Profit-Earning Facility Location

Adam Meyerson *
Department of Computer Science, Stanford University
Stanford CA 94305
awm@stanford.edu

ABSTRACT

We consider opening facilities in order to gain a profit. We are given a set of demand points, and we must open some set of facilities such that every demand may be satisfied from a local facility and the total profit gained in this process is maximized. This contrasts with previous work on facility location and k-center problems, where opening a facility incurred a cost. The profit gained by opening a facility is a function of the amount of demand the facility satisfies. We model the dependence of profit on demand by creating many different possible facilities at each location, each of which provides a certain profit if opened and requires at least a certain amount of demand in order to open. Our model captures problem instances where profits may be positive or negative, and also instances where it is not necessary to satisfy every demand. Our algorithms provide the optimum total profit, while stretching the definition of locality by a constant and violating the required demands by a constant. We prove that without this stretch, the problem becomes NP-Hard to approximate.

1. INTRODUCTION

Consider a company providing services to users. Every time the company opens a new facility, this facility should increase the total profit. The profitability of a facility depends on the number of users it serves; this dependence may not be linear but should be nondecreasing. The company will never open a facility with negative profit; it may be preferable to simply lose the potential revenue from some customers rather than guarantee service. However, if the company is constrained to serve every customer (a telephone company may be a good example of this), then the problem of negative profit facilities becomes interesting as well. We must also assign demands to open facilities; we constrain this assignment so that users will not have to travel more than a given distance D to reach their assigned facility. A greedy approach will not work well here; we can envision an instance where we would maximize profit by opening many small facilities, but we instead open a single large one and cannot open any others profitably (since so much demand is now dealt with).

One practical example of this problem is web caching, where "distance" represents internet latency and profits include revenue from advertising as well as the cost of the caches and revenue due to satisfied users. Another example is a dialup internet service provider, where "distance" represents the local area code and profit is determined by the number of customers and cost of servers. A third example is a chain of department stores, where distance represents driving time customers are willing to spend and profit represents sales minus upkeep. In these cases and many others, we would like to gain as much profit as possible. Since these problems are NP-Hard, we are interested in approximation algorithms. We observe that maximizing profit is very different from minimizing costs. A two-approximation on the costs may reduce the profit from positive to negative.

Our problem is a variant of facility location, in that we must select facilities to open and then map demand points to open facilities. We are given a bound on distance instead of optimizing for the average distance travelled, much as in the k-center problem [4]. However, known techniques for these problems fail to capture our goal of maximizing profit. Traditional facility location variants always assume that opening a facility is undesirable, and the algorithms try to open as few facilities as possible while still maintaining a facility reasonably near each user. In contrast, if we can obtain additional profits by opening more facilities, we will always do so. Traditional facility location variants also assume that the cost of a facility is independent of the number of users sent to that facility. Some recent work on capacitated facility location [5] has enabled facilities to take on concave costs. In our problem, the profit of a facility will be highly dependent on the amount of demand served. We do not require the profit to be concave or convex, only nondecreasing.

We call this problem *profit-earning facility location* and we will explore several variants of it. In the simplest version, every facility has a positive profit and a lower bound. In order to open the facility, we must send at least the lower bound in demand; if we do so, we obtain the profit. Our objective is to maximize profit. If we would like to capture some function $p(x)$ relating profit to demand, all we need to do is create a feasible facility with profit $p(x)$ and lower bound x for every integer x. Provided that each demand point has one unit of

*Research partially supported by ARO grant DAAG55-98-1-0170

STOC'01, July 6-8, 2001, Hersonissos, Crete, Greece.

demand, this introduces only polynomially many facilities. If we are not required to serve every demand, we can install artificial facilities with lower bound zero and profit zero (using these facilities corresponds to failing to satisfy the demand). We are required to make sure that each demand point is satisfied from a (real or artificial) facility within distance D. This problem is related to the *load-balanced facility location* problem introduced in [3, 6] because we must satisfy lower bounds on the demand served in order to open a facility. However, unlike the load-balanced problem, we are constrained by a maximum distance instead of minimizing total distance, and we gain a profit for facilities instead of assuming them to be free. Profit earning facility location is NP-HARD, so we will provide approximation algorithms.

Our approach is based on linear program rounding. The main difficulty comes in satisfying the lower bounds; if different facilities have different lower bounds then there is no straightforward way to merge facilities as was done in many previous papers. Previous results for the K-CENTER problem assume that when a facility opens we can send all demands within D to this facility. In our case, doing this may cause us to lose profit, since we could send a lesser amount of demand and still satisfy the lower bound, leaving enough demand free for an additional facility to open. We apply new techniques in the rounding of facilities to make sure that the lower bounds are approximately satisfied and no lower bound is more-than-satisfied unless it was so in the fractional solution. When rounding the assignment of demand points, we observe that the rounding used for Generalized Assignment [8] actually guarantees that the volume of any bin does not *decrease* by much in addition to guaranteeing that it does not *increase* by much, and so we apply this scheme (treating open facilities as bins, demand points as balls) to complete the rounding.

We consider an extension of profit earning facility location where some facilities have negative profit. Provided that we are required to satisfy every demand, this situation may well occur. Consider a telephone company, which is required to provide phone service to every customer regardless of profitability. Providing service to an isolated location may cost considerably more than the money paid by the customers in that location. Since service must still be provided, a negative profit service center will open. Of course, the telephone company still wants to maximize profits (rather than minimizing costs); even if total income is fixed at some dollar amount per user, a constant-approximation on costs will lead to a non-competitive profit total. We observe that this extension creates a *difference problem*; some facilities produce revenue while others have cost, and we would like to optimize for the total revenue minus the total cost. In general, difference problems are difficult to approximate because revenue and cost must scale together in order to guarantee any approximation factor.

Our algorithms obtain the optimum net profit. Since the problem is NP-HARD, we cannot (assuming $P \neq NP$) do this without violating some of the constraints. In particular, we will guarantee only that a demand point is never satisfied from a facility more than $3D$ away (instead of D) and that facility i is not opened unless at least a constant fraction of its lower bound demand is sent to i. This means we give an $O(1)$ approximation in a sense, though violating constraints by $O(1)$ is not equivalent to approximating the profit. We prove lower bounds which indicate that it is NP-HARD to provide an approximation which satisfies lower bounds exactly, or which does not extend the distances by a factor of three.

We consider several extensions, including the capacitated version of the problem for positive profits, and solutions which guarantee that some percentage of the overall demand is satisfied.

1.1 Previous Work

Facility Location and its variants have been studied extensively. In the traditional facility location problem, we would like to minimize the total cost of facilities plus the total distance from demands to facilities [9, 1, 5, 7]. In the K-MEDIAN problem [2, 1, 5], we are restricted to opening k facilities and we must minimize the total distance from demand to facilities. Our problem is more closely related to the K-CENTER problem; this problem makes guarantees on the *maximum* distance from any demand to to nearest open facility rather than the sum of distances. [4]. Load-balanced facility location, where lower bounds must be met in order to legally open facilities, was introduced independently by [3] and [6]. To our knowledge, we are the first to consider a model where opening facilities earns profit rather than incurring costs.

2. PROFIT EARNING FACILITY LOCATION

2.1 Problem Statement

We are given a graph $G = (V, E)$ with a length function $l : E \to R^+$ on the edges. Some subset of the vertices produce demand; the set of users U is the subset of vertices which produce nonzero demand, and we have a function $d : U \to R^+$ associating demand with the users. We are also given a set of feasible facilities F. Each facility i is represented by a *profit* p_i and a *lower bound* L_i. We must choose some subset of the facilities to open, and assign various users to each open facility such that every user is assigned to some open facility within a distance D, the total demand of users assigned to open facility i is at least L_i, and the total profit of open facilities is maximized.

The profit-earning facility location problem may be expressed by the following integer program. Here $a_i(u)$ indicates whether user u is assigned to facility i, and f_i indicates whether facility i is open.

- Maximize $\sum_{i \in F} p_i f_i$
- For all i: $\sum_{u \in U} a_i(u) d(u) \geq L_i f_i$
- For all i, u with $l(i, u) > D$: $a_i(u) = 0$
- For all i, u: $a_i(u) \leq f_i$
- For all $u \in U$: $\sum_{i \in F} a_i(u) = 1$
- $f_i \in \{0, 1\}$ for all i
- $a_i(u) \in \{0, 1\}$ for all i, u

The linear relaxation of this program replaces the last two constraints with $0 \leq f_i \leq 1$ and $0 \leq a_i(u) \leq 1$.

2.2 Approximation Algorithm

We will use the following algorithm to produce an approximate solution to the profit-earning facility location problem.

1. Solve the Linear Program relaxation
2. Set $I = \{i : 0 < f_i < 1\}$. I is the set of fractional facilities.
3. Select $i \in I$ which maximizes p_i/L_i.
4. Set $J = \{j : 0 < f_j < 1, l(i,j) \leq 2D\}$. J is the set of fractional facilities near i.
5. Set $\delta = \sum_{j \in J} \sum_{u \in U} a_j(u)d(u)$. This is the demand local to facility i.
6. Consider two cases based on the value of δ:
   - a If $\delta > L_i$ then for all $j \in J, j \neq i$ and all $u \in U$ increase $a_i(u)$ by $a_j(u)L_i/\delta$ and then multiply $a_j(u)$ by $(\delta - L_i)/\delta$. For all $j \in J, j \neq i$, multiply f_j by set $(\delta - L_i)/\delta$. Set $f_i = 1$.
   - b If $\delta \leq L_i$ then for all $j \in J, j \neq i$ and all $u \in U$ increase $a_i(u)$ by $a_j(u)$ and then set $a_j(u) = 0$. Then set $f_j = 0$ for all such j. Set $f_i = 1$.
7. If there still exist fractional f_j, return to step 2.
8. For each facility i which has $f_i = 1$, define δ_i to be the sum over $u \in U$ of $a_i(u)d(u)$. If $\delta_i < L_i/2$, then define V to be the set of users within D of facility i, and $d(V) = \sum_{u \in V}(1 - a_i(u))d(u)$. For all $u \in V$, set $a_i(u) = a_i(u) + (1 - a_i(u))(L_i/2 - \delta_i)/d(V)$. Reduce $a_j(u)$ for all $j \neq i$ accordingly.
9. Consider the facilities as "bins" and the demand points as "balls" and perform the rounding scheme for Generalized Assignment [8].

This algorithm is polynomial time, since we can solve a linear program in polynomial time and the number of loop iterations is linear in the number of possible facilities. We claim that no demand point has been sent more than $3D$ distance and that any open facility i satisfies at least $L_i/2 - \mu$ demand where μ is the largest demand of any user. The total profit of this solution is at least the optimum profit. We will prove these claims in section 2.2.1 and present approximation lower bounds in section 4.

2.2.1 Analysis

We would like to prove that the algorithm is in fact competitive. First, we observe that the equation requiring that the $a_i(u)$ sum to one is maintained as an invariant. Clearly when the algorithm terminates, all variables will be integers. We will first prove that the distance bound is approximately maintained; no user is sent to a facility more than $3D$ away.

LEMMA 2.1. *No user is sent to a facility which is more than $3D$ away.*

PROOF. We claim that at every stage of the algorithm, if $f_i < 1$ and $a_i(u) > 0$ then $l(i,u) \leq D$. This is clearly true of the fractional linear program solution. Suppose at some point it becomes untrue. This must be immediately after $a_i(u)$ was increased, which restricts us to considering stages 6 and 8 of the algorithm (we observe that step nine can only increase $a_i(u)$ if it was already nonzero). In step 6, we may increase the value of $a_i(u)$ for our selected i; the value $a_j(u)$ for other j will only decrease. But we will also set $f_i = 1$. So step six cannot cause the claim to become violated. In step eight, all facilities have f_i integer so the claim holds. On the other hand, consider what happens when we open a facility, setting $f_i = 1$. This happens only in stage six. At the time we open the facility, we may set $a_i(u)$ to become positive even though $l(i,u) > D$. However, any such u must have $a_j(u) > 0$ for some nearby $j \in J$. Since $j \in J$, we know that $f_j < 1$ and therefore $l(j,u) \leq D$. We also know that $l(i,j) \leq 2D$. Using the triangle inequality, it follows that $l(i,u) \leq 3D$ which proves the claim. □

We will next prove that all the lower bounds are approximately satisfied. Here the approximation ratio is $\frac{1}{2}$ before the Generalized Assignment rounding in the last step. After the last step, the approximation ratio becomes $\frac{1}{2} - \mu$ where μ is the maximum value of $d(u)/L_i$ for all choices of u and i with $L_i > 0$.

LEMMA 2.2. *If facility i is open, the total demand sent to facility i is at least $\frac{L_i}{2} - max_{u \in U} d(u)$.*

PROOF. Consider an open facility. If this facility was open in the fractional solution, then the lower bound was satisfied. The lower bound must remain satisfied until step 8, since flow cannot be removed from an open facility. If the facility had $f_i = 1$ set during step 6a, then we observe that when 6a completed the lower bound was satisfied. So these facilities too had lower bound satisfied when we reached step 8. Only the facilities opened during step 6b may pose a problem. Whenever step 6b opens a facility, all facilities in J have $f_j = 0$ set. It follows that no two 6b facilities may be $2D$ apart or closer. When we reach step 8, each facility may pull in users from D away. Only the 6b facilities have lower bounds unsatisfied, and no two of these will pull in the same users (since they are too far apart). Since they are only trying to satisfy $L_i/2$, and there must be enough demand within D to satisfy L_i, it follows that they pull in at most half the demand. Now consider a facility which is not from 6b. This facility can lose at most half its demand, and since its lower bound was previously satisfied it remains at least half-satisfied. So every facility which has $f_i = 1$ has at least half its lower bound in demand after step 8. The rounding scheme of [8] in step 9 guarantees that any facility which had $L_i/2$ demand before rounding still has at least $L_i/2 - \max_{u \in U} d(u)$ after rounding. □

We next prove that our total profit is at least the fractional optimum. Our approximation will violate *constraints* of the integer program while obtaining optimum profit. This is important for the case of negative profits to be considered in section 3.1.

LEMMA 2.3. *We obtain at least the fractional optimum profit.*

PROOF. We examine facilities in decreasing order of p_i/L_i. During step six, we move demand which was previously being sent to a fractionally-open facility of lesser p_i/L_i to one with larger. In 6a, the total demand sent to i after the step completes is exactly L_i. This means we have increased our profit by $(\delta - L_i)p_i/L_i$ and decreased it by the at most the same amount. Profit can only increase. In 6b, we moved all the demand to i. If we had increased f_i to be the total demand placed at i divided by the lower bound, the profit would have increased. But we instead set $f_i = 1$, an even larger value which guarantees that the profit must increase. So either way, the profit cannot decrease as we open facilities. □

We define an algorithm for profit-earning facility location to be (α, β, γ) competitive if it guarantees that distances increase by at most a factor of α, any open facility i which is open has at least βL_i demand, and the profit is at least the optimum profit times γ. The following theorem is direct from this definition and the lemmas above.

THEOREM 2.1. *The algorithm is* $(3, \frac{1}{2} - \mu, 1)$ *competitive, where* μ *is the maximum ratio of* $d(u)/L_i$ *for* $u \in U$, $i : L_i > 0$.

2.3 Tightening the Lower Bound

Suppose we would like to be closer to the lower bound. We can modify the algorithm in order to obtain a $(7, 1 - \mu, \frac{1}{2})$ approximation, sacrificing constants on the distance and profit in order to improve our approximation on the lower bound. We will need to modify the algorithm at steps 6b and 8 as shown below.

6b If $\delta \leq L_i$ then for each user which is assigned only to facilities in J, assign that user completely to i. For any user which is assigned at least fractionally outside of J, we increase $a_h(u)$ for h not in J with $a_h(u) > 0$ in order to send u completely outside of J. We then set $f_i = 1$ and $f_j = 0$ for $j \in J, j \neq i$.

8 Let B represent the set of open facilities with lower bounds unsatisfied. Let A represent the set of open facilities which are within $4D$ of a facility in B. We will close either all facilities of A or all facilities of B, whichever closing causes us to lose less profit. If we close the facilities of B, then we will redirect all demand points which were sent to those facilities to the nearest open facility. If we close the facilities of A, then we will send every demand point which is within D of some facility in B to that facility, and send all remaining unassigned demand points to the nearest facility.

The modified step 6b can only improve the profit, and guarantees that no two facilities opened in 6b can share any demand point (since the only demand points sent to such a facility will be sent there completely).

Again, define $\mu = \max_{u,i} d(u)/L_i$. We will prove the desired approximation ratio below.

THEOREM 2.2. *We can produce a* $(7, 1 - \mu, \frac{1}{2})$ *approximation.*

PROOF. When we arrive at step 8, our total profit is at least the fractional optimum, but we may have violated some lower bounds. When we choose the least costly of the two sets to close, we can guarantee that our new profit is at least half the fractional profit. Suppose we close the facilities of B. It is clear that lower bounds are now satisfied exactly. On the other hand, suppose we close the facilities of A. Any demand point which is within D of some facility in B is sent only to facilities in B, since any remaining facility not in B must be too far away. There is sufficient demand within D to open any facility, since otherwise this facility could never have $f_i > 0$. It follows that we are able to satisfy all lower bounds if A closes. Finally, we need to show that each demand point has a facility within $7D$. A demand point which was sent to A has a facility within $7D$ if A closes, because each facility in A is within $4D$ of some facility in B and this demand point will be sent a total of $3D + 4D = 7D$. On the other hand, suppose there is some demand point which is sent to B but is far from every facility in A. Why couldn't we satisfy the lower bound when the facility for this point was opened? It must be that some point which was fractionally sent to this facility was also fractionally sent to some other facility which was already open. But facilities in B cannot share demand points, so that other facility must have been in A, and it follows that every facility in B is within $4D$ of some facility in A. Thus no point is sent more than $7D$. The Generalized Assignment rounding stage in step 9 again causes us to lose μ on the lower bound. □

3. EXTENSIONS

3.1 Negative Profits

We consider the case where facilities may have negative profits. For example, suppose we are caching web content and would like to guarantee a local cache for every user. It costs us some amount of money to set up a cache, so profits of the caches may be negative. However, we also gain income by directing advertisements through our caching network. If sufficiently many users are local to a single cache, that cache may earn a net profit because of advertising. In this case some facilities will have costs and some (usually those with large lower bounds) will have profits. We would like to extend our algorithm to deal with the situation where there are both positive and negative profits. This is a *difference problem*, and it is important that we produce exactly the fractional optimum, since if we obtain half the profit and pay the full cost, we may take a loss instead of earning a profit.

The new algorithm treats the positive profit facilities in the same way, but we need to do something different for the negative profit facilities. The algorithm is as follows:

1. Solve the Linear Program relaxation

2. Set $I = \{i : 0 < f_i < 1\}$

3. Select $i \in I$ which maximizes p_i/L_i.

4. Set $J = \{j : 0 < f_j < 1, l(i,j) \leq 2D\}$. J is the set of facilities near i.

5. Set $\delta = \sum_{j \in J} \sum_{u \in U} a_j(u)d(u)$.

6. Consider two cases based on the value of d:

   a If $\delta > L_i$ then for all $j \in J, j \neq i$ and all $u \in U$ increase $a_i(u)$ by $a_j(u)L_i/\delta$ and then multiply $a_j(u)$ by $(\delta - L_i)/\delta$. For all $j \in J, j \neq i$, multiply f_j by set $(\delta - L_i)/\delta$. Set $f_i = 1$.

b If $\delta \leq L_i$ then for all $j \in J, j \neq i$ and all $u \in U$ increase $a_i(u)$ by $a_j(u)$ and then set $a_j(u) = 0$. Then set $f_j = 0$ for all such j. Set $f_i = 1$.

7. If there still exist fractional f_j with $p_j/L_j \geq 0$, return to step 2.

8. Find a user u such that $f_i < 1$ for all i with $a_i(u) > 0$. Select facility i which has the maximum (least negative) p_i and satisfies $a_i(u) > 0$. Set $f_i = 1$. For every nearby facility $j \neq i$ with $f_j < 1$ and $l(i,j) \leq 2D$, increase $a_i(v)$ by $a_j(v)$ for all $v \in U$ and then set $a_j(v) = 0$. Set $f_j = 0$. We are closing all these facilities and sending the demand to i. Repeat until no user with the required property exists.

9. For every facility i with $f_i < 1$ set $f_i = 0$. If $f_i = 0$ then set $a_i(u) = 0$ for all u.

10. For each user u, let $\delta(u) = 1 - \sum_i a_i(u)$. If $\delta(u) > 0$ then select some facility i with $f_i = 1$ and $l(i,u) < 3D$ and increase $a_i(u)$ by $\delta(u)$.

11. For each facility i which has $f_i = 1$, compute $\delta_i = \sum_{u \in U} a_i(u)d(u)$. If $\delta_i < L_i/2$, then define $V = \{u \in U : l(u,i) \leq D\}$, along with $d(V) = \sum_{u \in V}(1 - a_i(u))d(u)$. For all $u \in V$, set $a_i(u) = a_i(u) + (1 - a_i(u))(L_i/2 - \delta_i)/d(V)$. Reduce $a_j(u)$ for all $j \neq i$ accordingly.

12. Consider the facilities as "bins" and the demand points as "balls" and perform the rounding scheme for Generalized Assignment [8].

The first seven steps are identical to the original algorithm, dealing only with the positive profit facilities. In step 8, notice that the sum of f_i for the facilities with $a_i(u) > 0$ must be at least 1. When we open the least-costly facility, we will close all other facilities with $a_i(u) > 0$ since the fractionality of f_i guarantees that these facilities are within D of u and within $2D$ of each other. This means the cost will only improve. Any facility opened in step 8 may fail to meet its lower bound, but we guarantee that any two facilities which do not meet their lower bounds are more than $2D$ apart. In step 9 we are only closing negative-profit facilities so our profit improves. In step 10, notice that all remaining users are sent at least fractionally to some fully open facility within $3D$. It follows that there does exist an open facility within $3D$. Step 11 works just as before, since we again guarantee that no two open facilities with lower bounds unmet are within $2D$ of one another. Step 12 repeats the Generalized Assignment Rounding. It follows that we are again $(3, \frac{1}{2} - \max_{u,i} d(u)/L_i, 1)$ competitive with respect to optimum.

3.2 Capacitated Version

Suppose each facility has a capacity in addition to a lower bound. In order to open a facility i, we must have $L_i \leq \sum_u a_i(u)d(u) \leq C_i$. We can use this to define *any positive* function relating profit of a facility to the amount of demand sent to that facility. We can discard any facility with $C_i < L_i$ as invalid.

We will use the same algorithm as for the uncapacitated case of positive profits, except that we will add capacity constraints to the linear program. We observe that the demand placed at a facility which did not have $f_i = 1$ in the linear program solution will not exceed that facility's lower bound (and therefore cannot exceed the capacity) until the Generalized Assignment rounding step. It follows that the only way we can violate any capacity constraint is during the Generalized Assignment rounding, and the proof of [8] guarantees that this violation is at most by $\max_u d(u)$. The theorem below follows.

THEOREM 3.1. *If facility i is open, then $\frac{L_i}{2} - max_u(d(u)) \leq \sum_u a_i(u)d(u) \leq C_i + max_u(d(u))$.*

3.3 Linear Gains

Consider the case where profit of a facility grows linearly with the amount of demand sent there. This models, for example, the case of setting up web pages. Suppose we are setting up pages to sell something. We receive profit linear in the number of customers, plus an additional profit due to advertising revenue (this increases at various steps as we reach certain numbers of users), minus costs for maintaining servers (these costs also increase linearly with the number of users, but not so quickly as our profits).

We could approximate this problem by setting up a facility with each integer lower bound, and computing the overall profitability at that number of users. However, if users have variable "size" then the number of potential facilities may be exponentially large. We can get around this problem by again observing that facilities which were opened by the fractional solution remain open. If other facilities are sent more than their lower bound in demand, this will only help us because of the linear increases. In fact, this will work even in the case where some profits are negative. We build the linear function into the linear program formulation, and ignore it thereafter. This again obtains the same approximation as before.

3.4 User Satisfaction

Suppose we would like the maximum profit solution which provides service to at least some percentage of the users. This models many real life applications where we are allowed to refuse service to some users if providing such service would result in a net loss, but we cannot refuse service to a very high percentage of users.

We replace the linear program constraint that $\sum_i a_i(u) = 1$ with constraints that $\sum_i a_i(u) \leq 1$ and $\sum_u \sum_i a_i(u) \geq X$ for some constant X. After solving the linear program, we perform the identical rounding scheme given for the original algorithm. We observe that the only way a facility can disappear is when a facility within $2D$ is opened. It follows that every user with $\sum_i a_i(u) > 0$ has a facility within $3D$, so the number of users serviced can only increase. We provide a $(3, \frac{1}{2} - \mu, 1)$ approximation while guaranteeing at least the required percentage of users are served.

We observe that in the case where there are negative profits, this scheme does not work and in fact there is an arbitrarily large LP gap. Consider a case where all the users are near only a single potential facility, and this facility has profit -1 and lower bound zero. We have an additional facility, infinitely far away, with profit 1 and lower bound zero. We are required to satisfy some percentage $(0 < p < 100\%)$ of the users. The optimum integer solution has profit zero, since the negative facility must open. The optimum fractional solution has positive profit, since it can open the negative facility by p percent and then open the positive one completely, for a profit of $1-p>0$. Allowing finite stretch on the distances or lower bounds does not help with this problem.

4. LOWER BOUNDS

We prove several lower bounds on profit-earning facility location. We will first show the necessity of the factor three approximation on the distance, then show that we cannot meet lower bounds exactly.

Unless lower bounds are nonexistent (or completely ignored), the following theorem indicates that we need a factor of three stretch on the distances.

THEOREM 4.1. *It is* NP-HARD *to find an* (α, β, γ) *approximation to profit-earning facility location unless either* $\alpha \geq 3$ *or* $\beta = 0$.

PROOF. The proof is by reduction from set cover. Suppose we have a set cover instance we would like to solve. We create a profit-earning facility location with one node for each set and one node for each element. All "set" nodes are $2D$ from each other, all "element" nodes are $2D$ from each other. A set node is D from the element nodes corresponding to elements of that set; a set node is $3D$ from the element nodes not belonging to the set. We add an additional node which is D from every set node. It's not hard to see that this structure satisfies the triangle inequality. We set a demand of 1 at each element node, and we set k demands of $(n+1)/\beta$ at the additional node which is D from every set. A potential facility exists at every set node, with lower bound $(n+1)/\beta$. If there exists a set cover of size at most k, then this problem will have a feasible solution which opens facilities at the sets in the cover, sending one of the "big" demands to each set and sending each unit demand to the set covering that element. On the other hand, suppose we can find a solution which is (α, β, γ) approximate. If $\alpha < 3$, each of the demand one nodes must be sent to a facility which is within D (since all other facilities are $3D > \alpha D$ away). Each of these open facilities must have at least $\beta(n+1)/\beta = n+1$ units of demand. As there are only n unit demands, one of the "big" demands must have been sent here. Thus there are at most k open facilities. We conclude that this solution provides a set cover. Since set cover cannot be solved exactly in polynomial-time unless $P = NP$, we can conclude that it is NP-HARD to find an approximation with $\alpha < 3$ and $\beta > 0$. □

The following theorem indicates that we cannot obtain optimum profit without any error on the lower bounds, regardless of the amount by which we stretch the distances.

THEOREM 4.2. *It is* NP-HARD *to find an* (α, β, γ) *approximation to profit-earning facility location unless either* $\beta < 1$ *or* $\gamma \leq \frac{1}{2}$.

PROOF. The proof is by reduction from subset sum. Suppose we would like to determine whether some subset of a list of numbers sums to k. We create an instance of profit-earning facility location where there is only a single node. At this node we have a demand for each of the numbers on the list, with the demand equal to the number itself. We also have two potential facilities, one with lower bound k and one with lower bound $s-k$ where s is the sum of all numbers. Each facility has profit one. Assuming $\beta = 1$, we can open both facilities if and only if some subset of the numbers sums to k. If we open only one facility, then we must have $\gamma \leq \frac{1}{2}$. □

The following theorem indicates that for the difference problem (where profits may be negative), we cannot guarantee that lower bounds will be satisfied unless we completely ignore either the profit or the distances.

THEOREM 4.3. *It is* NP-HARD *to find an* (α, β, γ) *approximation to profit-earning facility location with negative profits, unless either* $\beta < 1$ *or* $\gamma = 0$ *or* $\alpha = \infty$.

PROOF. Again the proof is by reduction from subset sum. We construct a graph with two nodes, infinitely far apart. Node one has a demand of one, and a possible facility with lower bound zero and profit -1. Node two has a demand for each of the numbers from the subset sum problem, each demand being the size of the corresponding number. Node two has two potential facilities, one with lower bound k and one with $s-k$, where k is the subset sum and s is the sum of all the numbers from the subset sum instance. Each of these facilities has profit 1. Unless α is infinite, we must open the facility with negative profit. If there is a subset of numbers summing to k, then the optimum solution will have profit 1. Otherwise the optimum has profit 0. If we can guarantee any nonzero competitivity against the profit, we must open both facilities every time the optimum does so. It follows that $\beta < 1$ unless we can solve subset sum in polynomial time. □

5. ACKNOWLEDGMENTS

The author would like to thank Kamesh Munagala and Serge Plotkin for many useful discussions.

6. REFERENCES

[1] M. Charikar and S. Guha. Improved combinatorial algorithms for facility location and k-median problems. *Proceedings of 40th IEEE FOCS*, 1999.

[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, Meyerson A., and K. Munagala. Hierarchical placement and network design problems. *Proceedings of 41st IEEE FOCS*, 2000.

[4] Dorit Hochbaum and David Shmoys. A best possible heuristic for the k-center problem. *Math of Operations Research*, 10(2):180–184, 1985.

[5] K. Jain and V. Vazirani. Primal-dual approximation algorithms for metric facility location and k-median problems. *Proceedings of 40th IEEE FOCS*, 1999.

[6] D.R. Karger and M. Minkoff. Building steiner trees with incomplete global knowledge. *Proceedings of 41st IEEE FOCS*, 2000.

[7] J.-H. Lin and J. S. Vitter. ϵ-approximations with minimum packing constraint violations. *Proceedings of 24th ACM STOC*, 1992.

[8] D. Shmoys and É. Tardos. An approximation algorithm for the generalized assignment problem. *Proceedings of the 4th Annual ACM-SIAM SODA*, 1993.

[9] D. B. Shmoys, É. Tardos, and K. Aardal. Approximation algorithms for facility location problems. *Proceedings of 29th ACM STOC*, 1997.