CS684: Approximation and Network Algorithms	**10/26/01**
Primal-Dual Approximation for Metric Facility Location	**scribe: Yasser Mourtada**

1 Problem Description

The metric facility location problem can be defined as follows:

- A set of clients $C = \{1, .., n\}$
- A set of possible facilities $F = \{1, .., m\}$
- For every $i, j \in \{F \times C\}$, a $dist(i, j)$ function which specifies the metric distance between facilities and clients, and satisfies the following:
  $dist(i, j) > 0, \forall i, j \in \{F \times C\}$
  $dist(i, j) = dist(j, i), \forall i, j \in \{F \times C\}$
  $dist(i, j) + dist(j, k) \geq dist(i, k)$ (Triangle Inequality)
- $\forall i \in F$, a cost f_i denoting the cost of opening facility i.

The goal of the metric facility location problem is to assign every client to a facility by opening a subset $S \subseteq F$, as to minimize a cost function given by:

$$\sum_{i \in S} f_i + \sum_{j \in C} dist(j, i_j)$$

where i_j is the facility assigned to client j.

Other versions of the facility location problem require each client to be assigned to the closest facility in S, but we will relax this requirement in our case as long as each client is connected to a facility.
The facility location problem is known to be NP-complete. In this lecture, we will present a primal-dual based 3-approximation algorithm for it.

2 Primal-Dual Algorithm

The following primal-dual algorithm is due to Jain and Varizani. It consists of two phases: the first assigns every client to at least one facility, and the second adjusts the solution so that each client is assigned to exactly one facility. Before we present the algorithm however, we define rules for a feasible solution based on an economic interpretation of the problem:
Let p_j denote the total price paid by client $j \in C$.

1. Every client $j \in C$ is assigned to some open facility $i \in S$, and $p_j \geq 0 \forall j \in C$.

2. If client j is assigned to facility i, then $p_j \geq dist(i,j)$. The economic interpretation of this is that each client is solely responsible for its costs incurred to reach the facility.

3. Let X be the subset of clients assigned to facility i, then the following holds:

$$\sum_{j \in X} \{p_j - dist(i,j)\} = f_i$$

This indicates that the opening cost of a given facility is paid for by clients assigned to it.

4. Equilibrium condition: For any facility $i \in F$ and any subset of clients $X \subseteq C$, the following must be true:

$$\sum_{j \in X} \{p_j - dist(i,j)\} \leq f_i$$

In the economic sense, this means that no subset of clients can find an economic incentive of opening a new facility at a lower cost. In other words, the amount that they are paying now will not be enough to open any other facility.

2.1 The Algorithm

2.1.1 Phase I

In the first phase, all unassigned clients uniformly raise their prices. When their offered price becomes larger than the connection cost to a given facility, they start contributing towards the opening cost of the facility. Once the opening cost of a facility is fully paid for, the facility is opened and all clients who had been contributing to its opening cost are assigned to it. If an unassigned client reaches an open facility, it is assigned to it and does not have to contribute to its opening cost. Phase I continues until there are no unassigned clients left. Note that at the end of this phase, clients do not retract their contributions to a given facility even if they end up assigned to another one.

Start with $p_j = 0 \forall j \in C$.
While $\exists$ unassigned clients
    All unassigned clients j uniformly raise their prices p_j until they reach a facility i when $p_j \geq dist(i,j)$, then they start contributing to its opening cost by $p_j - dist(i,j)$
    If the opening cost of a facility i becomes fully paid for:
        $\sum_{j \in C} \{p_j - dist(i,j)\}^+ = f_i$, where $\alpha^+ = max\{0, \alpha\}$
        Open facility i and assign to it all unassigned clients j such that $p_j \geq dist(i,j)$.
    When a client is assigned to a facility, it stops raising its price.
    If an unassigned client reaches an open facility, it is assigned to it without contributing to its opening cost.
End While

Phase I

Properties of Phase I

Claim 2.1.1 *If client j is assigned to facility i, then* $p_j \geq$ *dist(i,j).*
Proof: This is clear from the algorithm since it assigns client j to facility i only if $p_j \geq$ dist(i,j).

Claim 2.1.2 *(Equilibrium Condition) For any subset* $X \subseteq C$ *and any facility* $i \in F$, $\sum_{j \in X}\{p_j - dist(i,j)\} \leq f_i$.
Proof: For any given facility $i \in F$, the best subset of clients capable of violating the inequality would be $X = \{j : p_j > dist(i,j)\}$. However, in the algorithm as clients raise their prices, if at any point $\sum_{j \in C}\{p_j - dist(i,j)\}^+ = f_i$, facility i is opened and the sum cannot increase anymore. This is true since all clients just assigned to the facility stop raising their prices once assigned. On the other hand, any clients that later reach the facility do not have to contribute to its opening cost.

Corollary 2.1.3 *If the equilibrium condition holds, then the cost of any solution will be* $\geq \sum_{j \in C} p_j$.
Proof: Consider a solution $S^* \in F$, in which facility i is assigned a subset of clients $X_i \in C$. From the equilibrium condition, we have $\sum_{j \in X_i} p_j \leq f_i + \sum_{j \in X_i} dist(i,j)$. Summing over all facilities $i \in S^*$, we get

$$\sum_{i \in S^*}\{f_i + \sum_{j \in X_i} dist(i,j)\} \geq \sum_{i \in S^*}\sum_{j \in X_i} p_j = \sum_{j \in C} p_j$$

2.1.2 Phase II - Cleanup Phase

Note that at this point in the algorithm, every client is assigned to one facility. However, since clients do not retract their contributions to the opening cost of a facility even if they end up assigned to another, the open facilities are not fully paid for at the end of Phase I. We need to cleanup the solution so that, every client contributes only to the facility it is assigned to, and the total solution cost is $\leq 3\sum_{j \in C} p_j$.The cleanup phase consists of two steps: the first resolves conflicts among facilities, and the second assigns clients who lose their assigned facility due to the first step so that the final solution has the desired cost upperbound.

We consider facilities i and $i^{'}$ in conflict if there exists a client j such that $p_j - dist(i,j) > 0$ and $p_j - dist(i^{'},j) > 0$. We resolve conflicts in the following way:
Consider the facilities in the order in which they were opened in the first phase, and close facility i if some previously considered conflicting facility is open. After this step, for any client $j \in C$, there can be at most one open facility j such that $p_j > dist(i,j)$, and we assign j to i. If however no such facility exists, but there exist an open facility i such that $p_j = dist(i,j)$, we assign j to i. Note that some clients may remain unassigned, but that will be taken care of in the second step. We can claim the following after resolving all conflicts:

Claim 2.1.4 *For all clients* $j \in C$, *there is at most one open facility* i *such that* $p_j > dist(i,j)$.
Proof: Suppose not. Then, for some client j, we can find more than one facility i such that $p_j > dist(i,j)$. However, the facility conflict resolution step only leaves only one such facility open - contradiction.

Claim 2.1.5 *For any open facility* i, $\sum_{j \in X}\{p_j - dist(i,j)\} = f_i$, *where* X *is the subset of clients assigned to facility* i.

Proof: From the first phase, a facility i is open only when $\sum_{j \in C} \{p_j - dist(i,j)\}^+ = f_i$. If the facility remains open after the conflict resolution step, then all clients j such that $p_j - dist(i,j) > 0$ are assigned to it, and thus the equality would still hold.

At this point, assigned clients pay for all open facilities and their own connection costs. Yet, there may still be unassigned clients, and we will assign them in such a way that the total solution cost remains $\leq 3 \sum_{j \in C} p_j$.

Claim 2.1.6 *If client j is unassigned, we can assign it to an open facility i' such that $dist(i', j) \leq 3p_j$.*

Proof: If j is unassigned, then the facility i to which it was assigned was closed due to a conflict. This implies that there exists another client j' that had i and some other facility i' in conflict. Since i was closed due to that, i' must have been opened first, which implies that $p_{j'} \leq p_j$ (j' was assigned to i' and stopped raising its price, while j kept raising its price). We know that:

$$
\begin{aligned}
&dist(j,i) \leq p_j \\
&dist(j',i) \leq p_{j'} \\
&dist(j',i') \leq p_{j'}
\end{aligned}
$$

By the triangular inequality, $dist(j,i') < dist(j,i) + dist(j',i) + dist(j',i')$
$\leq p_j + 2p_{j'} \leq 3p_j$.

Corollary 2.1.7 *The total solution cost is $\leq 3 \sum_{j \in C} p_j$.*

Proof: Let X_F be the subset of clients assigned by the end of the first phase. These clients pay for the opening cost of all facilities and their connection costs. As for any client j assigned by the second phase, its connection cost can be at most $3p_j$. Then we have:

$$\sum_{j \in X_F} p_j = \sum_{i \in S} f_i + \sum_{j \in X_F} dist(j, i_j)$$

$$3 \times \sum_{j \notin X_F} p_j \geq \sum_{j \notin X_F} dist(j, i_j)$$

which implies that the total solution cost, $\sum_{i \in S} f_i + \sum_{j \in C} dist(j, i_j) \leq 3 \times \sum_{j \in C} p_j$

Resolve conflicts among facilities by keeping the ones that were opened first in the first phase.
For any client $j \in C$
  If $\exists$ an open facility i such that $p_j > dist(i,j)$, assign j to i
  Else if $\exists$ an open facility i such that $p_j = dist(i,j)$, assign j to i
  Else assign j to the open facility i that was in conflict with j's assigned facility at the end of Phase I

Phase II