CMPUT 675: Approximation Algorithms **Fall 2014**

Lecture 15 (Oct 6): LP Duality

Lecturer: Zachary Friggstad *Scribe: Zachary Friggstad*

15.1 Introduction by Example

Given a linear program and a feasible solution $\mathbf{x}'$, it is easy to convince someone that $\mathbf{x}'$ is not optimal by simply providing an even better solution. However, if $\mathbf{x}'$ is optimal then is there any easy way to be convinced of this fact? Is there a short proof of the fact that no other feasible solution is better?

We saw earlier that there is always an optimal solution that is also an extreme point, so we could enumerate all possible extreme points and see if any one of them is better than $\mathbf{x}'$. However, this "proof" is very long as there could be exponentially many extreme points.

There is a much nicer way. Consider the following example.

$$\begin{array}{rrcl}
\text{minimize :} & 3x_1 + x_2 + 4x_3 & & \\
\text{subject to :} & x_1 + 2x_2 & \geq & 3 \\
& x_1 + 2x_3 & \geq & 2 \\
& 2x_1 + 3x_2 + x_3 & \geq & 4 \\
& x_1, x_2, x_3 & \geq & 0
\end{array} \qquad \textbf{(LP-Example)}$$

It is easy to check that the solution $\mathbf{x}' = (0, 3/2, 1)$ is feasible and has value $11/2$. Is there any better solution?

We can generate lower bounds by scaling some of the constraints and adding them together. For example, $\frac{1}{2}$(constraint # 2) + $\frac{1}{3}$(constraint # 3) is

$$\frac{7}{6}x_1 + x_2 + \frac{4}{3}x_3 \geq \frac{7}{3}.$$

Since we only took nonnegative multiples of the constraints (so we did not flip the inequality sign), this constraint must be satisfied by any feasible solution to (**LP-Example**). Note that the coefficient of each x_i in this constraint is at most its coefficient in the objective function. Since $x_i \geq 0$ is also a constraint, this means any feasible solution $\mathbf{x}''$ must satisfy

$$3x''_1 + x''_2 + 4x''_3 \geq \frac{3}{2}x''_1 + x''_2 + 2x''_3 \geq \frac{7}{3}.$$

That is, the optimum solution value to this linear program is at least $\frac{7}{3}$.

Is there an even better lower bound we can generate in this manner? Yes, in fact this process can be automated: we can model the problem of finding the best possible lower bound that can be obtained this way as another linear program that we call the *dual* of the original linear program.

Let y_1, y_2, y_3 denote the coefficients of the linear combination. The following linear program finds the best lower bound on the optimum solution value to (**LP-Example**).

$$\begin{array}{rrcl} \text{maximize :} & 3y_1 + 2y_2 + 4y_3 & & \\ \text{subject to :} & y_1 + y_2 + 2y_3 & \leq & 3 \\ & 2y_1 + 3y_3 & \leq & 1 \\ & 2y_2 + y_3 & \leq & 4 \\ & y_1, y_2, y_3 & \geq & 0 \end{array}$$

To recap, for any feasible solution $\mathbf{y}'$, the objective function value of $\mathbf{y}'$ is a lower bound on the optimum solution to (**LP-Example**). The constraints of this new linear program ensure that the coefficients of each x_i in $\sum_{j=1}^{3} y_j \cdot$ (constraint # i of (**LP-Example**)) are at most their corresponding coefficient in the objective function of (**LP-Example**). Finally, $y_i \geq 0$ ensures we only take a nonnegative linear combination of the constraints of (**LP-Example**).

The solution $\mathbf{y}' = (1/2, 2, 0)$ is feasible for this linear program and has value 11/2. That is,

$$\frac{1}{2} \cdot (\text{constraint } \# \ 1) + 2 \cdot (\text{constraint } \# \ 2) = \frac{5}{2}x_1 + x_2 + 4x_3 \geq \frac{11}{2}.$$

Thus, for any feasible solution $\mathbf{x}''$ to (**LP-Example**) we have

$$3x_1'' + 2x_2'' + 4x_3'' \geq \frac{5}{2}x_1'' + x_2'' + 4x_3'' \geq \frac{11}{2}$$

so the solution $\mathbf{x}'$ above is in fact optimal.

15.1.1 The General Recipe

In the previous section, we saw an example of how to prove optimality of a proposed LP solution by considering some appropriate linear combination of the constraints. We saw that finding the best lower bound that can be obtained using this method involved solving another LP that we called the dual LP.

Here is the general recipe. Consider a linear program in standard form.

$$\begin{array}{rrcl} \text{minimize :} & \mathbf{c}^T \cdot \mathbf{x} & & \\ \text{subject to :} & \mathbf{A} \cdot \mathbf{x} & \geq & \mathbf{b} \\ & \mathbf{x} & \geq & 0 \end{array} \qquad (\textbf{Primal})$$

For each of the i constraints of the form $\mathbf{A} \cdot \mathbf{x} \geq \mathbf{b}$, let y_i be a *dual variable* for that constraint. Recall that we want to compute the best lower bound possible by considering a nonnegative linear combination of the constraints. This is expressed by the following linear program.

$$\begin{array}{rrcl} \text{maximize :} & \mathbf{b}^T \cdot \mathbf{y} & & \\ \text{subject to :} & \mathbf{A}^T \cdot \mathbf{x} & \leq & \mathbf{c} \\ & \mathbf{y} & \geq 0 & \end{array} \qquad (\textbf{Dual})$$

Theorem 1 (Weak Duality) *Let $\overline{\mathbf{x}}$ be a feasible solution to (***Primal***) and $\overline{\mathbf{y}}$ be a feasible solution to (***Dual***). Then $\mathbf{c}^T \cdot \overline{\mathbf{x}} \geq \mathbf{b}^T \cdot \overline{\mathbf{y}}$.*

Proof.

$$
\begin{aligned}
\mathbf{c}^T \cdot \overline{\mathbf{x}} &= \sum_j c_j \cdot \overline{x}_j \\
&\geq \sum_j \left(\sum_i A_{i,j} \cdot \overline{y}_i \right) \cdot \overline{x}_j && (15.1) \\
&= \sum_i \left(\sum_j A_{i,j} \cdot \overline{x}_j \right) \cdot \overline{y}_i \\
&\geq \sum_i b_i \cdot \overline{y}_i && (15.2) \\
&= \mathbf{b}^T \cdot \overline{\mathbf{y}}
\end{aligned}
$$

The inequality (15.1) is justified because $\mathbf{x} \geq 0$ and $\mathbf{A^T} \cdot \mathbf{y} \leq \mathbf{c}$. The inequality (15.2) is justified because $\mathbf{y} \geq 0$ and $\mathbf{A} \cdot \mathbf{x} \geq \mathbf{b}$. ■

So, what is the best lower bound on the optimum primal solution we can prove using this approach? We can, in fact, always prove an optimum lower bound!

Theorem 2 (Strong Duality) *The primal LP has an optimum solution if and only if the dual LP has an optimum solution. In this case, their optimum solutions have the same value.*

See [KV12] for a proof.

Sometimes, our original (primal) LP is a maximization LP. Using this same approach (i.e. find the best upper bound by some appropriate combination of the constraints) leads to essentially the same results. If the primal LP is $\max\{\mathbf{c}^T \cdot \mathbf{x} : \mathbf{A} \cdot \mathbf{x} \leq \mathbf{b}, \mathbf{x} \geq 0\}$ then the dual is $\min\{\mathbf{b}^T \cdot \mathbf{y} : \mathbf{A}^T \cdot \mathbf{y} \geq \mathbf{c}, \mathbf{y} \geq 0\}$.

Theorems 1 and 2 hold in this setting in the sense that the value of any dual solution provides an upper bound on the value of any primal solution and that their optimum solutions have equal value. Notice that taking the dual of this dual LP then results in the primal again, so LPs come in pairs: the primal LP and its dual.

Note that the dual of the dual of an LP is the original LP itself: **LPs come in pairs**.

Question: What if some of the constraints in the primal LP are equality constraints?
Answer: If the i'th constraint of the primal is an equality constraint, then do not add the nonnegativity constraint $y_i \geq 0$ for the corresponding dual variable y_i.

Question: What if some of the primal variables are not restricted to be nonnegative?
Answer: If $x_j \geq 0$ is not present in the constraints, then make the j'th constraint of the dual an equality constraint.

Question: What if we have a $\leq$ constraint (e.g. $x_1 - 3x_2 \leq 4$) for a primal "min" LP or a $\geq$ constraint for a primal "max" LP?
Answer: Just negate it to switch the direction of the inequality before constructing the dual.

The rows of the following table summarize the rules for constructing the dual of an LP. To use it, identify the column that corresponds to the type of primal LP you start with (i.e. whether it is a min or max LP). For each constraint of the primal, check to see if it is an inequality or equality constraint to see if the corresponding dual variable should be nonnegative or unconstrained. For each variable of the primal, check to see if it is constrained to be nonnegative or not to determine if the corresponding dual constraint should be an inequality or equality constraint.

minimize	maximize
$\sum_j A_{i,j}x_i \geq b_i$	$y_i \geq 0$
$\sum_j A_{i,j}x_i = b_i$	y_i unconstrained
$x_i \geq 0$	$\sum_i A_{i,j}y_i \leq c_j$
x_i unrestricted	$\sum_i A_{i,j}y_i = c_j$

Question: Do we still have weak and strong duality for these more general LPs?
Answer: Yes, Theorems 1 and 2 hold in this more general setting. The proof of Theorem 1 is pretty much identical, except we have to say things like "because either $\overline{y}_i \geq 0$ **or** constraint i is an equality constraint" to justify the inequalities.

15.1.2 Application 1: Set Cover

Recall the Set Cover problem. We are given a finite set X and a collection $\mathcal{S}$ of subsets of X. Each $S \in \mathcal{S}$ has a cost $c(S) \geq 0$. The goal is to find the cheapest subcollection $\mathcal{C} \subseteq \mathcal{S}$ that covers X (i.e. $\cup_{S \in \mathcal{C}} S = X$).

Recall in our analysis of the greedy algorithm that we constructed a set $\mathcal{C}$ and values $\bar{z}_i$ for each item $i \in I$ with the following properties:

1. $\bar{z}_i \geq 0$ for each $i \in X$
2. $\sum_{i \in S} \bar{z}_i \leq H_{|S|} \cdot c(S)$ for each $S \in \mathcal{S}$ where $H_m = \sum_{a=1}^{m} \frac{1}{a} = \ln m + O(1)$ is the m'th harmonic number
3. $\sum_{i \in X} \bar{z}_i = \text{cost}(\mathcal{C})$

Setting $k = \max_{S \in \mathcal{S}} |S|$, we used these values to show that $\text{cost}(\mathcal{C}) \leq H_k \cdot OPT$.

Now consider the following LP relaxation for Set Cover:

$$\begin{array}{rrcll} \text{minimize}: & \sum_{S \in \mathcal{S}} c(S) \cdot x_S & & & \\ \text{subject to}: & \sum_{S : i \in S} x_S & \geq & 1 & \text{for each item } i \in X \\ & \mathbf{x} & \geq & 0 & \end{array} \qquad \textbf{(LP-SC)}$$

To construct the dual, note that the constraints of the primal correspond naturally to the items in X, so we will use a dual variable z_i for each $i \in X$. The $\mathbf{b}$ vector in the relaxation contains all 1s, so the objective of the dual is simply to minimize $\sum_i z_i$. Since we have a variable for each $S \in \mathcal{S}$, the dual will have a constraint for each $S \in \mathcal{S}$ and the right-hand side of this constraint will be $c(S)$.

Usually the most tricky part about constructing the dual is coming up with a nice way to express its constraint matrix. Note that in the constraint matrix A for the LP relaxation, each column corresponds to a subset S and the 1 entries in that column are in the rows/items i contained in S. So, the S'th row of the dual constraint matrix contains a 1 for each variable z_i where $i \in S$.

The dual LP is:

$$\begin{array}{rrcll} \text{maximize}: & \sum_{i \in X} z_i & & & \\ \text{subject to}: & \sum_{i \in S} z_i & \leq & c(S) & \text{for each subset } S \in \mathcal{S} \\ & \mathbf{z} & \geq & 0 & \end{array}$$

Now, the properties of the specific values $\bar{z}$ listed above show that $\bar{z}$ is almost a dual solution. In fact, the vector $\bar{\mathbf{z}}/H_k$ is in fact a feasible dual solution by properties 1 and 2 for $\bar{\mathbf{z}}$. Finally, let $\mathbf{x}^*$ be an optimum primal solution. By weak duality, we have

$$OPT_{LP} = \sum_S c(S) \cdot x_S^* \geq \frac{1}{H_k} \sum_i \bar{z}_i = \frac{1}{H_k} \cdot \text{cost}(\mathcal{C})$$

where the last equality is from property 3.

Theorem 3 *The integrality gap of the LP relaxation (***LP-SC***) for* Set Cover *is at most H_k where k is the size of the largest set $S \in \mathcal{S}$.*

Note that we proved this without rounding an optimal LP solution. LP Duality was the main tool.

15.1.3 Application 2: Generalized Max-Flow/Min-Cut Theorems

Recall the Multicut LP relaxation from Lecture 13.

$$\begin{array}{rrcll} \text{minimize :} & \displaystyle\sum_{e \in E} c_e \cdot x_e & & & \\ \text{subject to :} & \displaystyle\sum_{e \in P} x_e & \geq & 1 & \text{for each } 1 \leq i \leq k \text{ and each path } P \in \mathcal{P}_i \\ & \mathbf{x} & \geq & 0 & \end{array}$$

The dual of this LP has a variable y_P^i for each $1 \leq i \leq k$ and each $P \in \mathcal{P}_i$. It is

$$\begin{array}{rrcll} \text{maximize :} & \displaystyle\sum_{i=1}^{k} \sum_{P \in \mathcal{P}_i} y_P^i & & & \\ \text{subject to :} & \displaystyle\sum_{i,P : e \in P} y_P^i & \leq & c_e & \text{for each edge } e \in E \\ & \mathbf{y} & \geq & 0 & \end{array}$$

The dual LP solves what is called the Maximum Multicommodity Flow problem. Recall that an $s-t$ flow in an undirected graph consists of bidirecting each edge $e \in E$ (i.e. for $e = (u,v) \in E$, replacing e with the two directed edges (u,v) and (v,u)) and assigning a "flow" $0 \leq f_e$ to each directed edge so that the total flow entering some node $v \neq s,t$ equals the total flow exiting v. Furthermore, if e', e'' are the two directed copies of e then $f_{e'} + f_{e''} \leq f_e$. The value of this flow is the net flow exiting s, namely $\sum_{e \text{ exiting } s} f_e - \sum_{e \text{ entering } s} f_e$ (this also equals the net flow entering t).

A *multicommodity flow* for pairs $(s_1, t_1), \ldots, (s_k, t_k)$ consists of a flow f^i for each $1 \leq i \leq k$ such that for any edge e, the total flow sent by all pairs across e is at most c_e. The Maximum Multicommodity Flow problem is to find a multicommodity flow with maximum total value over all pairs.

There is a natural correspondence between the multicommodity flows and feasible solutions to the dual LP above. Namely, every feasible dual LP solution gives rise to a multicommodity flow with the same value. This follows through the correspondence of flows and their path decompositions (the details are omitted here). This is also why we can rewrite the primal LP for Minimizing Congestion using only polynomially many variables (i.e. using flows of value 1 instead of assigning weights to paths).

The well-known max-flow/min-cut theorem states that for any edge-capacitated graph $G = (V, E)$ and any two $s, t \in V$ that the value of a maximum $s-t$ flow equals the minimum capacity $s-t$ cut. We proved the integrality gap of the primal LP is at most $2\ln(k+1)$. Using LP duality, this can be interpreted as sort of a generalization of the max-flow/min-cut theorem for undirected graphs to multicommodity flows and multicuts.

Theorem 4 (Multicommodity Flow/Cut Gap) *Let f^* denote the value of a maximum multicommodity flow and C^* denote the cheapest multicut. Then $f^* \le C^* \le 2\ln(k+1) \cdot f^*$.*

15.2 Complementary Slackness

Suppose $\mathbf{x}^*$ and $\mathbf{y}^*$ are optimal solutions to the primal and dual. By strong duality, $\mathbf{c}^T \cdot \mathbf{x} = \mathbf{b}^T \cdot \mathbf{y}^*$ so both inequalities (15.1) and (15.2) in the proof of weak duality for these solutions must in fact hold with equality.

In the proof of weak duality, the inequalities held term-by-term. That is, the (15.1) was true because for **every** j, either

a) $c_j \ge \sum_i A_{i,j} \cdot y^*_i$ and x^*_j

or

b) The j'th constraint in the dual is an equality constraint.

So, if the inequality (15.1) to hold with equality it must be the case that either $x^*_j = 0$ or $c_j = \sum_i A_{i,j} \cdot y^*_i$ for each j. Similarly, if inequality (15.2) is to hold with equality it must be the case that either $y^*_i = 0$ or $b_i = \sum_j A_{i,j} \cdot x^*_j$ for each i. These are the so-called *complementary slackness* conditions.

The converse is true as well, if $\overline{\mathbf{x}}$ and $\overline{\mathbf{y}}$ are feasible primal and dual solutions and if these complementary slackness conditions hold, then following the proof of weak duality we see that inequalities (15.1) and (15.2) hold with equality so in fact $\mathbf{c}^T \cdot \overline{\mathbf{x}} = \mathbf{b}^T \cdot \overline{\mathbf{y}}$ so both $\overline{\mathbf{x}}$ and $\overline{\mathbf{y}}$ are optimal solutions for their respective LP.

Summarizing:

Theorem 5 (Complementary Slackness) *Let $\overline{\mathbf{x}}$ and $\overline{\mathbf{y}}$ be feasible primal and dual LP solutions. Then both $\overline{\mathbf{x}}$ and $\overline{\mathbf{y}}$ are optimal for their respective linear program if and only if the following two conditions hold:*

- *For every i, either $\overline{y}_i = 0$ or $\sum_j A_{i,j} \cdot \overline{x}_j = c_i$.*
- *For every j, either $\overline{x}_j = 0$ or $\sum_i A_{i,j} \cdot \overline{y}_i = b_i$.*

Note: If, say, only the first condition holds then we cannot conclude that at least one of $\overline{\mathbf{x}}$ or $\overline{\mathbf{y}}$ is optimal for its respective LP. They are all-or-nothing conditions.

15.3 The Uncapacitated Facility Location Problem

We finally get to our next facility location problem known as the UNCAPACITATED FACILIITY LOCATION problem. Here, we are given a set of clients C, a set of potential facility locations F, and metric costs $c(i,j)$ between these clients and facilities. Furthermore, each potential facility $i \in F$ has an *opening cost* f_i. The goal is to open some facilities and assign each client to an open facility to minimize the total opening and assignment cost. That is, we should find some nonempty $S \subseteq F$ to minimize:

$$\sum_{i \in S} f_i + \sum_{j \in C} d(j, S).$$

The term "uncapacitated" in the title means each facility can serve an unbounded number of clients. More complicated (and interesting) variants place restrictions on how many clients each facility can handle but we will not consider them here.

Consider the following LP relaxation. We have variables y_i indicating whether facility i is open or not and variables $x_{i,j}$ indicating whether client j is assigned to facility i. The relaxation is as follows:

$$
\begin{array}{lrcll}
\text{minimize :} & \sum_{i \in F} f_i \cdot y_i + \sum_{i \in F, j \in C} c(i,j) \cdot x_{i,j} & & & \\
\text{subject to :} & \sum_{i \in F} x_{i,j} & = & 1 & \text{for each } j \in C \\
 & y_i - x_{i,j} & \geq & 0 & \text{for each } i \in F, j \in C \\
 & \mathbf{x}, \mathbf{y} & \geq & 0 &
\end{array} \qquad \textbf{(FL-Primal)}
$$

The first constraints ensure every client is assigned to some facility and the second constraints ensure that a client can only be assigned to an open facility.

This LP is not in standard form as it involves equality constraints. However, the recipe for constructing duals mentioned earlier says this simply means the corresponding dual variable is not constrained to be nonnegative.

We use dual variables α_j for each constraint of the first type and $\beta_{i,j}$ for each constraint of the second type. The dual is then:

$$
\begin{array}{lrcll}
\text{maximize :} & \sum_{j \in C} \alpha_j & & & \\
\text{subject to :} & \alpha_j - \beta_{i,j} & \leq & c(i,j) & \text{for each } i \in F, j \in C \\
 & \sum_{j \in C} \beta_{i,j} & \leq & f_i & \text{for each } i \in F \\
 & \beta & \geq & 0 &
\end{array} \qquad \textbf{(FL-Dual)}
$$

The first set of constraints in the dual correspond to primal variables $x_{i,j}$ and the second set of constraints correspond to primal variables y_i. Nonnegativity of the α variables is ommited because the corresponding primal constraints are equality constraints.

The approximation we see here requires us to solve both the primal and the dual. The dual solution will help guide the execution of the algorithm to ensure we find a relatively cheap solution. The following lemma helps illustrate why this is.

From now on, let $\mathbf{x}^*, \mathbf{y}^*$ be an optimal primal solution and α^*, β^* be an optimal dual solution. Let $F_j = \{i \in F : x_{i,j} > 0\}$ be the set of facilities that partially serve i. Finally, let $i(j)$ denote the facility in F_j with the cheapest opening cost. That is, $f_{i(j)} = \min_{i \in F_j} f_i$.

Lemma 1 *For any $i \in F_j$, $c(i,j) \leq \alpha_j$.*

Proof. This crucially relies on the fact that both the primal and dual solutions are optimal. For $i \in F$ we have, by definition, $x^*_{i,j} > 0$. By complementary slackness (i.e. Theorem 5), the corresponding dual constraint holds with equality: $\alpha_j - \beta_{i,j} = c(i,j)$. Since $\beta_{i,j} \geq 0$ is also a constraint of the dual, we see in fact that $\alpha_j \geq c(i,j)$. ■

So, if we were to ensure that each client gets assigned to some facility in F_j then we know the total assignment cost is at most $\sum_{j \in C} \alpha^*_j = OPT_{\text{LP}}$. However, this does not account for the facility opening cost. It may be too expensive to open a facility in each F_j. However, if we open the cheapest facility in a collection of disjoint facilities sets F_j then the opening cost is bounded.

Lemma 2 *Suppose $C' \subseteq C$ is such that $F_j \cap F_{j'} = \emptyset$ for any distinct $j, j' \in C'$. Then $\sum_{j \in C'} f_{i(j)} \le OPT_{LP}$.*

Proof. First, for any $j \in C$ the constraints of LP (**FL-Primal**) show.

$$\begin{aligned} f_{i(j)} &= \sum_{i \in F_j} f_{i(j)} \cdot x^*_{i,j} \\ &\le \sum_{i \in F_j} f_i \cdot x^*_{i,j} \\ &\le \sum_{i \in F_j} f_i \cdot y^*_i \end{aligned}$$

Finally, since $F_j \cap F_{j'} = \emptyset$ for any $j, j' \in \mathcal{C}'$ then

$$\sum_{j \in C'} f_{i(j)} \le \sum_{j \in C'} \sum_{i \in F_j} f_i \cdot y^*_i \le \sum_{i \in F} f_i \cdot y^*_i \le OPT_{\text{LP}}.$$

■

Algorithm 1 An LP-based UNCAPACITATED FACILITY LOCATION approximation

```
x*, y*, α*, β* ← optimal primal and dual solutions
C' ← ∅
F_j ← {i ∈ F : x*_{i,j} > 0} for each j ∈ C
i(j) ← arg min_{i∈F_j} f_i for each j ∈ C
for each j ∈ C in increasing order of α*_j do
  if there is some j' ∈ C' such that F_j ∩ F_j' ≠ ∅ then
    assign j to i(j')
  else
    open facility i(j) and assign j to i(j)
    C' ← C' ∪ {j}
  end if
end for
```

Theorem 6 *The solution found by Algorithm 1 has cost at most $4 \cdot OPT_{LP}$.*

Proof. By construction, the only open facilities are those of the form $i(j)$ for some $j \in C'$. By construction, $F_j \cap F_{j'} = \emptyset$ for every two $j, j' \in C'$ so the total opening cost is at most OPT_{LP} by Lemma 2.

Now for the connection costs. For $j \in C'$, we assigned j to $i(j) \in F_j$ and $c(j, i(j)) \le \alpha^*_j$ by Lemma 1. If $j \notin C'$, then we connected j to some $i(j')$ for some j' that was in C' during the iteration for client j where $F_j \cap F_{j'} \neq \emptyset$. Let i be some facility in $F_j \cap F_{j'}$. Finally, we have $\alpha_{j'} \le \alpha_j$ because j' was considered before j in the loop.

By the triangle inequality and by using Lemma 1 again to bound both $c(i(j'), j')$ and $c(i, j')$ (noting that $i(j'), i \in F_{j'}$) we have

$$c(i(j'), j) \le c(i(j'), j') + c(i, j') + c(i, j) \le \alpha_{j'} + \alpha_{j'} + \alpha_j \le 3 \cdot \alpha_j.$$

Overall, the total connection cost is at most $\sum_{j \in C} 3\alpha_j = 3 \cdot OPT_{\text{LP}}$. Therefore, the total connection and facility opening cost is at most $4 \cdot OPT_{\text{LP}}$. ■

The UNCAPACITATED FACILITY LOCATION problem is fairly well understood, the gap between the known upper and lower bounds is quite small. Namely, there is a 1.488-approximation [L11] and there is no 1.463-approximation unless NP $\subseteq$ DTIME($n^{O(\log \log n)}$) [GK98].

References

- GK98 S. Guha and S. Khuller, Greedy strikes back: Improved facility location algorithms, Journal of Algorithms, 649–657, 1998.
- KV12 B. Korte and J. Vygen, *Combinatorial Optimization: Theory and Algorithms*, Springer, 2012.
- L11 S. Li, A 1.488-approximation algorithm for the uncapacitated facility location problem, in Proceedings of ICALP, 2011.