Network Design Foundations
Fall 2011
Lecture 10

Instructor: Mohammad T. Hajiaghayi
Scribe: Catalin-Stefan Tiseanu

November 2, 2011

1 Overview

We study constant factor approximation algorithms for **CONNECTED FACILITY LOCATION** and **METRIC UNCAPACITATED FACILITY LOCATION**.

2 Connected facility location

Suppose we are given a graph $G = (V, E)$, with a metric cost function c over the edges. (that is, a cost function which satisfies the triangle inequality). Suppose we have a set of clients $D \subseteq V$. Consider that facilities can be opened at any client vertex. Each client $i \in D$ has an associated demand d_i. There are no facility opening costs. The task is to connect each client to a facility. Let us call this the connection cost. Additionally, all the opened facilities need to be connected to each other, via higher-cost edges. Let us call this second cost the Steiner cost. Let a set of facilities $F \subseteq D$ be the set of opened facilities. Formally, we need to minimize the following objective function:

$$\text{Minimize} \sum_{j \in D} d_j \cdot l_c(i^*(j), j) + m \cdot c(T)$$

The first term in the objective function is the connection cost, and the second term is the Steiner cost. T is the Steiner tree used to connect the opened facilities. $i^*(j)$ is the closest opened facility to client $j \in D$. $l_c(v, u)$ is the minimum distance between vertex v and u. Finally, each edge used in T costs a multiplicative $m \geq 1$ more than an edge used for the connection cost. The real world intuition is the following: For connecting the facilities we buy our infrastructure (therefore it costs a factor m more). For the connection cost we can afford to rent the paths (since the clients may change with time).

Assume that all demands are 1 for now. Let us define the following algorithm:

Algorithm Connected facility location
1. For each $j \in D$ mark it with probability $\frac{1}{m}$. Let the set of marked vertices be $D' \subseteq D$.
2. Construct an ρ_{ST}-approximation Steiner tree of $F = D' \cup \{r\}$, where r is an arbitrary root vertex.
3. For each $j \in D$ assign it to the closest facility in F.

Note: ρ_{ST} represents the best constant for an approximation algorithm for **Steiner tree**, which at the time of the writing is 1.39.

Theorem 1 *The above algorithm is a* $2+\rho_{ST}$*-approximation for* **CONNECTED FACILITY LOCATION***. (shown in [1])*

The analysis will proceed in two steps. First, let $Z^* = C^* + S^*$ be an optimal cost solution, where C^* is the connection cost, and S^* is the Steiner cost in the optimal solution.

Lemma 1 *The expected Steiner cost* S *is at most* $\rho_{ST} \cdot Z^*$.

Proof: We will show that the expected cost of the optimal Steiner tree is at most Z^* (we pay a factor of ρ_{ST} to find it). Let T^* be the Steiner tree in the optimal solution. Construct a Steiner tree T on F as follows: For each vertex $v \in T^*$ open a facility and connect them using the edges in T^*; for the vertices $u \in F \setminus T^*$ use the shortest path from u to T^*. Now, for the vertices already in T^* we pay S^*, because we use the existing Steiner tree. Each of the other vertices $v \in F \setminus T^*$ needs to be connected to T with probability $\frac{1}{m}$. The expected cost of connecting them, assuming the worst case where $F \cap T = \emptyset$, is at most $\displaystyle\sum_{j \in D} \frac{1}{m} \cdot m \cdot l_c(j, i^*(j)) \leq C^*$. Therefore, the total cost of T is at most Z^*. However, we can only find an ρ_{ST}-approximation. The expected Steiner cost S is then at most $\rho_{ST} \cdot Z^*$.

Lemma 2 *The expected connection cost* C *is at most* $2Z^*$.

Proof: The key intuition is that the expected connection cost is independent of the the particular Steiner tree used in the previous step. Therefore let us use the minimum spanning tree on D (again, we consider the metric completion of the graph G). We will order the vertices in the order Prim's Minimum Spanning Tree Algorithm would visit them. Start with an arbitrary vertex v, and let $S = \{v\}$. Now, at each step consider $j \in D \setminus S$ such that $l_c(j, S)$ is minimized. Flip a coin:

- with probability $\frac{1}{m}$ add this vertex to the existing Steiner tree S, paying $m \cdot l_c(j, S)$. Add it to S and to a set A which hold the current Steiner tree.
- with probability $1 - \frac{1}{m}$ connect this vertex to the existing Steiner tree, paying $l_c(j, S)$. Add it to S and to a set B which holds all vertices for which we pay the connection cost.

For each $j \in B$, $\text{primcostB}_j = (1 - \frac{1}{m}) \cdot l_c(j, S) < l_c(j, S) = \frac{1}{m} \cdot m \cdot l_c(j, S) = \text{primcostA}_j$ is the cost, in expectation, paid for connecting it to the Steiner tree. Observe that it is smaller than the cost paid for adding it to the Steiner tree (primcostA_j). Therefore, we can bound the total connection cost $\sum_{j \in B} \text{primcostB}_j$ by the Steiner cost $\sum_{j \in D} \text{primcostA}_j$. Observe now that the Steiner cost in the algorithm is precisely the cost of the Minimum Spanning Tree. We know that the Minimum Spanning Tree is a 2-approximation for Steiner Tree. Therefore, $\sum_{j \in B} \text{primcostB}_j \leq 2 \cdot Z^*$.

Combining Lemma 1 and Lemma 2 we get a $2 + \rho_{ST}$ approximation algorithm for **CONNECTED FACILITY LOCATION**.

It remains to be said that this approach can be generalized to general demands by adjusting the probability for a vertex $j \in D$ to be marked to $\frac{d_j}{m}$.

3 Metric uncapacitated facility location

As in the previous problem, we are given a graph $G = (V, E)$, with a metric cost function on the edges $c(e)$, $e \in E$ (that is, a cost function which satisfies the triangle inequality). Suppose we have a set of clients $D \subseteq V$. Consider that we are given a set $F \subseteq V$ where we can open facilities. Each client $i \in D$ has an associated demand d_i. This time each facility $i \in F$ has an opening cost f_i. Additionally, we can have capacitated or uncapacitated facilities (here we only consider the uncapacitated version). The task is to connect each client to a facility, which needs to be opened beforehand. Let us call this the connection cost. Additionally, we pay an opening cost for all facilities which are used. Let us call this second cost the opening cost. Formally, we need to minimize the following objective function:

$$
\begin{aligned}
\text{Minimize} \quad & \sum_{i \in F} f_i + \sum_{i \in F} \sum_{j \in D} d_j c_{ij} x_{ij} \\
\text{subject to} \quad & \sum_{i \in F} x_{ij} = 1, \quad \forall j \in D \\
& x_{ij} \leq y_i, \quad \forall i \in F, j \in D \\
& x_{ij} \in \{0, 1\}, \quad \forall i \in F, j \in D \\
& y_i \in \{0, 1\}, \quad \forall i \in F
\end{aligned}
$$

What we gave above is an Integer Linear Program (ILP). Obviously, we cannot solve it efficiently. Therefore, we need to consider the fractional relaxation, where $x_{ij} \geq 0, \forall i \in F, j \in D$ and $y_i \geq 0, \forall i \in F$.

Theorem 2 **METRIC UNCAPACITATED FACILITY LOCATION** *has a constant factor approximation. (shown in [2])*

The key intuition is that figuring out the $y's$ is the hard part.

Definition 1 (g-closeness) *Given the values* $g_j, \forall j \in D$*, we say a feasible solution* (x, y) *is* g*-close if whenever* $x_{ij} > 0$ *we have that* $c_{ij} \leq g_j$, $\forall j \in D$.

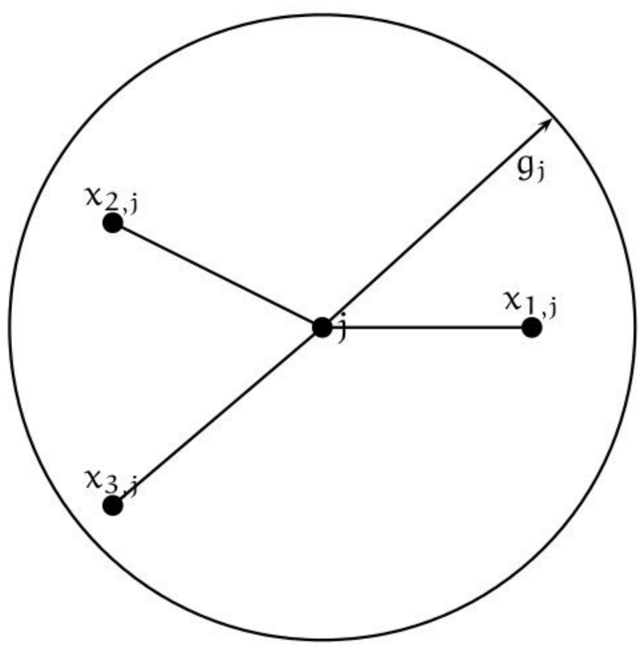

For a client $j \in D$ sort all the $i \in F$ facilities according to c_{ij}. Let $c_{\pi(1),j} \leq c_{\pi(2),j} \leq \ldots \leq c_{\pi(f),j}$ be the sorted order.

Definition 2 (α-point) *The* α*-point for a client* $j \in D$ *is defined as* $c_j(\alpha) = c_{\pi(i^*),j}$, *where* $i^* = \min\{i : \sum_{k=1}^{i} x_{\pi(k),j} \geq \alpha\}$.

Basically, $c_j(\alpha)$ gives us an α-fraction of the ball centered in j (use the above picture for visual reference). Let us get a bound on the α-point at $j \in D$, in terms of the cost spent over client j. Define $S = \{i : c_{ij} \geq c_j(\alpha)\}$. We have that $\sum_{i \in S} x_{ij} \geq 1 - \alpha$.

$$\sum_{i\in F} x_{ij}c_{ij} \geq \sum_{i\in S} x_{ij}c_{ij} \geq (1-\alpha)c_j(\alpha)$$
$$\implies c_j(\alpha) \leq \frac{1}{1-\alpha}\sum_{i\in F} x_{ij}c_{ij}$$

Thus, $c_j(\alpha)$ is intuitively the average cost connection. Now, let us use this bound to change feasible solutions (x, y) to feasible solutions $(\bar{x}, \bar{y})$ which are $c_j(\alpha)$-close without losing too much on the optimal solution. Formally,

Lemma 3 *For $\alpha \in (0, 1)$ we can change a feasible solution (x, y) to find a feasible solution $(\bar{x}, \bar{y})$ which is $c_j(\alpha)$-close and*

$$\sum_{i\in F} f_i\bar{y}_i \leq \frac{1}{\alpha}\sum_{i\in F} f_i y_i.$$

Proof: For each $j \in D$, let $\alpha_j = \sum_{i\in F, c_{ij}\leq c_j(\alpha)} x_{ij}$. Then, just set

$$\bar{x}_{ij} = \begin{cases} \frac{x_{ij}}{\alpha_j} & \text{if } c_{ij} \leq c_j(\alpha) \\ 0 & \text{otherwise} \end{cases}$$

Additionally, for each $i \in F$ set $\bar{y}_i = \min\{\frac{y_i}{\alpha}, 1\}$. Then, $(\bar{x}, \bar{y})$ will satisfy our requirement, because $\alpha_j \geq \alpha, \forall j \in D$.

■

Lemma 4 *Given a feasible, fractional g-close solution, we can find a feasible integer $3g$-close solution $(\hat{x}, \hat{y})$ such that*

$$\sum_{i\in F} f_i\hat{y}_i \leq \sum_{i\in F} f_i\bar{y}_i.$$

Proof: Let $\hat{F} = \{i \in F : 0 < \hat{y}_i < 1\}$ (this is the set of partially opened facilities). Let $\hat{D}$ be the set of clients assigned to facilities in $\hat{F}$. Note that $x_{ij} > 0 \implies i \in \hat{F}$. Also, $g_j = c_j(\alpha)$.

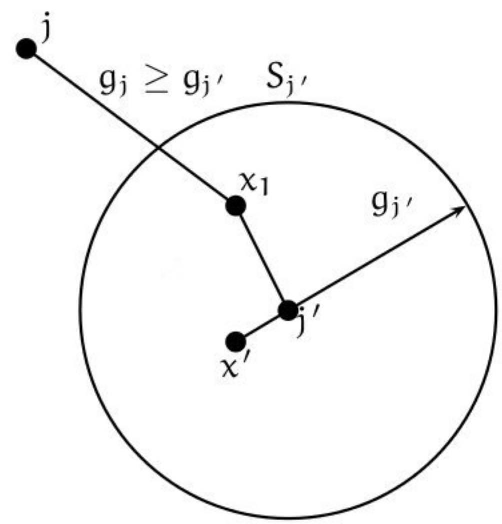

Now sort the clients by g_j. For client $j' \in \hat{D}$, taken in the sorted order do the following:

- **1.** Let $S_{j'} = \{i \in \hat{F} : \hat{x}_{ij'} > 0\}$ (this is the set of currently open facilities that j' is connected to. We will close all but one of these).

- **2.** Let i' be the facility with minimum opening cost. Assign j' completeley to it (set $\hat{y}_i' = 1$ and $\hat{y}_i'' = 0, \forall i'' \in S \setminus \{i'\}$). Adjust $\hat{x}_{ij'}$ accordingly. This does not increase the total facility opening cost.

- **3.** Let $T = \{j \in D : \exists i \in S \text{ such that } x_{ij} > 0\}$. T is exactly the set of clients which were assigned to facilities in S which are closed. Connect all these clients to i', by adjusting $\hat{x}_{i'j}$ accordingly.

The process ends when all clients $j \in D$ have $x_{ij} > 0$ coressponding to a facility i where $\hat{y}_i = 1$. Fully assign these clients to facility i (set $\hat{x}_{ij} = 1$).

Let's see how much the total cost went up.
In step 2 we started with a solution which was already $c_j(\alpha)$-close, therefore the total connection cost can increase by a factor of at most $\frac{1}{1-\alpha}$.
In step 3, we have for a client $j \in D$ that $g_j \geq g_j'$, due to the way the algorithm considers the clients. Now, the new connection cost, due the fact that that the cost distance function is metric, is less than $c_{i'j} = c_{i'j'} + c_{j'i} + c_{ij} \leq g_j + g_j' + g_j' \leq 3g_j$ (refer to the above picture for visual reference). Therefore the connection cost of elements in T increases by a factor of at most 3.

The final cost of $(\hat{x}, \hat{y})$, which is integral, is therefore

$$
\begin{aligned}
\mathsf{cost}(\hat{x},\hat{y}) &= \sum_{i\in F} f_i \hat{y}_i + \sum_{i\in F}\sum_{j\in D} d_j c_{ij}\hat{x}_{ij} \\
&\leq \frac{1}{\alpha}\sum_{i\in F} f_i \hat{y}_i + 3\sum_{j\in D} d_j c_j(\alpha) \\
&\leq \frac{1}{\alpha}\sum_{i\in F} f_i \hat{y}_i + 3\frac{1}{1-\alpha}\sum_{j\in D} d_j \sum_{i\in F} c_{ij}x_{ij} \\
&(\text{choose } \alpha = \frac{3}{1-\alpha} \implies \alpha = \frac{1}{4}) \\
&= 4\sum_{i\in F} f_i y_i + 4\sum_{i\in F}\sum_{j\in D} d_j c_{ij} x_{ij} \\
&= 4\cdot \mathsf{LP}_{\mathsf{OPT}}
\end{aligned}
$$

Therefore we obtained a 4-approximation, which concludes the proof.

■

By choosing α at random we obtain a 3.16 approximation (see the paper for details). The best current algorithm, due to Byrka, obtains a 1.5-approximation. The current best lower bound is 1.463 (so almost tight), assuming that $\mathsf{NP} \subsetneq \mathsf{DTIME}(n^{\lg\lg n})$.

References

[1] Anupam Gupta, Amit Kumar, Tim Roughgarden. Simpler and Better Approximation algorithms for Network Design. In the *Proc. of the 35th Annual ACM Symposium on the Theory of Computing*, 2003.

[2] D.B Shmoys, E. Tardos and K. Aardal. Approximation algorithms for the facility location problem. In the *Proc. of the 29th Annual ACM Symposium on the Theory of Computing*, pages 265-274, 1997.

[3] Lecture notes for Network Design and Game Theory course, held by Mohammad T. Hajiaghayi. Scribed by Mangesh Gupte, Lecture 5, Spring 2008, Rutgers University.

[4] Handwritten notes for Network Design and Game Theory course, held by Mohammad T. Hajiaghayi. Lecture 5, Spring 2008, Rutgers University.