601.435/635 Approximation Algorithms **Lecturer:** Michael Dinitz
Topic: Deterministic Rounding: Metric Uncapacitated Facility Location **Date:** 2/28/19
Scribe: Michael Dinitz

10.1 Introduction

Today we're going to be talking about a new, interesting problem, as well as a more involved algorithmic technique. Last class we analyzed essentially the simplest rounding scheme there is on a particularly simple problem. Today we have a more complicated problem, and a more complicated rounding scheme. Like last class, though, our rounding scheme will be deterministic. This is in contrast to the next couple of weeks, where we'll mostly be concerned with randomized techniques.

10.2 Uncapacitated Metric Facility Location (UFL): Definition

I actually mentioned this problem earlier, when we talked about k-center, but this is the first time we're going to define it formally. There are a ton of variations, but this is the most basic version, so sometimes it is just called *Facility Location*. If we want to distinguish it from some of the more popular variants, we can also call it Uncopacitated Metric Facility Location, or UFL.

Input: Metric Space (V, d), Facility opening costs $\{f_i\}_{i \in V}$

Feasible: Set $S \subseteq V$ of facilities, $S \neq \emptyset$

Objective: $\min_{S \subseteq V} Cost(S) = \sum_{i \in S} f_i + \sum_{j \in V} d(j, S)$, where $d(j, S) = \min_{x \in S} d(j, x)$

In other words, we want to open a set of facilities that minimize the cost of opening the facilities plus the distances from eveyr node to their closest open facility.

10.3 Integer Linear Programing formulation and LP relaxation

Variables:

$$Y_i = \begin{cases} 1 & \text{if } i \in S \\ 0 & \text{o/w} \end{cases}$$

$$X_{ij} = \begin{cases} 1 & \text{if } j \text{ is assigned to } i \\ 0 & \text{o/w} \end{cases}$$

ILP:

$$\begin{aligned}
\text{minimize:} \quad & \sum_{i\in V} Y_i f_i + \sum_{j\in V}\sum_{i\in V} d(i,j)X_{ij} && && (\textbf{UFL-ILP})\\
\text{subject to:} \quad & \sum_{i\in V} X_{ij} = 1 && \forall j\in V && (10.3.1)\\
& X_{ij} \le Y_i && \forall i,j\in V && (10.3.2)\\
& X_{ij} \in \{0,1\} && \forall i,j\in V && (10.3.3)\\
& Y_i \in \{0,1\} && \forall i\in V && (10.3.4)
\end{aligned}$$

The first set of constraints requires every vertex to be assigned to one opened facility, and the second set of constraints say that j can be assigned to i only if i is n opened facility. Clearly this is an exact formulation of UFL.

Now we can relax constraints 10.3.3 and 10.3.4 to get the following Linear Program:

$$\begin{aligned}
\text{minimize:} \quad & \sum_{i\in V} Y_i f_i + \sum_{j\in V}\sum_{i\in V} d(i,j)X_{ij} && && (\textbf{UFL-LP})\\
\text{subject to:} \quad & \sum_{i\in V} X_{ij} = 1 && \forall j\in V\\
& X_{ij} \le Y_i && \forall i,j\in V\\
& 0 \le X_{ij} \le 1 && \forall i,j\in V\\
& 0 \le Y_i \le 1 && \forall i\in V
\end{aligned}$$

Let $F(X,Y)=\sum_{i\in V} Y_i f_i$ be the total facility opening cost and $C(X,Y)=\sum_{j\in V}\sum_{i\in V} d(i,j)X_{ij}$ be the total connecting costs. We will also let $Z(X,Y)=F(X,Y)+C(X,Y)$ be the total cost of the (fractional) solution (X,Y). Now it is a polynomial size LP, so it can be solved in polynomial time such that:

$$OPT(LP) \le OPT(ILP) = OPT.$$

10.4 LP rounding

Theorem 10.4.1 [STA97] *Given feasible fractional solution (X,Y), there is an integer feasible solution $(\widehat{X},\widehat{Y})$ with $Z(\widehat{X},\widehat{Y}) \le 4\cdot Z(X,Y)$.*

Although Theorem (10.4.1) claims a 4-approximation, we will begin with a 6-approximation which is a little more intuitive. A proof of 4-approximation then can be easily constructed based on the idea of the proof of 6-approximation. This algorithm is split into two stages: filtering and rounding (although the filtering stage is more of a thought-experiment than an actual algorithmic step)

10.4.1 Stage 1: Filtering

The ideas behind filtering are due to Lin and Vitter [LV92]. Based on the fractional solution (X, Y) provided by LP, let's define *"fractional connection cost"* for node j as follows:

$$\Delta_j = \sum_{i \in V} d(i,j) X_{ij}.$$

Since for any $j \in V$, the values $\{X_{ij}\}_{i \in V}$ are non-negative and sum to 1 (constraint 10.3.1), we can think of them as a probability distribution over $i \in V$, so Δ_j is essentially the *expected* connection cost when the facility j connects to is drawn from this distribution. Such a view will help us later when we use Markov's inequality. Now let's define the ball B_j around node j as follows:

$$B_j = \{i \in V : d(i,j) \leq 2\Delta_j\}$$

Lemma 10.4.2 *Given fractional solution (X, Y), we can find another fractional solution (X', Y') such that:*

1. *$Z(X', Y') \leq 2Z(X, Y)$, and*
2. *If $X'_{ij} > 0$, then $i \in B_j$ (and hence $d(i,j) \leq 2\Delta_j$).*

Proof: Let j be an arbitrary node. We first claim that most of the X-value for j lies inside B_j. This is straightforward from the probabilistic interpretation and Markov's inequality, but we prove it here for completeness.

Claim 10.4.3 $\displaystyle\sum_{i \notin B_j} X_{ij} \leq \frac{1}{2}$

Proof: Suppose $\displaystyle\sum_{i \notin B_j} X_{ij} > \frac{1}{2}$. We prove the claim by way of contradiction as follows:

$$\begin{aligned} \Delta_j = \sum_{i \in V} d(i,j) X_{ij} &\geq \sum_{i \notin B_j} d(i,j) X_{ij} \\ &\geq \sum_{i \notin B_j} 2\Delta_j X_{ij} \\ &= 2\Delta_j \sum_{i \notin B_j} X_{ij} \\ &> \Delta_j \end{aligned}$$

This is clearly a contradiction, and hence $\sum_{i \notin B_j} X_{ij} \leq \frac{1}{2}$ as claimed. ∎

Now we can define new fractional variables X'_{ij} and Y'_i as follows:

$$\begin{aligned} X'_{ij} &= \begin{cases} 0 & \text{if } i \notin B_j \\ \frac{X_{ij}}{\sum_{i\in B_j} X_{ij}} & \text{if } i \in B_j \end{cases} \\ Y'_i &= \min\{1, 2Y_i\} \end{aligned}$$

Claim 10.4.4 *(X', Y') is a feasible solution to the LP.*

Proof: Clearly both the X'_{ij}'s and the Y'_i's are in the interval $[0, 1]$. It is also true by construction that for any $j \in V$, $\sum_{i\in V} X'_{ij} = 1$. So we simply need to prove that $X'_{ij} \le Y'_i$ for all $i, j \in V$. This is clearly true if $Y'_i = 1$, so without loss of generality assume that $Y'_i = 2Y_i$. Then

$$X'_{ij} = \frac{X_{ij}}{\sum_{i\in B_j} X_{ij}} \le \frac{Y_i}{1/2} = 2Y_i = Y'_i,$$

where we used Claim 10.4.3. ■

To finish the proof of Lemma 10.4.2, note that the second condition of the lemma is satisfied by construction. So we just need to prove that $Z(X', Y') \le 2Z(X, Y)$. To do this, note that by Claim 10.4.3 we know that $X'_{ij} \le 2X_{ij}$. Hence

$$\begin{aligned} Z(X', Y') &= \sum_i f_i Y'_i + \sum_j \sum_i d(i,j) X'_{ij} \\ &\le \sum_i 2 f_i Y_i + \sum_j \sum_i 2d(i,j) X_{i,j} \\ &= 2Z(X, Y) \end{aligned}$$

■

10.4.2 Stage 2: Rounding

We can now do the rounding: this is given in Algorithm 1. Note that this rounding starts with the LP solution (X, Y), not the filtered solution (X', Y'). The filtered solution appears in the analysis.

We first give a bound on the facility opening costs.

Lemma 10.4.5 $F(\widehat{X}, \widehat{Y}) \le F(X', Y') \le 2F(X, Y)$

Proof: We have already proved RHS. We only need to show LHS holds. We have:

$$\begin{aligned} F(\widehat{X}, \widehat{Y}) &= \sum_{\substack{j \text{ considered} \\ \text{by Alg}}} f_{a(j)} \\ &\le \sum_{\substack{j \text{ considered} \\ \text{by Alg}}} \sum_{i\in B_j} f_i Y'_i \\ &\le \sum_{i\in V} f_i Y'_i \\ &= F(X', Y') \end{aligned}$$

Algorithm 1 Rounding Algorithm for UFL

Initially all nodes are unassigned
  while there exists unassigned nodes **do**
    let j be unassigned node with minimum Δ_j
    open facility $a(j) \in B_j$ with smallest opening cost
    assign j to a(j)
    for any j' unassigned with $B_j \cap B_{j'} \neq \emptyset$ **do**
      assign j' to a(j)
    end for
  end while
  call this $(\widehat{X}, \widehat{Y})$ and facilities opened $\widehat{S}$

The second inequality is true because clearly for any two nodes j, j' considered by the algorithm, $B_j, B_{j'}$ are disjoint. The first inequality is true because

$$\sum_{i \in B_j} f_i Y_i' \geq \sum_{i \in B_j} f_{a(j)} Y_i' \geq f_{a(j)},$$

where we used the fact that $a(j)$ has the smallest opening cost of any node in B_j. ■

We can now begin to bound the connection costs.

Lemma 10.4.6 $d(j, \widehat{S}) \leq 3 \cdot \mathrm{Rad}(B_j) = 6\Delta_j$ *for all* $j \in V$.

Proof: We divide into cases depending on whether j was considered by the algorithm (i.e. a facility was opened up because of j) or whether it was assigned in the for loop of the algorithm.

Case 1: j considered by algorithm 1. Then a facility was opened up within B_j, and hence $d(j, \widehat{S}) \leq \mathrm{Rad}(B_j) = 2\Delta_j$.

Case 2: j not considered by algorithm 1. Then there exists j' considered by algorithm 1 such that $\Delta_{j'} \leq \Delta_j$ and j assigned to $a(j')$ and $B_j \cap B_{j'} \neq \emptyset$. Let $i' \in B_j \cap B_{j'}$. Then

$$\begin{aligned} d(j, \widehat{S}) &\leq d(j, a(j')) \\ &\leq d(j, i') + d(i', j') + d(j', a(j')) \\ &\leq \mathrm{Rad}(B_j) + 2\,\mathrm{Rad}(B_{j'}) \\ &\leq 3\,\mathrm{Rad}(B_j) = 6\Delta_j \end{aligned}$$

■

Using this lemma, we can easily bound the total connection costs.

Lemma 10.4.7 $C(\widehat{X}, \widehat{Y}) \leq 6 \cdot C(X, Y)$.

Proof:

$$C(\widehat{X}, \widehat{Y}) \;=\; \sum_j d(j, \widehat{S}) \leq \sum_j 6\Delta_j = 6 \sum_j \Delta_j = 6 \cdot C(X, Y)$$

■

Putting this all together, we get a 6-approximation:

$$\begin{aligned} Z(\widehat{X}, \widehat{Y}) &= F(\widehat{X}, \widehat{Y}) + C(\widehat{X}, \widehat{Y}) \\ &\leq 2F(X, Y) + 6C(X, Y) \\ &= 6Z(X, Y). \end{aligned}$$

To improve this to a 4-approximation, first note that the above bound is weak in the sense that it gives a factor 2 loss in the facility opening costs but a factor 6 loss in the connection costs. It turns out that we can balance these out more evenly, so we lose a factor of 4 on both. To do this, we can simply redo the whole analysis with B_j redefined to be $B_j = \{i \in V : d(i, j) \leq \frac{4}{3}\Delta_j\}$. This gives a 4-approximation.

References

[LV92] Jyh-Han Lin and Jeffrey Scott Vitter. epsilon-approximations with minimum packing constraint violation (extended abstract). In *Proceedings of the 24th Annual ACM Symposium on Theory of Computing (STOC)*, pages 771–782. ACM, 1992.

[STA97] David B. Shmoys, Éva Tardos, and Karen Aardal. Approximation algorithms for facility location problems (extended abstract). In *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing (STOC)*, pages 265–274. ACM, 1997.