The Facility Location Problem with General Cost Functions

M.T. Hajiaghayi*, M. Mahdian*, and V.S. Mirrokni*

Abstract

In this paper we introduce a generalized version of the facility location problem in which the facility cost is a function of the number of clients assigned to the facility. We focus on the case of concave facility cost functions. We observe that this problem can be reduced to the uncapacitated facility location problem. We analyze a natural greedy algorithm for this problem and show that its approximation factor is at most 1.861. We also consider several generalizations and variants of this problem.

Keywords: Approximation algorithms, facility location problem, generalized facility location, concave facility location.

1 Introduction

The facility location problem is a central problem in operations research. In this problem, we have a set of clients and a set of facilities, and we want to connect each client to a facility in a way that minimizes the total cost. There are two types of costs associated with a solution: the connection costs and the facility costs. The *connection cost* between client j and facility i is a given number c_{ij}. We will have to pay this amount if we connect client j to facility i. Each facility i also has a *facility cost* f_i, which is the amount we have to pay if we connect at least one client to it. The cost of a solution is the sum of the connection costs and the facility costs associated with the solution. We will generally assume that the connection costs obey the metric inequality.

The facility location problem is well studied in the field of approximation algorithms, and a number of different approximation algorithms have been proposed for this problem using a variety of techniques (see [23] for a survey). The first constant-factor approximation algorithm for this problem was given by Shmoys, Tardos, and Aardal [24] and was based on LP rounding and a filtering technique due to Lin and Vitter [18]. Later, the factor was improved by Chudak and Shmoys [3, 4] to $1 + 2/e$. Jain and Vazirani [13] gave a primal-dual algorithm for this problem, achieving a factor of 3. Strategies based on local search and greedy improvement for facility location problems have also been studied. The work of Korupolu et al. [15] shows that a simple local search heuristic proposed by Kuehn and Hamburger [16] yields a $(5 + \epsilon)$-approximation algorithm for any $\epsilon > 0$. This was later improved by Arya et al. [1] to 3. Charikar and Guha [2] improved the factor of the LP-based algorithm [3] slightly to 1.728 by combining the primal-dual algorithm [13] with greedy augmentation. They also combined

*Laboratory for Computer Science, MIT, Cambridge, MA, USA. Emails: {hajiagha,mahdian,mirrokni}@mit.edu.

greedy improvement and cost scaling to improve the factor of the primal-dual algorithm [13] to 1.853. Mahdian et al. [19] used a variant of the primal-dual schema, called the dual-fitting method, to give a combinatorial algorithm with an approximation ratio of 1.861. This was later improved by Jain et al. [12] and Mahdian et al. [20] to 1.61 and 1.52, respectively. The latter is the best approximation factor for this problem known to date. The LP rounding algorithm was improved by Sviridenko [25] to get a factor of 1.58. Guha and Khuller [9] showed that this problem is hard to approximate within a factor better than 1.463, assuming $NP \not\subseteq DTIME[n^{O(\log\log n)}]$.

A natural generalization of the facility location problem is to allow the cost of a facility to be an arbitrary function of the number of cities connected to it. This question was first asked by T. Leighton [17] and was motivated by its applications in placing servers on the Internet. Suppose a company needs to place servers on the Internet to serve a given distribution of demands. There is a connection cost between each demand point and each server. The cost of installing and maintaining a server is a function of the amount of demand served by it. Because of economies of scale, these cost functions are usually concave. Therefore, the case of concave cost functions is of special interest. However, other types of cost functions also occur in several applications. For example, many variants of the facility location problem such as soft or hard capacitated facility location [5, 6, 15, 22] and lower-bounded (a.k.a. load-balanced) facility location [10, 14] are special cases of this.

This paper is organized as follows. In Section 2, we give a formal definition of the problem, and observe that it can be reduced to both the capacitated and the uncapacitated facility location problems. In Section 3, we present a greedy 1.861-approximation algorithm for the concave facility location problem. Our algorithm is a natural generalization of the algorithm of Mahdian et al. [19]. The analysis, presented in Section 4, uses the techniques of dual-fitting and factor-revealing LPs as used in [19]. The dual-fitting approach (without the name "dual-fitting") was first used by Chvátal [7] to analyze the greedy algorithm for min-cost set cover. The technique of using factor-revealing LPs consists of expressing the constraints imposed by the problem statement and the algorithm as linear inequalities and bounding the approximation factor of the algorithm by solving a series of linear programs. This technique is similar to the idea of LP bounds in coding theory [21]. In the context of approximation algorithms, Goemans and Kleinberg [8] use a similar method in the analysis of their algorithm for the minimum latency problem. See Jain et al. [11] for a discussion of these techniques. In Section 5, we study generalizations and variants of the problem such as allowing the facility cost function to be slightly non-concave, facility location with convex cost functions, and relaxing the metric inequality for the connection costs.

2 The problem

We define the *generalized facility location* problem as follows. We are given a set $\mathcal{F}$ of *facilities* (a.k.a servers); a set $\mathcal{D}$ of *clients* (a.k.a. cities or demands); a non-decreasing *facility cost function* $f_i : N \to N$ for every facility $i \in \mathcal{F}$ with $f_i(0) = 0$, which specifies the cost of the facility as a function of the number of clients served by it; and finally *connection costs* c_{ij} between facility i and client j. We define n to be the total number of facilities and clients. We generally assume that the connection costs form a *metric* (i.e., are symmetric and satisfy the triangle inequality), unless otherwise stated. The objective of the problem is to find an assignment ψ of *all* clients to facilities (i.e., $\psi : \mathcal{D} \mapsto \mathcal{F}$) with minimum total cost. The total cost is the sum of the connection costs

$(\sum_{j\in\mathcal{D}} c_{\psi(j),j})$ and the facility costs $(\sum_{i\in\mathcal{F}} f_i(|\{j\in\mathcal{D}:\psi(j)=i\}|))$.

The focus of this paper is on the *concave facility location problem*, i.e., the generalized facility location problem where all facility cost functions are concave. Recall that a function $f: N \to N$ is called concave if and only if for each $x \geq 1$, $f(x+1)-f(x) \leq f(x)-f(x-1)$.

We observe that the concave facility location problem can be reduced to the *capacitated facility location*. In the capacitated facility location problem, each facility i also has a capacity u_i, which is the maximum number of clients that can be served by this facility. The problem has two variants; in the first one, the capacities are *hard* in that facility i can be opened at most once to serve at most u_i clients; in the second variant the capacities are *soft*; i.e., facility i may be opened k times to serve up to ku_i clients at a cost of kf_i. It is easy to observe that both of these problems are special cases of the generalized facility location problem. The capacitated facility location problem with soft capacities can be reduced to the uncapacitated facility location problem [20], and the best known approximation factor of 2.89 for this problem is obtained via this reduction [20]. However, hard capacities are known to be problematic for techniques that rely on linear programs (such as LP-rounding and primal-dual algorithms). For the case of uniform capacities, constant-factor algorithms were given in [6, 15]. Recently, Pal et al. [22] used local search techniques to give the first constant-factor approximation algorithm for this problem with non-uniform capacities, achieving a factor of $9+\epsilon$. In the following, we show that the concave facility location problem can be reduced to both variants of the capacitated facility location problem.

Proposition 1 *Assume there is a polynomial time α-approximation algorithm for the capacitated facility location problem (with hard or soft capacities). Then the concave facility location problem can be approximated within a factor of α in polynomial time.*

Proof: First consider hard capacities. For each facility $i \in \mathcal{F}$, we place $|\mathcal{D}|$ copies of this facility, where the jth copy has opening cost $f_i(j)$ and capacity j. We use the α-approximation algorithm for the capacitated facility location problem to solve this instance. Since the facility cost function is concave, we can assume without loss of generality that at most one of these $|\mathcal{D}|$ copies is opened in this solution. Therefore, the solution to the capacitated facility location instance can be easily transformed to a solution to the concave facility location problem with the same total cost. Now, we observe that it suffices to use soft capacities. The jth copy of facility i may now serve kj clients at a facility cost of $kf_i(j)$. However, since f_i is concave, $kf_i(j) \geq f_i(kj)$, so we are only overestimating the cost. $\square$

The above proposition together with the soft-capacitated facility location algorithm of Mahdian et al. [20] yields a 2.89-approximation algorithm for the concave facility location problem. It was noticed by J. Zhang [27] that it is in fact possible to reduce the concave facility location problem to the uncapacitated facility location problem. This reduction, together with the algorithm of Mahdian et al. [20], gives an algorithm with an approximation factor of 1.52 for the concave facility location problem with a running time of $O(n^6)$.

Proposition 2 *Any polynomial time α-approximation algorithm for the uncapacitated facility location problem yields a polynomial time α-approximation algorithm for the concave facility location problem.*

Proof: The main idea is that any concave function $f(x)$ can be represented by $\min_k\{L_k = (f(k) - f(k-1))x + kf(k-1) - (k-1)f(k)\}$, $1 \leq k \leq |\mathcal{D}|$. Here L_k is the line passing through points $(k-1, f(k-1))$ and $(k, f(k))$. Now, we take each facility i with concave cost function f_i, and replace it by multiple facilities $i_1, i_2, \ldots, i_{|\mathcal{D}|}$ such that facility i_k has opening cost $kf_i(k-1) - (k-1)f_i(k)$. In addition, each unit of demand routed to this location costs $f_i(k) - f_i(k-1)$ extra at the facility. This cost $f_i(k) - f_i(k-1)$ can be added to the connection costs still keeping them metric. This reduction gives an uncapacitated facility location instance. In addition, since function f_i is concave, the most efficient way to assign k cities to a facility i is to use the kth copy i_k in the uncapacitated facility location instance. $\square$

3 The greedy algorithm

In this section, we present a faster greedy 1.861-approximation algorithm for the *concave facility location problem*. Our method is an extension of the method used by Mahdian et al. [19] and Jain et al. [12]. We start with the following definition from [12].

Definition. A *star* consists of one facility and several clients. For star S, consisting of clients $a_1, a_2, \ldots, a_k$ and facility p, the cost of S, denoted by $c(S)$ or c_S, is defined as $\sum_{i=1}^{k} c_{p,a_i} + f_p(k)$, i.e., the sum of the connection costs of the clients to facility p plus the facility cost of p for k clients.

Suppose our algorithm finds a solution SOL of cost $cost(\text{SOL})$, and values α_j for every client $j \in \mathcal{C}$ such that $\sum_{j \in \mathcal{C}} \alpha_j = cost(\text{SOL})$ (one can think of α_j as the contribution of client j in the total cost). In addition, assume that there is a constant $\gamma \geq 1$ such that for every star S, $\sum_{j \in S \cap \mathcal{C}} \alpha_j \leq \gamma c_S$. We now consider an optimal solution OPT to the problem. Let i be a facility that is opened in OPT, D_i be the set of clients that are served by it, and S_i be the star consisting of i and D_i. By writing the inequalities $\sum_{j \in D_i} \alpha_j \leq \gamma c(S_i)$ for every i that is open in OPT, we obtain $cost(\text{SOL}) = \sum_{j \in \mathcal{C}} \alpha_i \leq \gamma \sum_{S \in \text{OPT}} c_S = \gamma \cdot cost(\text{OPT})$. Therefore if we can prove the statement $\sum_{j \in S \cap \mathcal{C}} \alpha_j \leq \gamma c_S$ for every star S, then γ is an upper bound on the approximation factor of the algorithm. This technique is called *dual-fitting*, and is explained in [11].

In the next section using the approach discussed above, we show that the following greedy algorithm is a 1.861-approximation algorithm for the concave facility location problem. In this algorithm, we use a notion of time, so that every event can be associated with the time at which it occurs. Also each client j has a budget from which it can offer some money to facilities; if j is unconnected and its budget is more than the cost of the connection to a facility i, it offers the extra budget to i (line 10).

Algorithm A: greedy algorithm for concave facility location
Input: Metric connection costs c_{ij} for each facility i and client j.
Concave non-decreasing facility cost functions $f_i : N \rightarrow N$ for each facility i.
Output: For each client j, a facility $\psi(j)$ to which j is assigned.
For each client j, contribution α_j of client j to the total cost.
begin
1 **let** $t = 0$
2 **for** each facility i **let** $\text{level}_i = 0$
3 **for** each client j

let $\psi(j) = $ **null**
let budget$(j) = 0$
end for
while there is an unconnected client increase time t
for each unconnected client j **let** budget$(j) = t$
for each client j and each facility i
if $\psi(j) = $ **null let** offer$(j, i) = \max(\text{budget}(j) - c_{ij}, 0)$
else let offer$(j, i) = 0$
end for
if there is a facility i and unconnected clients $a_1, \ldots, a_k$
s.t. $\sum_{j=1}^{k} \text{offer}(a_j, i) = f_i(\text{level}_i + k) - f_i(\text{level}_i)$ **then**
let level$_i$ = level$_i + k$
for each $1 \le j \le k$ **let** $\psi(a_j) = i$
end if
end while
for each client j
let α_j = budget(j) (i.e., the time that j first gets connected)
end for
end

The above algorithm is similar to the greedy algorithm of Mahdian et al. [19]. Using an inductive argument, we can prove that for any set S of unconnected clients, any facility i, and any time t the invariant $\sum_{j \in S} \text{offer}(j, i) \le f_i(l_i + |S|) - f_i(l_i)$ holds, where l_i is the level of facility i at time t. At the beginning, the condition is obviously correct. We only need to verify that after assigning a set S of clients to facility i, the above inequality still holds. Assume, to the contrary, that after assigning clients in S to facility i, there is a set S' of unconnected clients for which $\sum_{j \in S'} \text{offer}(j, i) > f_i(l_i + |S| + |S'|) - f_i(l_i + |S|)$. We know that $\sum_{j \in S} \text{offer}(j, i) = f_i(l_i + |S|) - f_i(l_i)$, when we assign S to i. It means $\sum_{j \in S \cup S'} \text{offer}(j, i) > f_i(l_i + |S| + |S'|) - f_i(l_i)$ before assigning S to facility i contradicting the induction hypothesis. Notice that if several events occur at the same time, we can process them in an arbitrary order.

The proof of the following lemma is clear from the algorithm and the discussion above.

Lemma 3 *The total cost of the solution found by Algorithm A is equal to the sum of the α_j's.*

By an analysis very similar to the analysis of Mahdian et al. [19], the running time of Algorithm A is at most $O(n^3 \log n)$. More precisely, in order to implement Algorithm A, we can use a heap data structure to maintain a list of events. There are two types of events: (a) City j starts contributing towards opening facility i; and (b) The number of cities assigned to facility i (i.e., level$_i$) is increased by at least one. The analysis is the same as Mahdian et al. [19], except here instead of one event for opening a facility i we might have n events for increasing level$_i$. Thus the running time is $O(n^3 \log n)$ instead of $O(n^2 \log n)$.

4 The approximation factor

In this section, we show that Algorithm A is a 1.861-approximation algorithm for the concave facility location problem. We prove this by showing that for each star S, the ratio of the sum of α_j's of all clients contained in S to the cost of S is at most 1.861. As in [19], our approach is to write various constraints imposed by the structure of the algorithm as linear inequalities, and find the maximum ratio of the sum of α_j's of all clients in a star to the cost of the star by solving a linear program called the *factor-revealing LP*. See [11] for a detailed explanation of this technique. Luckily, the factor-revealing LP that we get is the same as the one derived in [19]. This allows us to use their bound of 1.861 on the solution of the factor-revealing LP.

Consider a star S consisting of a facility p and k clients numbered 1 through k. Let d_j denote the connection cost between facility p and client j, and α_j denote the share of j in the total cost (see the definition of α in the algorithm). The cost of the star is $f_p(k) + \sum_{i=1}^{k} d_i$. For simplicity, we set $f = f_p(k)$. Without loss of generality, we assume $\alpha_1 \leq \alpha_2 \leq \cdots \leq \alpha_k$.

Lemma 4 *For every $i = 1, \ldots, k$, we have $\sum_{j=i}^{k} \max(\alpha_i - d_j, 0) \leq f$.*

Proof: Consider the time just before i gets connected for the first time, i.e., when $t = \alpha_i - \epsilon$ for very small ϵ. We call this the *critical time* for client i. At this time, the amount client j offers to facility f is $\max(t - d_j, 0)$ for $j \geq i$. From the algorithm and the discussion after that, the total offer of clients $i, \ldots, k$ to facility p can never exceed the incremental cost of facility p for serving $k - i + 1$ clients. Thus if l denotes the current level of p, then the total offer of clients $i, \ldots, k$ is at most $f_p(l + k - i + 1) - f_p(l) \leq f_p(k - i + 1)$ (by concavity) $\leq f_p(k) = f$, since otherwise all these clients would have been assigned to the facility p at an earlier time. We note that the last inequality holds since f_p is non-decreasing. This completes the proof of the lemma. □

So far, we have not used the triangle inequality of connection costs. We use this assumption in the next lemma.

Lemma 5 *For all clients i and j such that $j < i$, we have $\alpha_i \leq \alpha_j + d_i + d_j$.*

Proof: Consider the critical time for client i ($t = \alpha_i - \epsilon$). If j is not connected to a facility at time t, $\alpha_i = \alpha_j$; otherwise let p' be the facility that j is connected to at this time. Also, let l denote the value of $\text{level}_{p'}$ right after client j gets connected to p'. By concavity of $f_{p'}$, at any time after this, the cost of serving a new client by p' is at most $f_{p'}(l) - f_{p'}(l-1)$. Therefore, we must have $t \leq c_{p',i} + f_{p'}(l) - f_{p'}(l-1)$, since otherwise i should have been connected at a time earlier than t to p'. Therefore,

$$\alpha_i \leq c_{p',i} + f_{p'}(l) - f_{p'}(l-1). \tag{1}$$

Now, consider the critical time for j, i.e., just before it gets connected to p', and let q denote the value of $\text{level}_{p'}$ at this moment. This means that there are $l - q$ new clients $b_1, \ldots, b_{l-q}$ that are connected to p' at time α_j. For simplicity, let b_{l-q} be j. First, notice that at the critical time for j we have $\sum_{i=1}^{l-q} \text{offer}(b_i, p') = f_{p'}(l) - f_{p'}(q)$. Also, we must have $\sum_{i=1}^{l-q-1} \text{offer}(b_i, p') \leq f_{p'}(l-1) - f_{p'}(q)$, since otherwise the clients $b_1, \ldots, b_{l-q-1}$ will be connected to p' at an earlier time. Combining the above inequalities, we obtain $\text{offer}(j,p') \geq f_{p'}(l) - f_{p'}(l-1)$, which implies

$$\alpha_j - c_{p',j} \geq f_{p'}(l) - f_{p'}(l-1). \tag{2}$$

Inequalities (1) and (2) imply $\alpha_i \leq c_{p',i} + \alpha_j - c_{p',j}$. By the triangle inequality, $c_{p',i} \leq c_{p',j} + d_i + d_j$, and therefore $\alpha_i \leq \alpha_j + d_i + d_j$. □

The above lemmas imply that the maximum ratio of $\sum \alpha_i$ to the cost of S can be bounded by the solution of the following maximization program.

$$\begin{array}{ll} \text{maximize} & \dfrac{\sum_{i=1}^k \alpha_i}{f + \sum_{i=1}^k d_i} \\ \text{such that} & \forall\, 1 \leq i < k: \ \alpha_i \leq \alpha_{i+1} \\ & \forall\, 1 \leq j < i \leq k: \ \alpha_i \leq \alpha_j + d_i + d_j \\ & \forall\, 1 \leq i \leq k: \ \displaystyle\sum_{j=i}^{k} \max(\alpha_i - d_j, 0) \leq f \\ & \forall\, 1 \leq j \leq i \leq k: \ \alpha_j, d_j, f \geq 0 \end{array} \tag{3}$$

Therefore, as in [19], if we let γ be $\sup_k\{z_k\}$, where z_k is the solution of the above maximization program, then Algorithm A is a γ-approximation algorithm for the concave facility location problem. This linear program is the same as the one derived in [19], and therefore we can use the following result of [19].

Lemma 6 *For every k, the solution of the maximization program (3) is less than 1.861.*

This, together with the above discussion implies the following.

Theorem 7 *Algorithm A is a 1.861-approximation algorithm for the concave metric facility location problem.*

5 Generalizations

In this section, we consider more general variants of the problem. So far, we have shown approximation algorithms for the concave facility location problem. In this section, we consider the generalized facility location problem with more general facility cost functions. First, we consider the problem with non-metric connection costs. For this case, we show that the algorithm of Section 3 has an approximation factor of $\ln(n)$.

Theorem 8 *Algorithm A is a $\ln(n)$-approximation algorithm for the concave non-metric facility location problem.*

Proof: Recall the definitions of α_j, d_j, and f in Section 4. By Lemma 4, we have $\sum_{j=i}^{k}(\alpha_i - d_j) \leq f$ (notice that the metric assumption was not used in the proof of Lemma 4). Thus,

$$\alpha_i \leq \frac{1}{k-i+1}(f + \sum_{j=i}^{k} d_j) \leq \frac{1}{k-i+1}(f + \sum_{j=1}^{k} d_j).$$

It follows that

$$\sum_{i=1}^{k} \alpha_i \leq \sum_{i=1}^{k} \frac{1}{k-i+1}(f + \sum_{j=1}^{k} d_j) = H_k(f + \sum_{j=1}^{k} d_j) \leq (\ln n)c_S.$$

□

Now, we give a polynomial time algorithm for the generalized facility location problem with convex facility cost functions. Recall that a function f is called *convex* if for every $x \geq 1$, $f(x+1) - f(x) \geq f(x) - f(x-1)$. Here we assume without loss of generality that $f(0) = 0$ for each facility cost function f.

Theorem 9 *The convex facility location problem can be solved exactly in polynomial time.*

Proof: First we reduce the problem to the capacitated facility location problem with unit hard capacities. Next, we show how this problem can be solved in polynomial time. For each facility $i \in \mathcal{F}$, we place $|\mathcal{D}|$ copies of unit-capacity facilities where f_i^j, the opening cost of the jth copy of facility $i \in \mathcal{F}$, is $f_i(j+1) - f_i(j)$, $0 \leq j \leq |\mathcal{D}| - 1$. The correctness of the reduction follows from the fact that if we use a facility f_i^j, we should use all facilities f_i^k, $k \leq j$, since f_i is convex. We now solve the capacitated facility location problem with unit capacities using an algorithm for the minimum weight maximal matching problem on bipartite graphs. We construct a bipartite graph $G = (X \cup Y, E)$ as follows. For each client j, we place a vertex in the set X, for each facility i, we place a vertex in the set Y, and finally we place an edge $\{j, i\}$ in E between a client j and a facility i of weight $c_{ij} + f_i$. We can easily observe that by finding a minimum weight maximal matching [26] in G, one can find an optimal assignment for the original problem. □

The problem of finding a constant-factor approximation algorithm for general facility cost functions is open. However, it is easy to observe that if each facility cost function f_i can be approximated by a concave function g_i, i.e., there is a function g_i such that for every x, $g_i(x)/c \leq f_i(x) \leq g_i(x)$, then applying an approximation algorithm A for the concave facility location problem on an instance with cost functions g_i would give a solution with a factor at most c times the approximation factor of A.

6 Conclusions and open problems

The main contribution of this paper is the introduction of the generalized facility location problem and several interesting open questions related to this problem. One interesting open problem is to obtain a constant factor approximation for the generalized facility location problem with metric connection costs. Since the natural LP formulation of this problem has unbounded integrality gap, it is natural to consider a local search approach or a different LP formulation. Another open question is to find an $O(\ln n)$ approximation algorithm for the generalized facility location problem with non-metric connection costs. We observed that this is possible when the facility costs are concave. Also, it is well known that the non-metric uncapacitated (and therefore generalized) facility location cannot be approximated within a factor better than $\ln(n)$.

Acknowledgments. We would like to thank Tom Leighton, Rajmohan Rajaraman and Ravi Sundaram for introducing the problem. We especially thank Jiawei Zhang for his comments on an earlier draft of this paper. Finally, we would like to thank the anonymous referees for their helpful suggestions.

References

[1] V. Arya, N. Garg, R. Khandekar, K. Munagala, and V. Pandit, Local search heuristics for k-median and facility location problems, In *Proceedings of the 33rd ACM Symposium on Theory of Computing* 2001, 21-29.

[2] M. Charikar and S. Guha, Improved combinatorial algorithms for facility location and k-median problems, In *Proceedings of the 40th IEEE Symposium on Foundations of Computer Science* 1999, 378-388.

[3] F.A. Chudak, Improved approximation algorithms for uncapacited facility location, In R.E. Bixby, E.A. Boyd, and R.Z. Ríos-Mercado, editors, *Integer Programming and Combinatorial Optimization* 1998, volume 1412 of *Lecture Notes in Computer Science*, 180-194.

[4] F.A. Chudak and D. Shmoys, Improved approximation algorithms for the uncapacitated facility location problem, Unpublished manuscript 1998.

[5] F.A. Chudak and D. Shmoys, Improved approximation algorithms for the capacitated facility location problem, In *Proceedings of 10th ACM-SIAM Symposium on Discrete Algorithms* 1999, 875-876.

[6] F.A. Chudak and D.P. Williamson, Improved approximation algorithms for capacitated facility location problems, In *Integer Programming and Combinatorial Optimization* 1999, volume 1610 of *Lecture Notes in Computer Science*, 99-113.

[7] V. Chvátal, A greedy heuristic for the set-covering problem, *Math. Oper. Res.* **4** (1979), 233–235.

[8] M. Goemans and J. Kleinberg, An improved approximation ratio for the minimum latency problem, *Mathematical Programming* **82** (1998), 111–124.

[9] S. Guha and S. Khuller, Greedy strikes back: Improved facility location algorithms, *Journal of Algorithms* **31** (1999), 228–248.

[10] S. Guha, A. Meyerson, and K. Munagala, Hierarchical placement and network design problems, In *Proceedings of the 41st IEEE Symposium on Foundations of Computer Science* 2000, 603–612.

[11] K. Jain, M. Mahdian, E. Markakis, A. Saberi, and V.V. Vazirani, Approximation algorithms for facility location via dual fitting with factor-revealing LP, To appear in the Journal of the ACM.

[12] K. Jain, M. Mahdian, and A. Saberi, A new greedy approach for facility location problem, In *Proceedings of the 34rd ACM Symposium on Theory of Computing* 2002, 731–740.

[13] Kamal Jain and Vijay V. Vazirani, Approximation algorithms for metric facility location and k-median problems using the primal-dual schema and Lagrangian relaxation, *Journal of the ACM* **48** (2001), 274–296.

[14] D. R. Karger and M. Minkoff, Building Steiner trees with incomplete global knowledge, In *Proceedings of the 41st IEEE Symposium on Foundations of Computer Science* 2000, 613–623.

[15] Madhukar R. Korupolu, C. Greg Plaxton, and Rajmohan Rajaraman, Analysis of a local search heuristic for facility location problems, *Journal of Algorithms* **37** (2000), 146–188.

[16] A.A. Kuehn and M.J. Hamburger, A heuristic program for locating warehouses, *Management Science* **9** (1963), 643–666.

[17] T. Leighton, Internet Research Seminars at MIT, February 2002.

[18] J.H. Lin and J.S. Vitter, ϵ-approximations with minimum packing constraint violation, In *Proceedings of the 24th ACM Symposium on Theory of Computing* 1992, 771–782.

[19] M. Mahdian, E. Markakis, A. Saberi, and V.V. Vazirani, A greedy facility location algorithm analyzed using dual fitting, In *Proceedings of 5th International Workshop on Randomization and Approximation Techniques in Computer Science* 2001, volume 2129 of *Lecture Notes in Computer Science*, 127–137.

[20] M. Mahdian, Y. Ye, and J. Zhang, Improved approximation algorithms for metric facility location problems, In *The 5th International Workshop on Approximation Algorithms for Combinatorial Optimization* 2002, 229–242.

[21] R.J. McEliece, E.R. Rodemich, H. Rumsey Jr., and L.R. Welch, New upper bounds on the rate of a code via the Delsarte-Macwilliams inequalities, *IEEE Transactions on Information Theory* **23** (1977), 157–166.

[22] M. Pal, E. Tardos, and T. Wexler, Facility location with hard capacities, *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science* 2001, 329–338.

[23] D.B. Shmoys, Approximation algorithms for facility location problems, In K. Jansen and S. Khuller, editors, *Approximation Algorithms for Combinatorial Optimization* 2000, volume 1913 of *Lecture Notes in Computer Science*, 27–33.

[24] D.B. Shmoys, E. Tardos, and K.I. Aardal, Approximation algorithms for facility location problems, In *Proceedings of the 29th ACM Symposium on Theory of Computing* 1997, 265–274.

[25] M. Sviridenko, An 1.582-approximation algorithm for the metric uncapacitated facility location problem, In *The Ninth Conference on Integer Programming and Combinatorial Optimization* 2002, 240–257.

[26] Douglas B. West, *Introduction to Graph Theory*, Prentice Hall, Upper Saddle River, NJ, 1996.

[27] J. Zhang, Personal communication.