Deterministic Rounding for Facility Location[1]

- In this lecture, we look at LP-rouding algorithms for the (uncapacitated) facility location, aka UFL, problem. Recall, in the UFL problem, we are given a set F of facilities, a set C of clients, and a metric $d(\cdot,\cdot)$ in $F \cup C$. Each facility $i \in F$ has an opening cost f_i. The objective is to open $X \subseteq F$ and connect clients to the nearest open facility via assignment $\sigma : C \to X$ so as to minimize

$$\texttt{cost}(X) = \sum_{i \in X} f_i + \sum_{j \in C} d(\sigma(j), j) \tag{1}$$

- *LP Relaxation.* Here is a natural LP-relaxation for UFL.

$$\begin{aligned}
\mathsf{lp} := \text{minimize} \quad & \sum_{i \in F} f_i y_i + \sum_{i \in F, j \in C} d(i,j) x_{ij} && && \text{(UFL-LP)}\\
& \sum_{i \in F} x_{ij} \geq 1, && \forall j \in C && (2)\\
& y_i - x_{ij} \geq 0, && \forall i \in F,\ \forall j \in C && (3)\\
& x_{ij}, y_i \geq 0, && \forall i \in F, j \in C &&
\end{aligned}$$

Let (x,y) be a fractional solution to the above LP. We define some notation. Let $\mathsf{F_{LP}} := \sum_{i \in F} f_i y_i$. Let $C_j := \sum_{i \in F} d(i,j) x_{ij}$. Let $\mathsf{C_{LP}} = \sum_{j \in C} C_j$. Thus, $\mathsf{lp} = \mathsf{F_{LP}} + \mathsf{C_{LP}}$. We now show a rounding algorithm which returns a solution of cost at most $4(F_{LP} + C_{LP}) \leq 4\texttt{opt}$.

The rounding algorithm proceeds in two stages. The first stage is called *filtering* which will "take care" of the x_{ij}'s, the second stage is called *clustering* which will "take care" of the y_i's.

- *Filtering.* Given a client j, order the facilities in increasing order of $d(i,j)$. That is, $d(1,j) \leq c(2,j) \leq \cdots \leq d(n,j)$ where $n = |F|$. The fractional cost of connecting j to the facilities is C_j; our goal is to make sure that in the final solution, client j doesn't pay "much more" than C_j. To this end, given a parameter $\rho > 1$ (which we will set later) define

$$N_j(\rho) := \{i \in F \ : \ d(i,j) \leq \rho \cdot C_j\} \tag{4}$$

Note it is possible that $x_{ij} > 0$ for some $i \notin N_j(\rho)$. However, we can *massage* the solution (x,y) so that j is fractionally connected only to facilities in $N_j(\rho)$.

Define $(\widehat{x}, \widehat{y})$ as follows. $\widehat{y}_i = \frac{\rho}{\rho-1} \cdot y_i$ for all $i \in F$. For all $j \in C$, set

$$\widehat{x}_{ij} = \begin{cases} \frac{\rho}{\rho-1} \cdot x_{ij} & \text{if } i \in N_j(\rho) \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

Claim 1. $(\widehat{x}, \widehat{y})$ satisfies (2) and (3)

[1] *Lecture notes by Deeparnab Chakrabarty. Last modified : 3rd Jan, 2022*
These have not gone through scrutiny and may contain errors. If you find any, or have any other comments, please email me at deeparnab@dartmouth.edu. Highly appreciated!

Proof. Let us consider the sum $C_j = \sum_i d(i,j)x_{ij} = \sum_{i \in N_j(\rho)} d(i,j)x_{ij} + \sum_{i \notin N_j(\rho)} d(i,j)x_{ij}$. Since $d(i,j) > \rho \cdot C_j$ for all $i \notin N_j(\rho)$, we get $\sum_{i \notin N_j(\rho)} x_{ij} < \frac{1}{\rho}$, for otherwise the second term in the RHS above exceeds C_j. This implies $\sum_{i \in N_j(\rho)} x_{ij} \geq (1 - \frac{1}{\rho})$ implying $\sum_{i \in N_j(\rho)} \widehat{x}_{ij} \geq 1$. $(\widehat{x}, \widehat{y})$ satisfy (3) by definition.

- *Clustering.* In the next step, we *partition* the facilities such that the $\widehat{y}$-mass in each part is at least 1. The rounding algorithm would open the cheapest facility in each part. To find the partition, we proceed iteratively. Initially, all clients are uncovered and comprise the set U. In the beginning of each iteration, we choose the uncovered client $j \in U$ with the smallest C_j. We add this client j to a "representative set" R, and define $F_j := \{i \in F : \widehat{x}_{ij} > 0\}$. That is, F_j is the set of facilities which "serve" client j in the massaged solution $(\widehat{x}, \widehat{y})$. Next, and this is a crucial step, we remove any uncovered client $\ell \in U$ such that $\widehat{x}_{i\ell} > 0$ for any $i \in F_j$. In English, we remove any uncovered client which is fractionally served by *any* facility in F_j in the massaged solution $(\widehat{x}, \widehat{y})$. We continue till the set U becomes empty, that is, all clients are covered. Two key observations follow.

  Claim 2. The sets $\{F_j : j \in R\}$ are pairwise disjoint.

  Proof. Suppose not, and say $i \in F_j \cap F_\ell$ and ℓ entered the set R later. In that case, $\widehat{x}_{i\ell} > 0$ and $i \in F_j$. This is a contradiction as ℓ should have been removed from U in the iteration in which j was added to R.

  Claim 3. For any $j \in R$, $\sum_{i \in F_j} \widehat{y}_i \geq 1$.

  Proof. Once again, the key observation is that $\sum_{i \in F_j} \widehat{x}_{ij} \geq 1$. This follows from Claim 1 because F_j contains all the facilities i such that $\widehat{x}_{ij} > 0$. Otherwise, j would not be in R. And thus, since $\widehat{y}_i \geq \widehat{x}_{ij}$, the claim follows.

- *Algorithm.* The rounding algorithm is now almost complete : open the cheapest facility in each F_j for $J \in R$ and connect each client to the closest open facility.

procedure UFL-ROUNDING($F \cup C$, f_i, $d(i,j)$):
  Solve (UFL-LP) to obtain (x, y).
  Define $N_j(\rho)$ for all $j \in C$ as in (4). ▷ *$\rho = 4/3$ gives the 4-appx*
  Define $\widehat{x}$ as in (5)
  ▷ *Next form the partitions*
  $U \leftarrow C$, $R \leftarrow \emptyset$.
  while $U \neq \emptyset$ **do**:
    Find $j \in U$ with smallest C_j and $R \leftarrow R \cup j$.
    $F_j \leftarrow \{i \in F : \widehat{x}_{ij} > 0\}$.
    Remove all $\ell \in U$ such that $\widehat{x}_{i\ell} > 0$ for any $\ell \in F_j$. ▷ *We let $j \in R$ be **responsible** for these clients.*
  For each $j \in R$, open the facility $i \in F_j$ with smallest f_i.
  Every client connects to nearest facility.

Theorem 1. UFL-ROUNDING returns a 4-approximation for metric UFL when $\rho = 4/3$.

Proof. Let $\mathsf{F}_{\mathsf{alg}}$ and $\mathsf{C}_{\mathsf{alg}}$ be the facility opening and connection costs of the above algorithm. We show that $\mathsf{F}_{\mathsf{alg}} \leq \frac{\rho}{\rho-1}\mathsf{F}_{\mathsf{LP}}$ and $\mathsf{C}_{\mathsf{alg}} \leq 3\rho\mathsf{lp}$. When $\rho = 4/3$, we get $\mathsf{alg} = \mathsf{F}_{\mathsf{alg}} + \mathsf{C}_{\mathsf{alg}} \leq 4\mathsf{lp}$, and this explains the choice of ρ.

Since we open the cheapest facility in $i_j \in F_j$ and since $\sum_{i \in F_j} \widehat{y}_i \geq 1$, we get that $f_{i_j} \leq \sum_{i \in F_j} f_i \widehat{y}_i$. Since the F_j's are pairwise disjoint for $j \in R$, we get $\mathsf{F}_{\mathsf{alg}} \leq \sum_{i \in F} f_i \widehat{y}_i = \frac{\rho}{\rho-1}\mathsf{F}_{\mathsf{LP}}$ by definition of $\widehat{y}$.

Fix a client j. If $j \in R$, then indeed we open a client in $N_j(\rho)$. Therefore, the connection cost that j pays is at most ρC_j. Consider a client $\ell \notin R$. Let $j \in R$ be the representative responsible for ℓ. Firstly, we can assert $C_j \leq C_\ell$ because of Line 8. Let $i_\ell \in F_j$ be the facility such that $\widehat{x}_{i_\ell \ell} > 0$ which removed ℓ from U. Let $i_j \in F_j$ be the facility that is open; i_j may or may not be i_ℓ. Now note that the connection cost of ℓ is at most

$$d(\ell, i_j) \leq d(\ell, j) + d(j, i_j) \leq d(\ell, i_\ell) + d(i_\ell, j) + d(j, i_j)$$

where we have used the metric property of d. Now, $d(\ell, i_\ell) \leq \rho C_\ell$ since $i_\ell \in N_\ell(\rho)$. And the last two terms $d(i_\ell, j) \leq \rho C_j$ and $d(j, i_j) \leq \rho C_j$. And then using the fact that $C_j \leq C_\ell$, we get that the connection cost of client $\ell \leq 3\rho C_\ell$. Altogether, we get $\mathsf{C}_{\mathsf{alg}} \leq 3\rho\mathsf{C}_{\mathsf{LP}}$, proving the theorem. $\square$

Exercise: ☕☕

Recall the k-median problem: in this problem we are given the two sets $F \cup C$ and a metric connection costs $d(\cdot, \cdot)$ over these points. The objective is to open k facilities such that the sum of connection costs of clients to open facilities is minimized. Write a natural LP relaxation for the problem. Describe a rounding algorithm which is allowed to open αk facilities and has total connection cost at most $\beta\mathsf{lp}$, where α, β are some fixed constants (as small as possible).

Notes

The algorithm described here is the first constant factor approximation algorithm for UFL. This can be found in the paper [6] by Shmoys, Tardos, and Aardal. Indeed, the paper describes a better approximation factor of 3.16 which can be obtained by choosing ρ cleverly. The first constant factor approximation algorithm for the k-median problem follows a similar route as above and can be found in the paper [2] by Charikar, Guha, Shmoys, and Tardos. This paper gives a $6\frac{2}{3}$-approximation for the special case when $F = C$. The current best approximation factors for UFL is 1.488 in the paper [5] by Li, and for k-median is 2.625 in the paper [1] by Byrka, Pan, Rybicki, Srinivasan, and Trinh. It is known that unless $P = NP$, the approximation factors can't be below than 1.463 for UFL and 1.735 for k-median. These can be found in the papers [3] and [4], respectively.

References

[1] J. Byrka, T. Pensyl, B. Rybicki, A. Srinivasan, and K. Trinh. An improved approximation for k-median, and positive correlation in budgeted optimization. In *Proc., ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 737–756, 2014.

[2] M. Charikar, S. Guha, D. B. Shmoys, and É. Tardos. A Constant Factor Approximation Algorithm for the k-median Problem. *Proceedings of 31st ACM STOC*, 1999.

[3] S. Guha and S. Khuller. Greedy Strikes Back: Improved Facility Location Algorithms. In *Proc., ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 649–657, 1998.

[4] K. Jain, M. Mahdian, and M. Salavatipour. Packing Steiner Trees. *Proceedings of* 13^{th} *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 266–274, 2003.

[5] S. Li. A 1.488 approximation algorithm for the uncapacitated facility location problem. *Information and Computation*, 222:45–58, 2013.

[6] D. Shmoys, E. Tardos, and K. Aardal. Approximation algorithms for facility location problems. In *Proc., ACM Symposium on Theory of Computing (STOC)*, 1997.