of [6] with probability $p = 0.313$ and that of [3] with probability $1 - p$).

- Guha and Khuller [4] gives the best hardness result known so far: it is hard to approximate better than a factor of 1.463.

2.3 The k-median Problem

A problem closely related to the metric facility location problem is the *metric k-median problem* which differs in the following respects.

- it has an added constraint that limits the number of facilities that can be selected to at most k and
- all f_i are 0.

The best known approximation algorithm for the metric k-median problem is a $(3+\epsilon)$-approximation that runs in $O(n^{\frac{1}{\epsilon}})$ time, and the best hardness result known for the problem is a factor of $1 + 2/e$. Also, approximating the non-metric version of this problem by any finite factor C is NP-hard as the following reduction from the dominating set problem shows:

Given an instance $\mathcal{G} = (G = (V, E), k), V = \{1, 2, \ldots, n\}$ of the dominating set problem, construct instance $\mathcal{H}$ of k-median problem, where k medians are to be chosen as follows:

- Facility set $F = \{1, 2, \ldots, n\}$, where $n = |V|$.
- Client set $D = \{1, 2, \ldots, n\}$.
- for all $i \in F, j \in D$, the distance function,

$$d_{ij} = \begin{cases} 0, & \text{if } i = j \\ 1, & \text{if } (i,j) \in E \\ (1 + C)(n - k) + 1, & \text{if } (i,j) \notin E \end{cases} \tag{2}$$

If there is a dominating set of size k in instance $\mathcal{G}$ there is a solution to $\mathcal{H}$ that costs no more than $(n - k)$ and a factor-C approximation algorithm would give a solution of cost at most $C(n - k)$. If there is no solution to instance $\mathcal{G}$, then every solution to instance $\mathcal{H}$ will be of cost at least $C(n - k) + 1$, thus helping us determine whether $\mathcal{G}$ had a dominating set.

3 An LP rounding algorithm for UFL

Balinski [2] proposed the following Integer Linear Program (ILP) formulation of the metric facility location problem:

$$minimize : \sum_i f_i y_i + \sum_{i,j} d_{ij} x_{ij} \tag{3}$$

$$constraints : \sum_i x_{ij} \geq 1 \quad \forall j \in D \tag{4}$$

$$y_i \geq x_{ij} \quad \forall j \in D, i \in F \tag{5}$$

$$x_{ij} \in \{0,1\} \quad \forall j \in D \tag{6}$$

$$y_i \in \{0,1\} \quad \forall i \in F \tag{7}$$

Indeed, it is easy to verify that optimum solutions for this integer program correspond to optimal solutions for the facility location instance. The natural linear programming relaxation of this integer program relaxes the last two constraints to $x_{ij} \geq 0$ and $y_i \geq 0$. As every solution to the Integer program is also a solution to the linear program, the minimum value of the linear program lower bounds the minimum value of the integer program and therefore, the optimal value to the problem.

The LP rounding algorithm starts with an optimal solution x^*, y^* to the above LP: let $cost(LP)$ be the cost of this optimal solution, and let $F^* = \sum_i f_i y_i^*$ be the *fractional facility cost* and $C^* = \sum_{ij} d_{ij} x_{ij}^*$ be the *fractional connection cost*. The rounding algorithm then executes the following two steps:

Step 1 (Filtering): Let D_j denote $\sum_i d_{ij} x_{ij}$. Consider all the facilities to which a client j is fractionally assigned. Among these, let N_j denote the subset of "near" facilities located within a distance $2D_j$ of client j. The proof of the following claim (saying that the fractional solution assigns at least half of every client to "near" facilities) is very similar to the proof of Markov's inequality, and will be proved later.

Claim 3.1. *For any client j, $\sum_{i \in N_j} x_{ij}^* \geq 1/2$.*

We now construct a new solution x', y' by *filtering* as follows: Set x'_{ij}s corresponding to facilities i not in N_j to 0, and let all other x'_{ij} be the same as x^*_{ij}. Since the new solution x' may not satisfy the condition $\sum_i x_{ij} \geq 1$, we fix this problem by increasing the values of non-zero x'_{ij}s by a factor of 2 one-by-one until the constraint is satisfied (note that increasing all non-zero x'_{ij}s by a factor of 2 will always satisfy the constraint because the sum of non-zero x'_{ij}s is at least $1/2$). Set $y'_i = max_j\, x'_{ij}$. Note that $y'_i \leq 2\, y_i$. (This filtering technique was first used by Lin and Vitter [8].)

> At the end of the filtering step, the fractional solution (x', y') satisfies the following property: *If $x'_{ij} > 0$ then $d_{ij} \leq 2D_j$, and also $y'_i \leq 2y^*_i$.*

Step 2 (Rounding the Filtered Solution): Among all clients, select the client j with minimum D_j. Among the facilities in N_j, select the facility with the minimum f_i, say $k \in N_j$. Form a new

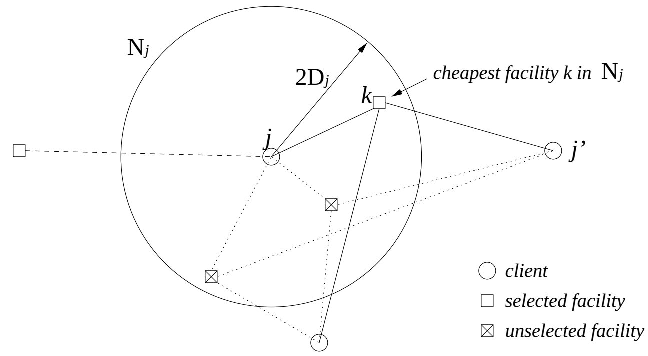

Figure 1: LP rounding

solution x'', y'' in which the client j is connected to the cheapest facility k in N_j and to no other facility ($x''_{kj} = 1$, $x''_{ij} = 0\ \forall i \in F \setminus \{k\}$). Facility k is selected into the solution ($y''_k = 1$) and all others facilities in N_j are dropped ($y''_i = 0\ \forall i \in N_j \setminus \{k\}$). Define the extended neighborhood of client j, EN_j, to be the set of all clients which are fractionally assigned to any facility in N_j. Allot all $j' \in EN_j$ to facility k alone, i.e., set $x''_{kj'} = 1$, $x''_{ij'} = 0\ \forall i \in F \setminus \{k\}$. Clients belonging to the set $EN_j \cup \{j\}$ are now considered allotted. The cost of connecting the client $j' \in EN_j$ to facility k is not greater than $(2\,D_j + 2 \times (2\,D_{j'})) \leq 6\,D_{j'}$ (as $D_j \leq D_{j'}$ by our choice of j). The cost of selecting facility k and dropping all other facilities in N_j is not greater than $\sum_{i \in N_j} f_i y'_i$.

Another round of allotments is started by selecting the unallotted client j with minimum D_j and allotting j and all unallotted clients in the extended neighborhood of j using the same procedure. The algorithm continues till all clients are allotted to some facility. It is easy to see that all facilities are touched exactly once in this algorithm and are either selected or dropped.

Analysis: Let $CostLPR$ be the cost of the final integral solution x'', y''.

$$\begin{aligned} CostLPR &\leq \sum_{i \in F} f_i y'_i + 6 \sum_{j \in D} D_j && (8)\\ &\leq 2 \sum_{i \in F} f_i y^*_i + 6 \sum_{i \in F, j \in D} d_{ij} x^*_{ij} && (9)\\ &= 2\,F^* + 6\,C^* \leq 6\,cost(LP). && (10) \end{aligned}$$

which gives us the claimed 6-approximation.

3.1 Tightening the Analysis

To see how to improve the analysis of the previous section: but first let us see a proof of (a generalization of) Claim 3.1. Define $N_j(\beta)$ denote the subset of “near” facilities located within a distance βD_j of client j: hence the N_j in the previous section is the same as $N_j(2)$. As mentioned above, the proof is very similar to Markov’s inequality.

Claim 3.2. $\sum_{i \in N_j(\beta)} x^*_{ij} \geq 1 - \frac{1}{\beta}$.

Proof. Let $F_j(\beta)$ denote the set of facilities to which a client j is fractionally assigned, but are farther than βD_j from j. Also, let $s_j = \sum_i x^*_{ij}$. For any feasible solution, $s_j \geq 1$. Suppose if possible, that $\sum_{i \in F_j(\beta)} x^*_{ij} > \frac{s_j}{\beta}$. Then,

$$\begin{aligned} D_j &= \sum_{i \in F_j(\beta)} d_{ij} x^*_{ij} + \sum_{i \notin F_j(\beta)} d_{ij} x^*_{ij} \\ &\geq \sum_{i \in F_j} d_{ij} x^*_{ij} \\ &\geq \beta D_j \sum_{i \in F_j} x^*_{ij} \\ &> \beta D_j \left(\frac{s_j}{\beta} \right) \\ &\geq D_j, \end{aligned}$$

which is a contradiction. Therefore $\sum_{i \in F_j} x^*_{ij} \leq \frac{s_j}{\beta}$ and $\sum_{i \in D_j} x^*_{ij} \geq s_j(1 - \frac{1}{\beta}) \geq 1 - \frac{1}{\beta}$. $\square$

Setting $\beta = 2$ gives the proof of Claim 3.1.

3.1.1 The improved algorithm

The filtering and rounding algorithm changes in the following way:

- We now set $x'_{ij} = 0$ if $i \notin N_j(\beta)$; to regain feasibility, we might have to scale up the other x^*_{ij} and y^*_i values by at most $\frac{\beta}{\beta-1}$.

  Hence, the invariant at the end of the modified filtering step ensures that if $x'_{ij} > 0$ then $d_{ij} \leq \beta D_j$, and also $y'_i \leq \frac{\beta}{\beta-1} \cdot y^*_i$.

- The rounding step is the same as above, and gives a solution whose cost is at most $\frac{\beta}{\beta-1} F^* + 3\beta C^*$.

To balance the two quantities, we set $\beta = 4/3$ and get a 4-approximation.

References

[1] V. Arya, N. Garg, R. Khandekar, K. Munagala, and V. Pandit. Local search heuristic for k-median and facility location problems. In *Proceedings of the Symposium on Theory of Computing*, pages 21–29, 2001.

[2] M. Balinski. Integer programming: Methods, uses, computation. *Management Science*, 12:253–313, 1965.

[3] J. Byrka. An optimal bifactor approximation algorithm for the metric uncapacitated facility location problem. In *International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, 2007.

[4] S. Guha and S. Khuller. Greedy strikes back: Improved algorithms for facility location. *Journal of Algorithms*, 31(1):228–248, 1999.

[5] K. Jain, M. Mahdian, E. Markakis, A. Saberi, and V. Vazirani. A greedy facility location algorithm analyzed using dual-fitting with factor-revealing lp. *Journal of the ACM*, 50(6):795–824, 2003.

[6] K. Jain, M. Mahdian, and A. Saberi. A new greedy approach for facility location problem. In *Proceedings of the Symposium on Theory of Computing*, 2002.

[7] K. Jain and V. V. Vazirani. Primal-dual approximation algorithms for metric facility location and k-median problems. In *Proceedings of the Annual Symposium on Foundation of Compute Science*, pages 2–13, 1999.

[8] J.-H. Lin and J. S. Vitter. Approximation algorithms for geometric median problems. Technical Report CS-92-37, Brown University, 1992.

[9] M. Mahdian, Y. Yea, and J. Zhang. Approximation algorithms for metric facility location problems. *SIAM Journal on Computing*, 36(2):411–432, 2006.

[10] D. B. Shmoys, É. Tardos, and K. Aardal. Approximation algorithms for facility location problems. In *Proceedings of the Symposium on Theory of Computing*, pages 265–274, 1997.