Facility Location with Demand Dependent Costs and Generalized Clustering

Sudipto Guha* Adam Meyerson† Kamesh Munagala‡

May 12, 2000

Abstract

We solve the vaiant of facility location problem in which the costs of facilities depend on the demand served, more specifically decrease with the demand served. We show application of this problem to generalized clustering problems which does not penalize large clusters.

1 Introduction

In this paper we investigate the facility location problem where the costs of facilities are dependent on the demands served. We will concentrate on the subproblem where the costs actually *decrease* with the demand served. Some techniques are known for the complimentary case where the costs of facilities increase with demand served, but none of these techniques extend to the decreasing case. An interesting application of the problem with decreasing is the case of negotiable facility costs, where when we are leasing facilities we can negotiate a lower price due to associated business generated by attracting larger number of customers. This is easy to observe in cases of renting spaces from retail chains, or opening facilities in locations where there are possibility of associated business.

As another interesting application of the facility location problem with decreasing costs we investigate the problem of clustering. Most known clustering criterion for which approximate heuristics are known penalize large clusters, the minsum clustering for example. The facility location with decreasing costs allow us to solve for clustering scenarios that prefer larger clusters.

*Department of Computer Science, Stanford University CA 94305. Research Supported by IBM Research Fellowship, NSF Grant IIS-9811904 and NSF Award CCR-9357849, with matching funds from IBM, Mitsubishi, Schlumberger Foundation, Shell Foundation, and Xerox Corporation. Email: sudipto@cs.stanford.edu.

†Department of Computer Science, Stanford University CA 94305. Supported by ARO DAAG-55-97-1-0221. Email: awm@cs.stanford.edu.

‡Department of Computer Science, Stanford University CA 94305. Supported by ONR N00014-98-1-0589. Email: kamesh@cs.stanford.edu.

1.1 Previous Results

Several techniques were developed for the case where the costs increase with demand served. However in all the capacitated facility location problems we require that more than one copy of a facility be placed, or the capacity constraint relaxed. It can also be shown that without such the integrality gap of the problem even on very simple problems is unbounded. One such example where multiple copies are allowed, is solved in [5]. Although the paper does not explicitly mention such, their techniques extend to "buy at bulk" versions of the facility location problem, in which buying two facilities at the same location may be cheaper than twice the cost of one facility, but more than the cost of a single facility. However we will be interested in is the case that buying two copies is actually *cheaper* than buying a single copy. No previously known technique can be applied to this version.

1.2 Organization of the paper

We will require using a facility location variant which involves lower bounds on the demand serviced by a facility. This problem was solved in [4] in the context of hierarchical network design problems. We present the results mostly for the sake of completeness. Also in [4] we did not require the lower bounds to be different for each facility, but we do in the present application. The algorithm however extends to the more general case. This problem is discussed in Section 2.

The application to facility location is dicussed in Section 3. In Section 4 we discussion the application to clustering.

2 Facility Location with Lower Bounds

This problem is a variant of the classical facility location problem. We are given a network $G(V, E)$ with a distance function $d(\cdot)$ on the edges and a set of demand points. The cost of opening a facility at location i is f_i. In addition, there is a lower bound of L_i on the demand a facility opened at i must satisfy. The goal is to open facilities and allocate the demands to the open facilities so that an open facility at i has at least L_i demand routed to it. The cost of our solution is the sum of the average distance traveled by the demands and the cost of the open facilities. We wish to minimize this cost.

Since this problem is a generalization of the classical facility location problem, it is Max-SNP hard [3]. We are therefore interested in finding an approximation algorithm for this problem. Also note that if we could obtain a solution in which we satisfied the lower bounds exactly, we could solve the partition problem, which is NP-hard. Hence, we have to approximate the lower bounds as well.

Definition 2.1 *An approximation algorithm for load balanced facility location is a (α, β) approximation for some $\alpha \geq 1$ and $\beta \geq 1$ if the cost of the solution is within α times the optimal cost and facility i, if opened, serves at least $\frac{L_i}{\beta}$ demand.*

Let us denote by r the best known approximation ratio for classical facility location. We present a $(2r, 3)$ approximation to this problem. This generalizes to a $(\frac{1+\alpha}{1-\alpha}r, \frac{1}{\alpha})$ approximation for $\alpha < 1$.

$$\text{Minimize} \sum_i \sum_j d_j c_{ij} x_{ij} + \sum_i f_i y_i$$

$$\begin{array}{rcll} \sum_i x_{ij} & \geq & 1 & \forall j \\ x_{ij} & \leq & y_i & \forall i, j \\ \sum_j d_j x_{ij} & \geq & L_i y_i & \forall i \\ x_{ij}, y_i & \in & \{0, 1\} & \forall i, j \end{array}$$

We can write an integer program for this problem. Unlike facility location [7, 2, 5], the lower bound makes it hard to round the linear relaxation directly. This arises from the fact that the filtering steps of Lin and Vitter in [6] do not work. Thus fractional solutions cannot be rounded by previous approaches.

2.1 The Algorithm

The algorithm proceeds in two basic steps. Our transformations work for the fractional solution obtained from the linear relaxation of the integer program discussed above, so our final approximation guarantee is against the fractional solution.

Facility Location: For facility i, we add the cheapest way to route at least L_i units of demand to i to the facility cost f_i. We next solve regular facility location with these facility costs. Finally, we show that the optimum solution to this problem is within a factor 2 of the optimum to the original problem.

Rounding to Remove Facilities: Consider any open facility i that serves less than $\frac{L_i}{3}$ amount of demand. We close the facility and route the demands it serves to their closest open facilities. This transformation does not increase the cost of our solution.

2.2 Analysis

We now describe the two steps of the algorithm in detail.

Firstly, we construct a regular facility location instance from this problem. Each potential facility location i is now assigned a new cost f'_i, which is the sum of f_i and the minimum cost of routing exactly L_i amount of demand to that location. For doing this, we take demand points in increasing order of distance to i till we have collected L_i amount of demand.

Lemma 2.1 *Consider any feasible fractional solution to the load balanced facility location problem of cost C. We can construct a feasible instance of the regular facility location problem of cost at most $2C$.*

Proof: Look at any fractional facility i. Since the feasible solution is routing at least L_i amount of demand to any open facility, the facility cost we assign in the new problem is at most the routing cost of the demand connected to that facility. Thus the total additional facility cost is at most C. ■

We now solve the facility location instance mentioned above. The cost of the solution we obtain is within a factor of $r = 1.728$ to the optimal solution for that instance.

Therefore the total cost in the solution we compute is bounded in terms of the routing cost of the original fractional solution to within a factor of $2r$. Also note that facility location guarantees that each demand point goes to the closest open facility.

We now consider the open facilities in arbitrary order. Suppose facility i serves less than $\frac{L_i}{3}$ amount of demand, we close the facility and route the demands it serves to their closest open facilities. At the end of this process, we are guaranteed that each open facility i serves at least $\frac{L_i}{3}$ amount of demand, and each demand goes to the closest open facility.

We have to show that removing a facility does not increase the total facility plus routing cost of the solution. For this, we show a feasible way to route the demands it serves so that the cost does not increase.

Lemma 2.2 *Removing a facility i serving less than $\frac{L_i}{3}$ amount of demand cannot increase the cost of our solution.*

Proof: Suppose we are closing facility i. Consider the closest demand point j which does not send demand to this facility. Suppose $d(i, j) = D$, where d is the distance metric. If j is being served by i', $d(i', j) < D$, as each demand point goes to the closest open facility.

Also, the facility cost $f'_i \geq \frac{2L_i}{3} D$. This follows because the facility serves only $\frac{L_i}{3}$ amount of demand, while the facility cost f'_i is f_i plus the cost of serving at least L_i units of demand.

When we close the facility, we can afford to use f'_i towards re-routing the demand it serves. We send the demand to i', the facility serving j. The extra cost for doing this is at most the cost of taking the demand from i to j and from there to i'. This distance is at most $2D$, and the demand is at most $\frac{L_i}{3}$, and so the total re-routing cost is at most $\frac{2L_i}{3} D$. ■

The above can be summarized in the following theorem,

Theorem 2.1 *The load balanced facility location problem has a $(2r, 3)$ approximation where each demand is served by its closest open facility.*

We can scale the facility costs to improve the lower bounds. We will state the following tradeoff theorem.

Theorem 2.2 *The load balanced facility location problem has a $(\frac{1+\alpha}{1-\alpha} r, \frac{1}{\alpha})$ approximation for $\alpha < 1$ where each demand is served by its closest open facility.*

Proof: We start off by adding λ times the cheapest way to serve L_i units of demand to facility i to its cost. It is immediate that the approximation is $((\lambda + 1)r, \frac{1}{\alpha})$, where $\lambda = \frac{2\alpha}{1-\alpha}$. ■

3 Decreasing Facility Costs

We will now consider a variant of facility location where the facility cost is a *non-increasing* positive function of the amount of demand the facility serves. This model is appropriate in some situations where there is a startup operating cost for a facility (for example, a supermarket), but this cost decreases because of profits made with increasing number of customers.

Let us denote the cost function for facility i as f_i so that the cost of serving demand d is $f_i(d)$. Fix constants α and β larger than 1.

We assume the function f_i satisfies the following "nice" property. Let α and β be some constants larger than 1. For all demands d, $\frac{f_i(d)}{f_i(\alpha d)} \leq \beta$. Note that all rational functions satisfy this property.

We also assume that all demands are larger than 1.

We can obtain a constant approximation for this problem by converting it to an instance of load balanced facility location. For each facility i, we create many copies $i_0, i_1, \ldots$ of this facility. Copy i_j has lower bound α^j and cost $f_i(\alpha^j)$. The metric and the demand points remain the same.

We now solve the modified instance as a load balanced facility location instance. In our solution, we will lose a factor of β because of scaling of the demands. Also, if we insist the lower bound is at least $\frac{1}{\alpha}$ times the original bounds, our approximation ratio for the algorithm is $\frac{\alpha+1}{\alpha-1}r$. Since the actual cost is the cost of routing $\frac{1}{\alpha}$ fraction less demand to the facility, our facility costs go up by another factor of β. Therefore the final approximation ratio is $\frac{\beta^2(\alpha+1)}{\alpha-1}r$, which is a constant.

4 Generalized Clustering Problems

In this section, we define a generalization of the regular clustering problem, and show how to solve it using load balanced facility location.

In *generalized k-clustering*, we are given n points in a metric space with distance function d. The goal is to partition the points into k clusters, $C_1, C_2, \ldots, C_k$ so that the following objective function is minimized:

$$\sum_{l=1}^{k} \frac{1}{|C_l|^\gamma} \sum_{i,j \in C_l} d(i,j)$$

It is easy to see that an equivalent formulation to a factor of 2 is the following. We wish to partition the points into k clusters $C_1, C_2, \ldots, C_k$ with centers $c_1, c_2, \ldots, c_k$ so that the following objective function is minimized:

$$\sum_{l=1}^{k} \frac{1}{|C_l|^{\gamma-1}} \sum_{i \in C_l} d(i, c_l)$$

Note that for $\gamma = 0$, this is *min-sum* clustering, while for $\gamma = 1$, this is the k-median problem upto constant factors.

Definition 4.1 *A bicriteria approximation algorithm for generalized k-clustering is a (p, q) approximation if the cost of the solution is within a factor of p of the optimal cost, and the number of clusters produced is qk.*

The technique in [1] gives a constant bicriteria approximation for the case when γ is a constant less than 1. We show how to extend the result to constant $\gamma \geq 1$, thereby generalizing it to all constant γ. Our approximation factor will however be exponential in γ.

For every point i, we create n copies of the point and denote them $i^{(1)}, i^{(2)}, \ldots, i^{(n)}$. We will use the copy $i^{(r)}$ if the point i is the center of a cluster of size r. We can now formulate generalized k-clustering for $\gamma \geq 1$ as an integer program. Let $y_i^{(r)}$ denote whether i is the center of a cluster of size r. Let $x_{ij}^{(r)}$ be set to 1 if point j is in a cluster of size r with center at i.

$$\text{Minimize} \sum_i \sum_j \sum_r x_{ij}^{(r)} d(i,j) \frac{1}{r^{\gamma-1}}$$

$$\begin{array}{rcll} x_{ij}^{(r)} & \leq & y_i^{(r)} & \forall i, j, r \\ \sum_i \sum_r x_{ij}^{(r)} & \geq & 1 & \forall j \\ \sum_j x_{ij}^{(r)} & \geq & r y_i^{(r)} & \forall i, r \\ \sum_i \sum_r y_i^{(r)} & \leq & k & \end{array}$$

We can now formulate this as a facility location problem with lower bounds. Every point j has demand 1. Facility location $i^{(r)}$ has cost 1 and lower bound r. The new integer program formulation is therefore:

$$\text{Minimize} \sum_i \sum_j \sum_r x_{ij}^{(r)} d(i,j) \frac{1}{r^{\gamma-1}} + \sum_i y_i^{(r)}$$

$$\begin{array}{rcll} x_{ij}^{(r)} & \leq & y_i^{(r)} & \forall i, j, r \\ \sum_i \sum_r x_{ij}^{(r)} & \geq & 1 & \forall j \\ \sum_j x_{ij}^{(r)} & \geq & r y_i^{(r)} & \forall i, r \end{array}$$

Let F^* and S^* denote the optimum facility cost and service cost respectively. Note that $F^* = k$ for our problem.

For facility $i^{(r)}$, we add β times the cheapest way to ship r amount of demand to that facility. The optimum solution pays $F^* + \beta S^*$ for the facility cost, and S^* for the service cost.

Note that the distance from j to $i^{(r)}$ is $\frac{d(i,j)}{r^{\gamma-1}}$, and this is not a metric. But, note that the distances are multiplied by a number that depends only on the facility and not on the demand point. It is shown in [1] that the Jain-Vazirani algorithm [5] can be modified so that it works even if the metric is multiplied by a number that depends just on the facility.

We can therefore perform the first step of the load balanced facility location algorithm just as before. For the second step where we close facilities, we consider facilities which do not satisfy the lower bound in *increasing* order of r.

Lemma 4.1 *If we close facilities in Step 2 of the load balanced facility location algorithm in increasing order of r, then, the cost of our solution does not go up.*

Proof: Suppose we close a facility $i^{(r)}$ because it did not have enough demand coming to it. Let D be the distance in the modified metric to the closest point j which was not connected to $i^{(r)}$. Then, j went to some facility $l^{(r')}$ which was less than D distance in the modified metric, and $r' \geq r$. We therefore have: $d(i,j) = Dr^{\gamma-1}$ and $d(l,j) \leq Dr'^{\gamma-1}$.

Let the cost of routing the demands to i be C, and that to l be C'. If x is the amount of demand routed to i, we have:

$$C' \leq C \left(\frac{r}{r'}\right)^{\gamma-1} + xD + xD\left(\frac{r}{r'}\right)^{\gamma-1} \leq C + 2Dx$$

Therefore, when we add β times the cheapest service cost for satisfying the lower bound, we can obtain a constant approximation so that every open facility $i^{(r)}$ serves at least $r\frac{\beta}{2+\beta}$ demand. ■

Now, the algorithm in [5] gives the following approximation:

$$3F + S \leq 3(F^* + S^*)$$

We have modified the facility costs so that F^* becomes $F^* + \beta S^*$. We therefore have:

$$3F + S \leq 3(F^* + S^*(1+\beta))$$

Also, every open facility $i^{(r)}$ serves at least $r\frac{\beta}{2+\beta}$ demand. Now, we apply Lagrangian relaxation. We scale the facility costs by a factor of λ and then solve the modified problem. In the final solution, we scale back the facility costs. We will therefore have:

$$3\lambda F + S \leq 3(\lambda F^* + S^*(1+\beta))$$

We now guess the value of λ as $\frac{S^*}{\alpha F^*}$. We therefore have:

$$\begin{aligned} F &\leq (1 + (1+\beta)\alpha)F^* \\ S &\leq 3(1 + \beta + \tfrac{1}{\alpha})S^* \end{aligned}$$

Now, we are paying a routing cost which is scaled for a cluster of size r, while the actual cluster size could be $r\frac{\beta}{2+\beta}$. This causes our service cost to go up by a factor of $\left(1+\frac{2}{\beta}\right)^{\gamma-1}$. We therefore have the following theorem:

Theorem 4.1 *For any $\gamma \geq 1$, generalized k-clustering has a $\left(3(1+\beta+\frac{1}{\alpha})(1+\frac{2}{\beta})^{\gamma-1}, 1+(1+\beta)\alpha\right)$ bicriteria approximation.*

References

[1] Moses Charikar. Approximation algorithms for min-sum clustering. *manuscript*, 2000.

[2] Moses Charikar and Sudipto Guha. Improved combinatorial algorithms for facility location and k-median problems. *Proceedings of the Twenty-Ninth Annual IEEE Symposium on Foundations of Computer Science*, 1999.

[3] Sudipto Guha and Samir Khuller. Greedy strikes back: Improved facility location algorithms. *Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 649–657, 1998.

[4] Sudipto Guha, Adam Meyerson, and Kamesh Munagala. Hierarchical placement and network design problems. *Stanford University Tech. Note*, 2000.

[5] Kamal Jain and Vijay Vazirani. Primal-dual approximation algorithms for metric facility location and k-median problems. *Proceedings of the Twenty-Ninth Annual IEEE Symposium on Foundations of Computer Science*, 1999.

[6] J.-H. Lin and J. S. Vitter. ϵ-approximations with minimum packing constraint violations. *Proceedings of the Twenty-Fourth Annual ACM Symposium on Theory of Computing*, 1992.

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