

Greedy Facility Location Algorithms Analyzed using Dual Fitting with Factor-Revealing LP

Kamal Jain*

Mohammad Mahdian[†]

Evangelos Markakis[‡]

Amin Saberi[‡]

Vijay V. Vazirani[‡]

Abstract

In this paper, we will formalize the method of dual fitting and the idea of factor-revealing LP. This combination is used to design and analyze two greedy algorithms for the metric uncapacitated facility location problem. Their approximation factors are 1.861 and 1.61, with running times of $O(m \log m)$ and $O(n^3)$, respectively, where n is the total number of vertices and m is the number of edges in the underlying complete bipartite graph between cities and facilities. The algorithms are used to improve recent results for several variants of the problem.

1 Introduction

A large fraction of the theory of approximation algorithms, as we know it today, is built around the theory of linear programming, which offers the two fundamental algorithm design techniques of rounding and the primal–dual schema (see [44]). Interestingly enough, the LP-duality based analysis [30, 10] for perhaps the most central problem of this theory, the set cover problem, did not use either of these techniques. Moreover, the analysis used for set cover does not seem to have found use outside of this problem and its generalizations [38], leading to a somewhat unsatisfactory state of affairs.

In this paper¹, we formalize the technique used for analyzing set cover as the *method of dual fitting*, and we also introduce the idea of using a *factor-revealing LP*. Using this combination we analyze two greedy algorithms for the metric uncapacitated facility location problem. Their approximation factors are 1.861 and 1.61, with running times of $O(m \log m)$ and $O(n^3)$ respectively, where m and n denote the total number of edges and vertices in the underlying complete bipartite graph between cities and facilities. In other words, $m = n_c \times n_f$ and $n = n_c + n_f$, where n_c is the number of cities and n_f is the number of facilities.

1.1 Dual fitting with factor-revealing LP

The set cover problem offers a particularly simple setting for illustrating most of the dominant ideas in approximation algorithms (see [44]). Perhaps the reason that the method of dual fitting was not

*Microsoft Research, One Microsoft Way, Redmond, WA 98052. E-mail: kamalj@microsoft.com

[†]Department of Mathematics, MIT, Cambridge, MA 02139, USA. E-mail: mahdian@math.mit.edu.

[‡]College of Computing, Georgia Tech, Atlanta, GA 30332, USA. E-mails: {vangelis, saberi, vazirani}@cc.gatech.edu.

¹This paper is based on the preliminary versions [31] and [21].

clear so far was that the set cover problem did not require its full power. However, in retrospect, its salient features are best illustrated again in the simple setting of the set cover problem – we do this in Section 9.

The method of dual fitting can be described as follows, assuming a minimization problem: The basic algorithm is combinatorial – in the case of set cover it is in fact a simple greedy algorithm. Using the linear programming relaxation of the problem and its dual, one first interprets the combinatorial algorithm as a primal-dual-type algorithm – an algorithm that is iteratively making primal and dual updates. Strictly speaking, this is not a primal-dual algorithm, since the dual solution computed is, in general, infeasible (see Section 9 for a discussion on this issue). However, one shows that the primal integral solution found by the algorithm is fully paid for by the dual computed. By *fully paid for* we mean that the objective function value of the primal solution is bounded by that of the dual. The main step in the analysis consists of dividing the dual by a suitable factor, say γ , and showing that the shrunk dual is feasible, i.e., it *fits* into the given instance. The shrunk dual is then a lower bound on OPT, and γ is the approximation guarantee of the algorithm.

Clearly, we need to find the minimum γ that suffices. Equivalently, this amounts to finding the worst possible instance – one in which the dual solution needs to be shrunk the most in order to be rendered feasible. For each value of n_c , the number of cities, we define a factor-revealing LP that encodes the problem of finding the worst possible instance with n_c cities as a linear program. This gives a family of LP’s, one for each value of n_c . The supremum of the optimal solutions to these LP’s is then the best value for γ . In our case, we do not know how to compute this supremum directly. Instead, we obtain a feasible solution to the dual of each of these LP’s. An upper bound on the objective function values of these duals can be computed, and is an upper bound on the optimal γ . In our case, this upper bound is 1.861 for the first algorithm and 1.61 for the second one. In order to get a closely matching tight example, we numerically solve the factor-revealing LP for a large value of n_c .

The technique of factor-revealing LPs is similar to the idea of LP bounds in coding theory. LP bounds give the best known bounds on the minimum distance of a code with a given rate by bounding the solution of a linear program. (cf. McEliece et al. [33]). In the context of approximation algorithms, Goemans and Kleinberg [12] use a similar method in the analysis of their algorithm for the minimum latency problem.

1.2 The facility location problem

In the (uncapacitated) facility location problem, we have a set \mathcal{F} of n_f *facilities* and a set \mathcal{C} of n_c *cities*. For every facility $i \in \mathcal{F}$, a nonnegative number f_i is given as the *opening cost* of facility i . Furthermore, for every facility $i \in \mathcal{F}$ and city $j \in \mathcal{C}$, we have a *connection cost* (a.k.a. service cost) c_{ij} between facility i and city j . The objective is to open a subset of the facilities in \mathcal{F} , and connect each city to an open facility so that the total cost is minimized. We will consider the *metric* version of this problem, i.e., the connection costs satisfy the triangle inequality.

This problem has occupied a central place in operations research since the early 60’s [3, 26, 28, 40, 41], and has been studied from the perspectives of worst case analysis, probabilistic analysis, polyhedral combinatorics and empirical heuristics (see [11, 35]). Although the first approximation algorithm for this problem, a greedy algorithm achieving a guarantee of $O(\log n)$ in the general (non-metric) case due to Hochbaum [20], dates back to almost 20 years ago, renewed interest in recent years has resulted

in much progress. Recently, the problem has found several new applications in network design problems such as placement of routers and caches [16, 29], agglomeration of traffic or data [1, 17], and web server replications in a content distribution network (CDN) [25, 36, 37].

The first constant factor approximation algorithm for this problem was given by Shmoys, Tardos, and Aardal [39]. Later, the factor was improved by Chudak and Shmoys [9] to $1 + 2/e$. Both these algorithms were based on LP-rounding, and therefore had high running times.

Jain and Vazirani [22] gave a primal–dual algorithm, achieving a factor of 3, and having the same running time as ours (we will refer to this as the JV algorithm). Their algorithm was adapted for solving several related problems such as the fault-tolerant and outlier versions, and the k -median problem [22, 23, 8]. Mettu and Plaxton [34] used a restatement of the JV algorithm for the on-line median problem.

Strategies based on local search and greedy improvement for facility location problem have also been studied. The work of Korupolu et al. [27] shows that a simple local search heuristic proposed by Kuehn and Hamburger [28] yields a $(5 + \epsilon)$ -approximation algorithm with a running time of $O(n^6 \log n / \epsilon)$, for any $\epsilon > 0$. Charikar and Guha [6] improved the factor slightly to 1.728 by combining the JV algorithm, greedy augmentation, and the LP-based algorithm [9]. They also combined greedy improvement and cost scaling to improve the factor of the JV algorithm to 1.853. For a metric defined by a sparse graph, Thorup [43] has obtained a $(3 + o(1))$ -approximation algorithm with running time $\tilde{O}(|E|)$. Regarding hardness results, Guha and Khuller [15] showed that the best approximation factor possible for this problem is 1.463, assuming $NP \not\subseteq DTIME[n^{O(\log \log n)}]$.

Since the publication of the first draft of the present paper, two new algorithms have been proposed for the facility location problem. The first algorithm, due to Sviridenko [42], uses the LP-rounding method to achieve an approximation factor of 1.58. The second algorithm, due to Mahdian, Ye, and Zhang [32], combines our second algorithm with the idea of cost scaling to achieve an approximation factor of 1.52, which is currently the best known factor for this problem.

1.3 Our results

Our first algorithm is quite similar to the greedy set cover algorithm: iteratively pick the most cost-effective choice at each step, where cost-effectiveness is measured as the ratio of the cost incurred to the number of new cities served. In order to use LP-duality to analyze this algorithm, we give an alternative description which can be seen as a modification of the JV algorithm – when a city gets connected to an open facility, it withdraws whatever it has contributed towards the opening cost of other facilities. This step of withdrawing contribution is important, since it ensures that the primal solution is fully paid for by the dual.

The second algorithm has a minor difference with the first one: A city might change the facility to which it is connected and connect to a closer facility. If so, it offers this difference toward opening the latter facility.

The approximation factor of the algorithms are 1.861 and 1.61, with running times of $O(m \log m)$ and $O(n^3)$ respectively where n is the total number of vertices and m is the number of edges in the underlying complete bipartite graph between cities and facilities.

We have experimented our algorithms on randomly generated instances as well as instances obtained from the Operations Research library [4] and GT-ITM Internet topology generator [45]. The cost of the integral solution found is compared against the solution of the LP-relaxation of the problem,

rather than OPT (computing which would be prohibitively time consuming). The results are encouraging: The average error of our algorithms is about 3% and 1% respectively, and is a significant improvement over the JV algorithm which has an error of even 100% in some cases.

The primal-dual algorithm of Jain and Vazirani [22] is versatile in that it can be used to obtain algorithms for many variants of the facility location problem, such as k -median [22], a common generalization of k -median and facility location [22], capacitated facility location with soft capacities [22], prize collecting facility location [8], and facility location with outliers [8]. In Section 8, we apply our algorithms to several variants of the problem. First, we consider a common generalization of the facility location and k -median problems. In this problem, which we refer to as the k -facility location problem, an instance of the facility location problem and an integer k are given and the objective is to find the cheapest solution that opens at most k facilities. The k -median problem is a special case of this problem in which all opening costs are 0. The k -median problem is studied extensively [2, 6, 7, 22] and the best known approximation algorithm for this problem, due to Arya et al. [2], achieves a factor of $3 + \epsilon$. The k -facility location problem has also been studied in operations research [11], and the best previously known approximation factor for this problem was 6 [22].

Next, we show an application of our algorithm to the facility location game. We also use our algorithm to improve recent results for some other variants of the problem. In the facility location problem with outliers we are not required to connect all cities to open facilities. We consider two versions of this variant: In the robust version, we are allowed to leave l cities unconnected. In facility location with penalties we can either connect a city to a facility, or pay a specified penalty. Both versions were motivated by commercial applications, and were proposed by Charikar et al. [8]. In this paper we will modify our algorithm to obtain a factor 2 approximation algorithm for these versions, improving the best known result of factor 3 [8].

In the fault tolerant variant, each city has a specified number of facilities it should be connected to. This problem was proposed in [23] and the best factor known is 2.47 [18]. We can achieve a factor of 1.61 when all cities have the same connectivity requirement. In addition, we introduce a new variant which can be seen as a special case of the concave cost version of this problem: the cost of opening a facility at a location is specified and it can serve exactly one city. In addition, a *setup cost* is charged the very first time a facility is opened at a given location.

2 Algorithm 1

In the following algorithm we use a notion of cost effectiveness. Let us say that a *star* consists of one facility and several cities. The cost of a star is the sum of the opening cost of the facility and the connection costs between the facility and all the cities in the star. More formally, the cost of the star (i, C') , where i is a facility and $C' \subseteq C$ is a subset of cities, is $f_i + \sum_{j \in C'} c_{ij}$. The cost effectiveness of the star (i, C') is the ratio of the cost of the star to the size of C' , i.e., $(f_i + \sum_{j \in C'} c_{ij}) / |C'|$.

Algorithm 1

1. Let U be the set of unconnected cities. In the beginning, all cities are unconnected i.e. $U := C$ and all facilities are unopened.
2. While $U \neq \emptyset$:

- Among all stars, find the most cost-effective one, (i, C') , open facility i , if it is not already open, and connect all cities in C' to i .
- Set $f_i := 0$, $U := U \setminus C'$.

Note that a facility can be chosen again after being opened, but its opening cost is counted only once since we set f_i to zero after the first time the facility is picked by the algorithm. As far as cities are concerned, every city j is removed from C , when connected to an open facility, and is not taken into consideration again. Also, notice that although the number of stars is exponentially large, in each iteration the most cost-effective pair can be found in polynomial time. For each facility i , we can sort the cities in increasing order of their connection cost to i . It can be easily seen that the most cost-effective star will consist of a facility and a set, containing the first k cities in this order, for some k .

The idea of cost effectiveness essentially stems from a similar notion in the greedy algorithm for the set cover problem. In that algorithm, the cost effectiveness of a set S is defined to be the cost of S over the number of uncovered elements in S . In each iteration, the algorithm picks the most cost-effective set until all elements are covered. The most cost-effective set can be found either by using direct computation, or by using the dual program of the linear programming formulation for the problem. The dual program can also be used to prove the approximation factor of the algorithm. Similarly, we will use the LP-formulation of facility location to analyze our algorithm. As we will see, the dual formulation of the problem helps us to understand the nature of the problem and the greedy algorithm.

The facility location problem can be captured by an integer program due to Balinski [3]. For the sake of convenience, we give another equivalent formulation for the problem. Let \mathcal{S} be the set of all stars. The facility location problem can be thought of as picking a minimum cost set of stars such that each city is in at least one star. This problem can be captured by the following integer program. In this program, x_S is an indicator variable denoting whether star S is picked and c_S denotes the cost of star S .

$$\begin{aligned}
& \text{minimize} && \sum_{S \in \mathcal{S}} c_S x_S && (1) \\
& \text{subject to} && \forall j \in \mathcal{C} : \sum_{S: j \in S} x_S \geq 1 \\
& && \forall S \in \mathcal{S} : x_S \in \{0, 1\}
\end{aligned}$$

The LP-relaxation of this program is:

$$\begin{aligned}
& \text{minimize} && \sum_{S \in \mathcal{S}} c_S x_S && (2) \\
& \text{subject to} && \forall j \in \mathcal{C} : \sum_{S: j \in S} x_S \geq 1 \\
& && \forall S \in \mathcal{S} : x_S \geq 0
\end{aligned}$$

The dual program is:

$$\begin{aligned}
& \text{maximize} && \sum_{j \in \mathcal{C}} \alpha_j && (3)
\end{aligned}$$

$$\begin{aligned} \text{subject to } \quad & \forall S \in \mathcal{S} : \sum_{j \in S \cap \mathcal{C}} \alpha_j \leq c_S \\ & \forall j \in \mathcal{C} : \alpha_j \geq 0 \end{aligned}$$

There is an intuitive way of interpreting the dual variables. We can think of α_j as the contribution of city j , or its share toward the total expenses. Note that the first inequality of the dual can also be written as $\sum_{j \in C} \max(0, \alpha_j - c_{ij}) \leq f_i$ for every facility i . We can now see how the dual variables can help us find the most cost-effective star in each iteration of the greedy algorithm: if we start raising the dual variables of all unconnected cities simultaneously, the most cost-effective star will be the first star (i, C') for which $\sum_{j \in C'} \max(0, \alpha_j - c_{ij}) = f_i$. Hence we can restate Algorithm 1 based on the above observation. This is in complete analogy to the greedy algorithm and its restatement using LP-formulation for set-cover.

Restatement of Algorithm 1

1. We introduce a notion of time, so that each event can be associated with the time at which it happened. The algorithm starts at time 0. Initially, each city is defined to be unconnected ($U := C$), all facilities are unopened, and α_j is set to 0 for every j .
2. While $U \neq \emptyset$, increase the time, and simultaneously, for every city $j \in U$, increase the parameter α_j at the same rate, until one of the following events occurs (if two events occur at the same time, we process them in arbitrary order).
 - (a) For some unconnected city j , and some open facility i , $\alpha_j = c_{ij}$. In this case, connect city j to facility i and remove j from U .
 - (b) For some unopened facility i , we have $\sum_{j \in U} \max(0, \alpha_j - c_{ij}) = f_i$. This means that the total contribution of the cities is sufficient to open facility i . In this case, open this facility, and for every unconnected city j with $\alpha_j \geq c_{ij}$, connect j to i , and remove it from U .

In each iteration of algorithm 1 the process of opening a facility and/or connecting some cities will be defined as an *event*. It is easy to prove the following lemma by induction.

Lemma 1 *The sequence of events executed by Algorithm 1 and its restatement are identical.*

Proof: By induction. □

This restatement can also be seen as a modification of JV algorithm [22]. The only difference is that in JV algorithm cities, when connected to an open facility, are not excluded from U , hence they might contribute towards opening several facilities. Due to this fact they have a second cleanup phase in which some of the already open facilities will be closed down.

Also, it is worth noting that despite the similarity between Algorithm 1 and Hochbaum's greedy algorithm for facility location (which is equivalent to the set cover algorithm applied on the set of stars), they are not equivalent. This is because we set f_i to zero after picking a set containing f_i . As the following example shows, the approximation factor of Hochbaum's algorithm is $\Omega(\frac{\log n}{\log \log n})$ on instances with metric inequality: Consider k facilities with opening cost p^k located in the same place Also $k - 1$ groups of cities S_1, S_2, \dots, S_{k-1} . The group S_i consists of p^{k-i+1} cities with

distance $\sum_{j=1\dots i} p^{j-1}$ from the facilities. Other distances are obtained from the triangle inequality. Hochbaum's algorithm opens all facilities and therefore its solution costs more than kp^k . The optimum solution is $p^k + \sum_{i=1\dots k-1} \sum_{j=1\dots i} p^{j-1}$. It is easy to show that with a careful choice of k , the ratio of these two expressions is $\Omega(\frac{\log n}{\log \log n})$. We do not know whether the approximation factor of Hochbaum's algorithm on metric instances is strictly less than $\log n$ or not.

3 Analysis of Algorithm 1

In this section we will give an LP-based analysis of the algorithm. As stated before, the contribution of each city goes towards opening at most one facility and connecting the city to an open facility. Therefore, the total cost of the solution produced by our algorithm will be equal to the sum $\sum_j \alpha_j$ of the contributions. However, α is not a feasible dual solution as it was in JV algorithm. The reason is that in every iteration of the restatement of Algorithm 1, we exclude a subset of cities and withdraw their contribution from all facilities. So at the end, for some facility i , $\sum_j \max(\alpha_j - c_{ij}, 0)$ can be greater than f_i and hence the corresponding constraints of the dual program is violated.

However, if we find an γ for which α/γ is feasible, $\sum_j \alpha_j/\gamma$ would be a lower bound to the optimum and therefore the approximation factor of the algorithm would be at most γ . This observation motivates the following definition.

Definition Given α_j ($j = 1, \dots, n_c$), a facility i is called at most γ -overtight if and only if

$$\sum_j \max(\alpha_j/\gamma - c_{ij}, 0) \leq f_i.$$

Using the above definition, it is trivial that α/γ is a feasible dual if and only if each facility is at most γ -overtight. Now, we want to find such an γ . Note that in the above sum we only need to consider the cities j for which $\alpha_j \geq \gamma c_{ij}$. Let us assume without loss of generality that it is the case only for the first k cities. Moreover, assume without loss of generality that $\alpha_1 \leq \alpha_2 \leq \dots \leq \alpha_k$. The next two lemmas express the constraints on α imposed by the problem or our algorithm. The first lemma mainly captures metric property and the second one expresses the fact that the total contribution offered to a facility at any time during the algorithm is no more than its cost.

Lemma 2 For every two cities j, j' and facility i , $\alpha_j \leq \alpha_{j'} + c_{ij'} + c_{ij}$.

Proof: If $\alpha_{j'} \geq \alpha_j$, the inequality obviously holds. Assume $\alpha_j > \alpha_{j'}$. Let i' be the facility that city j' is connected to by our algorithm. Thus, facility i' is open at time $\alpha_{j'}$. The contribution α_j cannot be greater than $c_{i'j}$ because in that case city j could be connected to facility i' at some time $t < \alpha_j$. Hence $\alpha_j \leq c_{i'j}$. Furthermore, by triangle inequality, $c_{i'j} \leq c_{i'j'} + c_{ij'} + c_{ij} \leq \alpha_{j'} + c_{ij'} + c_{ij}$. \square

Lemma 3 For every city j and facility i , $\sum_{l=j}^k \max(\alpha_j - c_{il}, 0) \leq f_i$.

Proof: Assume, for the sake of contradiction, that for some j and some i the inequality does not hold, i.e., $\sum_{k=j}^{n_c} \max(\alpha_j - c_{ik}, 0) > f_i$. By the ordering on cities, for $k \geq j$, $\alpha_k \geq \alpha_j$. Let time $t = \alpha_j$. By the assumption, facility i is fully paid for before time t . For any city k , $j \leq k \leq n_c$ for

which $\alpha_j - c_{ik} > 0$ the edge (i, k) must be tight before time t . Moreover, there must be at least one such city. For this city, $\alpha_k < \alpha_j$, since the algorithm will stop growing α_k as soon as k has a tight edge to a fully paid for facility. The contradiction establishes the lemma. \square

Subject to the constraints introduced by Lemmas 2 and 3, we want to find the minimum γ for which $\sum_{j=1}^k (\alpha_j/\gamma - c_{ij}) \leq f_i$. In other words, we want to find the maximum of the ratio $\frac{\sum_{j=1}^k \alpha_j}{f + \sum_{j=1}^k d_j}$. We can define variables f , d_j , and α_j , corresponding to facility cost, distances, and contributions respectively and write the following maximization program:

$$\begin{aligned}
z_k = \quad & \text{maximize} && \frac{\sum_{j=1}^k \alpha_j}{f + \sum_{j=1}^k d_j} \\
& \text{subject to} && \alpha_j \leq \alpha_{j+1} && \forall j \in \{1, \dots, k-1\} \\
& && \alpha_j \leq \alpha_l + d_j + d_l && \forall j, l \in \{1, \dots, k\} \\
& && \sum_{l=j}^k \max(\alpha_j - d_l, 0) \leq f && \forall j \in \{1, \dots, k\} \\
& && \alpha_j, d_j, f \geq 0 && \forall j \in \{1, \dots, k\}
\end{aligned} \tag{4}$$

It's not difficult to prove that z_k (the maximum value of the objective function of program 4) is equal to the optimal solution of the following linear program which we call the *factor-revealing LP*.

$$\begin{aligned}
z_k = \quad & \text{maximize} && \sum_{j=1}^k \alpha_j \\
& \text{subject to} && f + \sum_{j=1}^k d_j \leq 1 \\
& && \alpha_j \leq \alpha_{j+1} && \forall j \in \{1, \dots, k-1\} \\
& && \alpha_j \leq \alpha_l + d_j + d_l && \forall j, l \in \{1, \dots, k\} \\
& && x_{jl} \geq \alpha_j - d_l && \forall j, l \in \{1, \dots, k\} \\
& && \sum_{l=j}^k x_{jl} \leq f && \forall j \in \{1, \dots, k\} \\
& && \alpha_j, d_j, f \geq 0 && \forall j \in \{1, \dots, k\}
\end{aligned} \tag{5}$$

Lemma 4 *Let $\gamma = \sup_{k \geq 1} \{z_k\}$. Every facility is at most γ -overtight*

Proof: Consider facility i . We want to show that $\sum_j \max(\alpha_j/\gamma - c_{ij}, 0) \leq f_i$. Suppose without loss of generality that the subset of cities j such that $\alpha_j \geq \gamma c_{ij}$ is $\{j = 1, 2, \dots, k\}$ for some k . Moreover $\alpha_1 \leq \alpha_2 \leq \dots \alpha_k$. Let $d_j = c_{ij}$, $j = 1, \dots, k$, and $f = f_i$. By Lemmas 2 and 3 it follows immediately that the constraints of program 4 are satisfied. Therefore, α_i, d_i, f constitute a feasible solution of program 4. Consequently $\frac{\sum_{j=1}^k \alpha_j}{f_i + \sum_{j=1}^k c_{ij}} \leq z_k$. \square

By what we said so far, we know that the approximation factor of our algorithm is at most $\sup_{k \geq 1} \{z_k\}$. In the following theorem, we prove, by demonstrating an infinite family of instances, that the approximation ratio of Algorithm 1 is not better than $\sup_{k \geq 1} \{z_k\}$.

Theorem 5 *The approximation factor of our algorithm is precisely $\sup_{k \geq 1} \{z_k\}$.*

Proof: Consider an optimum feasible solution of program 4. We construct an instance of the

facility location problem with k cities and $k + 1$ facilities as follows: The cost of opening facility i is

$$f_i = \begin{cases} 0 & \text{if } 1 \leq i \leq k \\ f & \text{if } i = k + 1 \end{cases}$$

The connection cost between a city j and a facility i is:

$$c_{ij} = \begin{cases} \alpha_j & \text{if } 1 \leq i = j \leq k \\ d_j & \text{if } 1 \leq j \leq k, i = k + 1 \\ d_i + d_j + \alpha_i & \text{otherwise} \end{cases}$$

It is easy to see that the connection costs satisfy the triangle inequality. On this instance, our algorithm connects city 1 to facility 1, then it connects city 2 to facility 2, and finally connects city k to facility k . (The inequality $\sum_{l=j}^k \max(\alpha_j - d_l, 0) \leq f$ guarantees that city i can get connected to facility i before facility $k + 1$). Therefore, the cost of the restatement of Algorithm 1 is equal to $\sum_{j=1}^k c_{jj} + \sum_{i=1}^k f_i = \sum_{j=1}^k \alpha_j = z_k$.

On the other hand, the optimal solution for this instance is to connect all the cities to facility $k + 1$. The cost of this solution is equal to $\sum_{j=1}^k c_{k+1,j} + f_{k+1} = f + \sum_{j=1}^k d_j \leq 1$.

Thus, our algorithm outputs a solution whose cost is at least z_k times the cost of the optimal solution. \square

The only thing that remains is to find an upper bound on $\sup_{k \geq 1} \{z_k\}$. By solving the factor-revealing LP for any particular value of k , we get a lower bound on the value of γ . In order to prove an upper bound on γ , we need to present a general solution to the dual of the factor-revealing LP. Unfortunately, this is not an easy task in general. (For example, performing a tight asymptotic analysis of the LP bound is still an open question in coding theory). However, here empirical results can help us: we can solve the dual of the factor-revealing LP for small values of k to get an idea of how the general optimal solution looks like. Using this, it is usually possible (although sometimes tedious) to prove a close-to-optimal upper bound on the value of z_k . We have used this technique to prove an upper bound of 1.861 on γ .

Lemma 6 For every $k \geq 1$, $z_k \leq 1.861$.

Proof: Let $r = 1.8609$. By doubling a feasible solution of 4 it is easy to show that $z_k \leq z_{2k}$ so we can assume, without loss of generality that k is sufficiently large. Consider a feasible solution of the program 4. It is clear from the third inequality that for every j, j' we have

$$\sum_{i=j}^{j'} (\alpha_j - d_i) \leq f. \tag{6}$$

Now, we define l_j and θ_j as follows:

$$l_j = \begin{cases} p_2 k & \text{if } j \leq p_1 k \\ k & \text{if } j > p_1 k \end{cases}$$

$$\theta_j = \begin{cases} \frac{r+1}{p_2 k} & \text{if } j \leq p_1 k \\ \frac{(r+1)(p_2 - p_1)}{p_2(1-p_1)k} & \text{if } p_1 k < j \leq p_2 k \\ 0 & \text{if } j > p_2 k \end{cases}$$

where $p_1 = 0.1991$ and $p_2 = 0.5696$. We consider Inequality 6 for every $j \leq p_2k$ and $j' = l_j$, and multiply both sides of this inequality by θ_j . By adding up all these inequalities, we obtain

$$\sum_{j=1}^{p_1k} \sum_{i=j}^{p_2k} \theta_j(\alpha_j - d_i) + \sum_{j=p_1k+1}^{p_2k} \sum_{i=j}^k \theta_j(\alpha_j - d_i) \leq \left(\sum_{j=1}^{p_2k} \theta_j \right) f. \quad (7)$$

The coefficient of f in the right-hand side of the above inequality is equal to $\sum_{j=1}^{p_2k} \theta_j = \frac{r+1}{p_2k} p_1k + \frac{(r+1)(p_2-p_1)}{p_2(1-p_1)k} (p_2k - p_1k) \approx 1.8609 < 1.861$. Also, the coefficients of α_j and d_j in the left-hand side of Inequality 7 are equal to

$$\text{coeff}[\alpha_j] = \begin{cases} (p_2k - j + 1)\theta_j & j \leq p_1k \\ (k - j + 1)\theta_j & j > p_1k \end{cases} \quad (8)$$

$$\text{coeff}[d_j] = \begin{cases} \sum_{i=1}^j \theta_i & j \leq p_2k \\ \sum_{i=p_1k+1}^j \theta_i & j > p_2k \end{cases} \quad (9)$$

Notice that the sum of coefficients of α_j 's is equal to

$$\begin{aligned} \sum_{j=1}^k \text{coeff}[\alpha_j] &= \sum_{j=1}^{p_1k} \frac{r+1}{p_2k} (p_2k - j + 1) + \sum_{j=p_1k+1}^{p_2k} \frac{(r+1)(p_2-p_1)}{p_2(1-p_1)k} (k - j + 1) \\ &> (r+1) \left(p_1 - \frac{p_1^2}{2p_2} + \frac{(p_2-p_1)^2}{p_2(1-p_1)} - \frac{(p_2-p_1)^2(p_1+p_2)}{2p_2(1-p_1)} \right) k \\ &\approx 1.00004k \\ &> k \end{aligned}$$

Now, we use the inequality $\alpha_i \geq \alpha_j - d_j - d_i$ on the expression on the left hand side of inequality 7 to reduce the coefficients of α_j 's that are greater than 1, and increase the coefficient of α_j 's that are less than 1. Since the sum of these coefficients is greater than k , using this inequality and the inequality $\alpha_j \geq 0$ we can obtain an expression E that is less than or equal to the left hand side of inequality 7, and in which all α_j 's have coefficient 1. The coefficient of d_j in this expression will be equal to its coefficient in the left hand side of inequality 7, plus the absolute value of the change in the coefficient of the corresponding α_j . Therefore, by equations 8 and 9 this coefficient is equal to:

$$\text{coeff}_E[d_j] = \begin{cases} \sum_{i=1}^j \theta_i + |(p_2k - j + 1)\theta_j - 1| & j \leq p_1k \\ \sum_{i=1}^j \theta_i + |(k - j + 1)\theta_j - 1| & p_1k < j \leq p_2k \\ \sum_{i=p_1k+1}^j \theta_i + |(k - j + 1)\theta_j - 1| & j > p_2k \end{cases}$$

If $j \leq p_1k$, we have $(p_2k - j + 1)\theta_j > (p_2k - p_1k) \frac{r+1}{p_2k} = (r+1)(p_2-p_1)/p_2 \approx 1.8609 > 1$ Therefore,

$$\begin{aligned} \text{coeff}_E[d_j] &= \sum_{i=1}^j \theta_i + (p_2k - j + 1)\theta_j - 1 \\ &= r + O\left(\frac{1}{k}\right) \\ &< 1.861 \end{aligned}$$




Figure 1: The approximation ratio of Algorithm 1 is at least 1.5

Similarly, if $p_1k < j \leq p_2k$, we have $(k - j + 1)\theta_j > (k - p_2k)\frac{(r+1)(p_2-p_1)}{p_2(1-p_1)k} = \frac{(r+1)(p_2-p_1)(1-p_2)}{p_2(1-p_1)} \approx 1.00003 > 1$. Therefore,

$$\begin{aligned} \text{coeff}_E[d_j] &= \sum_{i=1}^j \theta_i + (k - j + 1)\theta_j - 1 \\ &= r + O\left(\frac{1}{k}\right) \\ &< 1.861 \end{aligned}$$

Finally, if $j > p_2k$, the coefficient of d_j is equal to

$$\begin{aligned} \text{coeff}_E[d_j] &= \sum_{i=p_1k}^j \theta_i + |0 - 1| \\ &= \frac{(r+1)(p_2-p_1)}{p_2(1-p_1)k}(p_2k - p_1k) + 1 \\ &\approx 1.8609 \\ &< 1.861 \end{aligned}$$

Therefore, in each case, the coefficient of d_j is less than or equal to 1.861. Thus, we have proved that

$$\sum_{j=1}^k \alpha_j - \sum_{j=1}^k 1.861d_j < 1.861f.$$

This clearly implies that $z_k < 1.861$. □

Figure 1 shows a tight example for $k = 2$, for which the approximation factor of the algorithm is 1.5. The cost of the missing edges is given by triangle inequality. Numerical computations using the software CPLEX show that $z_{300} \approx 1.81$. Thus, the approximation factor of our algorithm is between 1.81 and 1.861. We do not know the exact approximation ratio.

4 Algorithm 2

Algorithm 2 is similar to the restatement of Algorithm 1. The only difference is that in Algorithm 1 cities stop offering money to facilities as soon as they get connected to a facility, but here they still offer some money to other facilities. The amount that an already-connected city offers to a facility j is equal to the amount that it would save in connection cost by switching its facility to j . As we will see in the next section, this change reduces the approximation factor of the algorithm from 1.861 to 1.61.

Algorithm 2

1. We introduce a notion of time. The algorithm starts at time 0. At this time, each city is defined to be unconnected ($U := C$), all facilities are unopened, and α_j is set to 0 for every j .
At every moment, each city j offers some money from its contribution to each *unopened* facility i . The amount of this offer is computed as follows: If j is unconnected, the offer is equal to $\max(\alpha_j - c_{ij}, 0)$ (i.e., if the contribution of j is more than the cost that it has to pay to get connected to i , it offers to pay this extra amount to i); If j is already connected to some other facility i' , then its offer to facility i is equal to $\max(c_{i'j} - c_{ij}, 0)$ (i.e., the amount that j offers to pay to i is equal to the amount j would save by switching its facility from i' to i).
2. While $U \neq \emptyset$, increase the time, and simultaneously, for every city $j \in U$, increase the parameter α_j at the same rate, until one of the following events occurs (if two events occur at the same time, we process them in an arbitrary order).
 - (a) For some unopened facility i , the total offer that it receives from cities is equal to the cost of opening i . In this case, we open facility i , and for every city j (connected or unconnected) which has a non-zero offer to i , we connect j to i . The amount that j had offered to i is now called the *contribution* of j toward i , and j is no longer allowed to decrease this contribution.
 - (b) For some unconnected city j , and some open facility i , $\alpha_j = c_{ij}$. In this case, connect city j to facility i and remove j from U .

Clearly the main issue in the facility location problem is to decide which facilities to open. Once this is done, each city should be connected to the closest open facility. Observe that Algorithm 2 makes greedy choices in deciding which facilities to open and once it opens a facility, it does not alter this decision. In this sense, it is also a greedy algorithm.

5 Analysis of Algorithm 2

The following fact should be obvious from the description of Algorithm 2.

Lemma 7 *The total cost of the solution found by Algorithm 2 is equal to the sum of α_j 's.*

Now, as in the analysis of Algorithm 1, we need to find a number γ , such that for every star S , $\sum_{j \in S \cap C} \alpha_j \leq \gamma c_S$. Such a γ will be an upper bound on the approximation ratio of the algorithm, since if for every facility i that is opened in the optimal solution and the collection A of cities that are connected to it, we write the inequality $\sum_{j \in A} \alpha_j \leq \gamma(f_i + \sum_{j \in A} c_{ij})$ and add up these inequalities, we will obtain that the cost of our solution is at most γ times the cost of the optimal solution.

5.1 Deriving the factor-revealing LP

Our proof follows the methodology of Section 3: express various constraints that are imposed by the problem or by the structure of the algorithm as inequalities and get a bound on the value of γ defined above by solving a series of linear programs.

Consider a star S consisting of a facility having opening cost f (with a slight misuse of the notation, we call this facility f), and k cities numbered 1 through k . Let d_j denote the connection cost between facility f and city j , and α_j denote the contribution of the city j at the end of Algorithm 2. We may assume without loss of generality that

$$\alpha_1 \leq \alpha_2 \leq \dots \leq \alpha_k. \quad (10)$$

We need more variables to capture the execution of Algorithm 2. For every i ($1 \leq i \leq k$), consider the situation of the algorithm at time $t = \alpha_i - \epsilon$, where ϵ is very small, i.e., just a moment before city i gets connected for the first time. At this time, each of the cities $1, 2, \dots, i-1$ might be connected to a facility. For every $j < i$, if city j is connected to some facility at time t , let $r_{j,i}$ denote the connection cost between this facility and city j ; otherwise, let $r_{j,i} := \alpha_j$. The latter case occurs if and only if $\alpha_i = \alpha_j$. It turns out that these variables (f , d_j 's, α_j 's, and $r_{j,i}$'s) are enough to write down some inequalities to bound the ratio of the sum of α_j 's to the cost of S (i.e., $f + \sum_{j=1}^k d_j$).

First, notice that once a city gets connected to a facility, its contribution remains constant and it cannot revoke its contribution to a facility, so it can never get connected to another facility with a higher connection cost. This implies that for every j ,

$$r_{j,j+1} \geq r_{j,j+2} \geq \dots \geq r_{j,k}. \quad (11)$$

Now, consider time $t = \alpha_i - \epsilon$. At this time, the amount city j offers to facility f is equal to

$$\begin{aligned} \max(r_{j,i} - d_j, 0) & \quad \text{if } j < i, \text{ and} \\ \max(t - d_j, 0) & \quad \text{if } j \geq i. \end{aligned}$$

Notice that by the definition of $r_{j,i}$ this holds even if $j < i$ and $\alpha_i = \alpha_j$. It is clear from Algorithm 2 that the total offer of cities to a facility can never become larger than the opening cost of the facility. Therefore, for all i ,

$$\sum_{j=1}^{i-1} \max(r_{j,i} - d_j, 0) + \sum_{j=i}^k \max(\alpha_i - d_j, 0) \leq f. \quad (12)$$

The triangle inequality is another important constraint that we need to use. Consider cities i and j with $j < i$ at time $t = \alpha_i - \epsilon$. Let f' be the facility j is connected to at time t . By the triangle inequality and the definition of $r_{j,i}$, the connection cost $c_{f'i}$ between city i and facility f' is at most $r_{j,i} + d_i + d_j$. Furthermore, $c_{f'i}$ can not be less than t , since if it is, our algorithm could have connected the city i to the facility f' at a time earlier than t , which is a contradiction. Here we need to be careful with the special case $\alpha_i = \alpha_j$. In this case, $r_{j,i} + d_i + d_j$ is not more than t . If $\alpha_i \neq \alpha_j$, the facility f' is open at time t and therefore city i can get connected to it, if it can pay the connection cost. Therefore for every $1 \leq j < i \leq k$,

$$\alpha_i \leq r_{j,i} + d_i + d_j. \quad (13)$$

The above inequalities form the following factor-revealing LP.

$$\begin{aligned}
& \text{maximize} && \frac{\sum_{i=1}^k \alpha_i}{f + \sum_{i=1}^k d_i} && (14) \\
& \text{subject to} && \forall 1 \leq i < k : \alpha_i \leq \alpha_{i+1} \\
& && \forall 1 \leq j < i < k : r_{j,i} \geq r_{j,i+1} \\
& && \forall 1 \leq j < i \leq k : \alpha_i \leq r_{j,i} + d_i + d_j \\
& && \forall 1 \leq i \leq k : \sum_{j=\overline{k}}^{i-1} \max(r_{j,i} - d_j, 0) \\
& && \quad + \sum_{j=i} \max(\alpha_i - d_j, 0) \leq f \\
& && \forall 1 \leq j \leq i \leq k : \alpha_j, d_j, f, r_{j,i} \geq 0
\end{aligned}$$

Notice that although the above optimization program is not written in the form of a linear program, it is easy to change it to a linear program by introducing new variables and inequalities.

Lemma 8 *If z_k denotes the solution of the factor-revealing LP, then for every star S consisting of a facility and k cities, the sum of α_j 's of the cities in S in Algorithm 2 is at most $z_k c_S$.*

Proof: Inequalities 10, 11, 12, and 13 derived above imply that the values $\alpha_j, d_j, f, r_{j,i}$ that we get by running Algorithm 2 constitute a feasible solution of the factor-revealing LP. Thus, the value of the objective function for this solution is at most z_k . \square

Lemmas 7 and 8 imply the following.

Lemma 9 *Let z_k be the solution of the factor-revealing LP, and $\gamma := \sup_k \{z_k\}$. Then Algorithm 2 solves the metric facility location problem with an approximation factor of γ .*

5.2 Solving the factor-revealing LP

As mentioned earlier, the optimization program (14) can be written as a linear program. This enables us to use an LP-solver to solve the factor-revealing LP for small values of k , in order to compute the numerical value of γ . Table 1 shows a summary of results that are obtained by solving the factor-revealing LP using CPLEX. It seems from the experimental results that z_k is an increasing sequence that converges to some number close to 1.6 and hence $\gamma \approx 1.6$.

We are using the same idea as Lemma 6 in Section 3 to prove the upper bound of 1.61 on z_k .

Lemma 10 *Let z_k be the solution to the factor-revealing LP. Then for every k , $z_k \leq 1.61$.*

Proof: Using the same argument as in Lemma 6, we can assume, without loss of generality, that k is sufficiently large. Consider a feasible solution of the factor-revealing LP. Let $x_{j,i} := \max(r_{j,i} - d_j, 0)$. The fourth inequality of the factor-revealing LP implies that for every $i \leq i'$,

$$(i' - i + 1)\alpha_i \leq \sum_{j=i}^{i'} d_j + f - \sum_{j=1}^{i-1} x_{j,i}. \tag{15}$$

k	$\max_{i < k} z_i$
10	1.54147
20	1.57084
50	1.58839
100	1.59425
200	1.59721
300	1.59819
400	1.59868
500	1.59898

Table 1: Solution of the factor-revealing LP

Now, we define l_i as follows:

$$l_i = \begin{cases} p_2 k & \text{if } i \leq p_1 k \\ k & \text{if } i > p_1 k \end{cases}$$

where p_1 and p_2 are two constants (with $p_1 < p_2$) that will be fixed later. Consider Inequality 15 for every $i \leq p_2 k$ and $i' = l_i$, and divide both sides of this inequality by $(l_i - i + 1)$. By adding up these inequalities we obtain

$$\sum_{i=1}^{p_2 k} \alpha_i \leq \sum_{i=1}^{p_2 k} \sum_{j=i}^{l_i} \frac{d_j}{l_i - i + 1} + \left(\sum_{i=1}^{p_2 k} \frac{1}{l_i - i + 1} \right) f - \sum_{i=1}^{p_2 k} \sum_{j=1}^{i-1} \frac{x_{j,i}}{l_i - i + 1}. \quad (16)$$

Now for every $j \leq p_2 k$, let $y_j := x_{j,p_2 k}$. The second inequality of the factor-revealing LP implies that $x_{j,i} \geq y_j$ for every $j < i \leq p_2 k$ and $x_{j,i} \leq y_j$ for every $i > p_2 k$. Also, let $\zeta := \sum_{i=1}^{p_2 k} \frac{1}{l_i - i + 1}$. Therefore, inequality 16 implies

$$\sum_{i=1}^{p_2 k} \alpha_i \leq \sum_{i=1}^{p_2 k} \sum_{j=i}^{l_i} \frac{d_j}{l_i - i + 1} + \zeta f - \sum_{i=1}^{p_2 k} \sum_{j=1}^{i-1} \frac{y_j}{l_i - i + 1}. \quad (17)$$

Consider the index $\ell \leq p_2 k$ for which $2d_\ell + y_\ell$ has its minimum (i.e., for every $j \leq p_2 k$, $2d_\ell + y_\ell \leq 2d_j + y_j$). The third inequality of the factor-revealing LP implies that for $i = p_2 k + 1, \dots, k$,

$$\alpha_i \leq r_{\ell,i} + d_i + d_\ell \leq x_{\ell,i} + 2d_\ell + d_i \leq d_i + 2d_\ell + y_\ell. \quad (18)$$

By adding Inequality 18 for $i = p_2 k + 1, \dots, k$ with Inequality 17 we obtain

$$\begin{aligned} \sum_{i=1}^k \alpha_i &\leq \sum_{i=1}^{p_2 k} \sum_{j=i}^{l_i} \frac{d_j}{l_i - i + 1} + (2d_\ell + y_\ell)(1 - p_2)k + \sum_{j=p_2 k + 1}^k d_j - \sum_{i=1}^{p_2 k} \sum_{j=1}^{i-1} \frac{y_j}{l_i - i + 1} + \zeta f \\ &= \sum_{j=1}^{p_2 k} \zeta d_j - \sum_{j=1}^{p_2 k} \sum_{i=j+1}^{p_2 k} \frac{d_j + y_j}{l_i - i + 1} + \sum_{j=p_2 k + 1}^k \left(1 + \sum_{i=p_1 k + 1}^{p_2 k} \frac{1}{k - i + 1} \right) d_j \\ &\quad + (2d_\ell + y_\ell)(1 - p_2)k + \zeta f \\ &\leq \sum_{j=1}^{p_2 k} \zeta d_j + \sum_{j=p_2 k + 1}^k \left(1 + \sum_{i=p_1 k + 1}^{p_2 k} \frac{1}{k - i + 1} \right) d_j + \zeta f \end{aligned}$$

$$+(2d_\ell + y_\ell) \left((1 - p_2)k - \frac{1}{2} \sum_{j=1}^{p_2k} \sum_{i=j+1}^{p_2k} \frac{1}{l_i - i + 1} \right),$$

where the last inequality is a consequence of the inequality $2d_\ell + y_\ell \leq 2d_j + y_j \leq 2d_j + 2y_j$ for $j \leq p_2k$. Now, let $\zeta' := 1 + \sum_{i=p_1k+1}^{p_2k} \frac{1}{k-i+1}$ and $\delta := (1 - p_2) - \frac{1}{2k} \sum_{j=1}^{p_2k} \sum_{i=j+1}^{p_2k} \frac{1}{l_i - i + 1}$. Therefore, the above inequality can be written as follows:

$$\sum_{i=1}^k \alpha_i \leq \sum_{j=1}^{p_2k} \zeta d_j + \sum_{j=p_2k+1}^k \zeta' d_j + \zeta f + \delta(2d_\ell + y_\ell)k, \quad (19)$$

where

$$\zeta = \sum_{i=1}^{p_2k} \frac{1}{l_i - i + 1} = \ln \frac{p_2(1 - p_1)}{(p_2 - p_1)(1 - p_2)} + o(1), \quad (20)$$

$$\zeta' = 1 + \sum_{i=p_1k+1}^{p_2k} \frac{1}{k - i + 1} = 1 + \ln \frac{1 - p_1}{1 - p_2} + o(1), \quad (21)$$

$$\begin{aligned} \delta &= 1 - p_2 - \frac{1}{2k} \sum_{j=1}^{p_2k} \sum_{i=j+1}^{p_2k} \frac{1}{l_i - i + 1} \\ &= \frac{1}{2} \left(2 - p_2 - p_2 \ln \frac{p_2}{p_2 - p_1} - \ln \frac{1 - p_1}{1 - p_2} \right) + o(1). \end{aligned} \quad (22)$$

Now if we choose p_1 and p_2 such that $\delta < 0$, and let $\gamma := \max(\zeta, \zeta')$ then inequality 19 implies that

$$\sum_{i=1}^k \alpha_i \leq (\gamma + o(1)) \left(f + \sum_{i=1}^k d_j \right).$$

Using equations 20, 21, and 22, it is easy to see that subject to the condition $\delta < 0$, the value of γ is minimized when $p_1 \approx 0.439$ and $p_2 \approx 0.695$, which gives us $\gamma < 1.61$. \square

Also, as in the proof of Theorem 5, we can use the optimal solution of the factor-revealing LP that is computed numerically (see Table 1) to construct an example on which our algorithm performs at least z_k times worse than the optimum. These results imply the following.

Theorem 11 *Algorithm 2 solves the facility location problem in time $O(n^3)$, where $n = \max(n_f, n_c)$, with an approximation ratio between 1.598 and 1.61.*

6 The tradeoff between facility and connection costs

We defined the cost of a solution in the facility location problem as the sum of the facility cost (i.e., total cost of opening facilities) and the connection cost. We proved in the previous section that Algorithm 2 achieves an overall performance guarantee of 1.61. However, sometimes it is useful to get different approximation guarantees for facility and connection costs. The following theorem gives such a guarantee. The proof is similar to the proof of Lemma 9.




Figure 2: The tradeoff between γ_f and γ_c

Theorem 12 Let $\gamma_f \geq 1$ and $\gamma_c := \sup_k \{z_k\}$, where z_k is the solution of the following optimization program.

$$\begin{aligned}
& \text{maximize} && \frac{\sum_{i=1}^k \alpha_i - \gamma_f f}{\sum_{i=1}^k d_i} && (23) \\
& \text{subject to} && \forall 1 \leq i < k : \alpha_i \leq \alpha_{i+1} \\
& && \forall 1 \leq j < i < k : r_{j,i} \geq r_{j,i+1} \\
& && \forall 1 \leq j < i \leq k : \alpha_i \leq r_{j,i} + d_i + d_j \\
& && \forall 1 \leq i \leq k : \sum_{j=1}^{i-1} \max(r_{j,i} - d_j, 0) \\
& && \quad + \sum_{j=i}^k \max(\alpha_i - d_j, 0) \leq f \\
& && \forall 1 \leq j \leq i \leq k : \alpha_j, d_j, f, r_{j,i} \geq 0
\end{aligned}$$

Then for every instance \mathcal{I} of the facility location problem, and for every solution SOL for \mathcal{I} with facility cost F_{SOL} and connection cost C_{SOL} , the cost of the solution found by Algorithm 2 is at most $\gamma_f F_{SOL} + \gamma_c C_{SOL}$.

We have computed the solution of the optimization program 23 for $k = 100$, and several values of γ_f between 1 and 3, to get an estimate of the corresponding γ_c 's. The result is shown in the diagram in Figure 2. Every point (γ_f, γ_c) on the thick line in this diagram represents a value of γ_f , and the corresponding estimate for the value of γ_c . The dashed line shows the following lower bound, which can be proved easily by adapting the proof of Guha and Khuller [15] for hardness of the facility location problem.

Theorem 13 Let γ_f and γ_c be constants with $\gamma_c < 1 + 2e^{-\gamma_f}$. Assume there is an algorithm \mathcal{A} such that for every instance \mathcal{I} of the metric facility location problem, \mathcal{A} finds a solution whose cost is not more than $\gamma_f F_{SOL} + \gamma_c C_{SOL}$ for every solution SOL for \mathcal{I} with facility and connection costs F_{SOL} and C_{SOL} . Then $\mathbf{NP} \subseteq \text{DTIME}[n^{O(\log \log n)}]$.

Similar tradeoff problems are considered by Charikar and Guha [6]. However, an important advantage that we get here is that all the inequalities $ALG \leq \gamma_f F_{SOL} + \gamma_c C_{SOL}$ are satisfied by a *single* algorithm. In Section 8, we will use the point $\gamma_f = 1$ of this tradeoff to design algorithms for other

n_c	n_f	Random Points on a Grid			Random Graphs		
		JV	ALG 1	ALG 2	JV	ALG 1	ALG 2
50	20	1.0927	1.0083	1.0004	1.0021	1.0007	1.0001
100	20	1.0769	1.0082	1.0004	1.0014	1.0022	1.0
100	50	1.2112	1.0105	1.0013	1.0225	1.0056	1.0005
200	50	1.159	1.0095	1.001	1.0106	1.0094	1.0002
200	100	1.301	1.0105	1.0016	1.0753	1.0178	1.0018
300	50	1.1151	1.0091	1.0011	1.0068	1.0102	1.0002
300	80	1.1787	1.0116	1.001	1.0259	1.0171	1.0004
300	100	1.2387	1.0118	1.0014	1.0455	1.0185	1.0009
300	150	1.327	1.0143	1.0015	1.1365	1.0249	1.0018
400	50	1.0905	1.0092	1.0005	1.0044	1.012	1.0
400	100	1.8513	1.0301	1.0026	1.0313	1.0203	1.0003
400	150	1.8112	1.0299	1.0023	1.1008	1.0234	1.0009

Table 2: Random Graphs and Random Points on a Grid

variants of the facility location problem. Other points of this tradeoff can also be useful in designing other algorithms based on our algorithm. For example, Mahdian, Ye, and Zhang [32] use the point $\gamma_f = 1.1$ of this tradeoff to obtain a 1.52-approximation algorithm for the metric facility location problem.

7 Experimental Results

We have implemented our algorithms, as well as the JV algorithm, using the programming language C. We have made four kinds of experiments. In all cases the solution of the algorithms is compared to the optimal solution of the LP-relaxation, computed using the package CPLEX to obtain an upper bound on the approximation factor of the algorithms.

The test bed of our first set of experiments consists of randomly generated instances on a $10,000 \times 10,000$ grid: In each instance, cities and facilities are points, drawn randomly from the grid. The connection cost between a city and a facility is set to be equal to the euclidean distance of the corresponding points. Furthermore, the opening cost of each facility is drawn uniformly at random from the integers between 0 and 9999.

For the second set of experiments, we have generated random graphs (according to the distribution $G(n, p)$) and assigned uniform random weights on the edges. Cities and facilities correspond to the nodes of this graph, and the connection cost between a city and a facility is defined to be the shortest path between the corresponding nodes. The opening costs of facilities are generated at random.

The instance sizes in both of the above types vary from 50 cities and 20 facilities to 400 cities and 150 facilities. For each size, 15 instances are generated and the average error of the algorithm (compared to the LP lower bound) is computed. The results of these experiments are shown in Table 2.

An Internet topology generator software, namely GT-ITM, is used to generate the third set of instances. GT-ITM is a software package for generating graphs that have a structure modeling the

n_c	n_f	JV	ALG 1	ALG 2
100	20	1.004	1.0047	1.0001
160	20	1.5116	1.0612	1.0009
160	40	1.065	1.0063	1.0
208	52	2.2537	1.074	1.019
240	60	1.0083	1.0045	1.0001
300	75	1.8088	1.0478	1.0006
312	52	1.7593	1.0475	1.0008
320	32	1.0972	1.0015	1.0
400	100	1.0058	1.0048	1.0
416	52	1.0031	1.0048	1.0

Table 3: GT-ITM Model

topology of the Internet [45]. This model is used because of the applications of facility location problems in network applications such as placing web server replicas [37]. In this model we consider transit nodes as potential facilities and stub nodes as cities. The connection cost is the distance produced by the generator. The opening costs are again random numbers. We have generated 10 instances for each of the 10 different instance sizes. The results are shown in Table 3.

We also tested all algorithms on 15 instances from [4], which is a library of test data sets for several operations research problems. Our results are shown in Table 4.

As we can see from the tables, Algorithm 2 behaves extremely well, giving almost no error in many cases. Algorithm 1 has an error of 7% on the worst instance and an average error of 2-3%. On the other hand, the JV algorithm has much larger error, sometimes as high as 100 %. We should also note that the running times of the three algorithms did not vary significantly. In the biggest instances of 1000 cities and 100 facilities all the algorithms ran in approximately 1-2 seconds. The implementation of the algorithms as well as all the data sets are available upon request. For other experimental results see [5].

8 Variants of the problem

In this section, we show that our algorithms can also be applied to several variants of the metric facility location problem.

8.1 The k -median problem

The k -median problem differs from the facility location problem in two respects: there is no cost for opening facilities, and there is an upper bound k , that is supplied as part of the input, on the number of facilities that can be opened. The k -facility location problem is a common generalization of k -median and the facility location problem. In this problem, we have an upper bound k on the number of facilities that can be opened, as well as costs for opening facilities. The k -median problem is studied extensively [2, 6, 7, 22] and the best known approximation algorithm for this problem, due to Arya et al. [2], achieves a factor of $3 + \epsilon$. It is also straightforward to adapt the proof of

n_c	n_f	JV	ALG 1	ALG 2
50	16	1.0642	1.0156	1.0
50	16	1.127	1.0363	1.0
50	16	1.1968	1.0258	1.0
50	16	1.2649	1.0258	1.0022
50	25	1.1167	1.006	1.0028
50	25	1.2206	1.0393	1.0
50	25	1.3246	1.0277	1.0
50	25	1.4535	1.0318	1.0049
50	50	1.3566	1.0101	1.0017
50	50	1.5762	1.0348	1.0061
50	50	1.7648	1.0378	1.0022
50	50	2.0543	1.0494	1.0075
1000	100	1.0453	1.0542	1.0023
1000	100	1.0155	1.0226	1.0
1000	100	1.0055	1.0101	1.0

Table 4: Instances from Operations Research library

hardness of the facility location problem [15] to show that there is no $(1 + \frac{2}{e} - \epsilon)$ -approximation algorithm for k -median, unless $\mathbf{NP} \subseteq \text{DTIME}[n^{O(\log \log n)}]$. Notice that this proves that k -median is a strictly harder problem to approximate than the facility location problem because the latter can be approximated within a factor of 1.61.

Jain and Vazirani [22] reduced the k -median problem to the facility location problem in the following sense: Suppose \mathcal{A} is an approximation algorithm for the facility location problem. Consider an instance \mathcal{I} of the problem with optimum cost OPT , and let F and C be the facility and connection costs of the solution found by \mathcal{A} . We call algorithm \mathcal{A} a Lagrangian Multiplier Preserving α -approximation (or LMP α -approximation for short) if for every instance \mathcal{I} , $C \leq \alpha(OPT - F)$. Jain and Vazirani [22] show that an LMP α -approximation algorithm for the metric facility location problem gives rise to a 2α -approximation algorithm for the metric k -median problem. They have noted that this result also holds for the k -facility location problem.

Lemma 14 [22] *An LMP α -approximation algorithm for the facility location problem gives a 2α -approximation algorithm for the k -facility location problem.*

Here we use Theorem 12 together with the scaling technique of Charikar and Guha [6] to give an LMP 2-approximation algorithm for the metric facility location problem based on Algorithm 2. This will result in a 4-approximation algorithm for the metric k -facility location problem, whereas the best previously known was 6 [22].

Lemma 15 *Assume there is an algorithm \mathcal{A} for the metric facility location problem such that for every instance \mathcal{I} and every solution SOL for \mathcal{I} , \mathcal{A} finds a solution of cost at most $F_{SOL} + \alpha C_{SOL}$, where F_{SOL} and C_{SOL} are facility and connection costs of SOL , and α is a fixed number. Then there is an LMP α -approximation algorithm for the metric facility location problem.*

Proof: Consider the following algorithm: The algorithm constructs another instance \mathcal{I}' of the problem by multiplying the facility opening costs by α , runs \mathcal{A} on this modified instance \mathcal{I}' , and outputs its answer. It is easy to see that this algorithm is an LMP α -approximation. \square

Now we only need to prove the following. The proof of this theorem follows the general scheme that is explained in Section 9.

Theorem 16 *For every instance \mathcal{I} and every solution SOL for \mathcal{I} , Algorithm 2 finds a solution of cost at most $F_{SOL} + 2C_{SOL}$, where F_{SOL} and C_{SOL} are facility and connection costs of SOL .*

Proof: By Theorem 12 we only need to prove that the solution of the factor-revealing LP 23 with $\gamma_f = 1$ is at most 2. We first write the maximization program 23 as the following equivalent linear program.

$$\begin{aligned}
& \text{maximize} && \sum_{i=1}^k \alpha_i - f && (24) \\
& \text{subject to} && \sum_{i=1}^k d_i = 1 \\
& && \forall 1 \leq i < k : \alpha_i - \alpha_{i+1} \leq 0 \\
& && \forall 1 \leq j < i < k : r_{j,i+1} - r_{j,i} \leq 0 \\
& && \forall 1 \leq j < i \leq k : \alpha_i - r_{j,i} - d_i - d_j \leq 0 \\
& && \forall 1 \leq j < i \leq k : r_{j,i} - d_i - g_{i,j} \leq 0 \\
& && \forall 1 \leq i \leq j \leq k : \alpha_i - d_j - h_{i,j} \leq 0 \\
& && \forall 1 \leq i \leq k : \sum_{j=1}^{i-1} g_{i,j} + \sum_{j=i}^k h_{i,j} - f \leq 0 \\
& && \forall i, j : \alpha_j, d_j, f, r_{j,i}, g_{i,j}, h_{i,j} \geq 0
\end{aligned}$$

We need to prove an upper bound of 2 on the solution of the above LP. Since this program is a maximization program, it is enough to prove the upper bound for any relaxation of the above program. Numerical results (for a fixed value of k , say $k = 100$) suggest that removing the second, third, and seventh inequalities of the above program does not change its solution. Therefore, we can relax the above program by removing these inequalities. Now, it is a simple exercise to write down the dual of the relaxed linear program and compute its optimal solution. This solution corresponds to multiplying the third, fourth, fifth, and sixth inequalities of the linear program 24 by $1/k$, and the first one by $(2 - 1/k)$, and adding up these inequalities. This gives an upper bound of $2 - 1/k$ on the value of the objective function. Thus, for $\gamma_f = 1$, we have $\gamma_c \leq 2$. In fact, γ_c is precisely equal to 2, as shown by the following solution for the program 23.

$$\begin{aligned}
\alpha_i &= \begin{cases} 2 - \frac{1}{k} & i = 1 \\ 2 & 2 \leq i \leq k \end{cases} \\
d_i &= \begin{cases} 1 & i = 1 \\ 0 & 2 \leq i \leq k \end{cases} \\
r_{j,i} &= \begin{cases} 1 & j = 1 \\ 2 & 2 \leq j \leq k \end{cases}
\end{aligned}$$

$$f = 2(k - 1)$$

This example shows that the above analysis of the factor-revealing LP is tight. \square

Lemma 15 and Theorem 16 provide an LMP 2-approximation algorithm for the metric facility location problem. This result improves all the results in Jain and Vazirani [22], and gives straightforward algorithms for some other problems considered by Charikar et al [8].

Notice that Theorem 13 shows that finding an LMP $(1 + \frac{2}{\epsilon} - \epsilon)$ -approximation for the metric facility location problem is hard. Also, the integrality gap examples found by Guha [14] show that Lemma 14 is tight. This shows that one cannot use Lemma 14 as a black box to obtain a smaller factor than $2 + \frac{4}{\epsilon}$ for k -median problem. Note that $3 + \epsilon$ approximation is already known [2] for the problem. Hence if one wants to beat this factor using the Lagrangian relaxation technique then it will be necessary to look into the underlying LMP algorithm as already been done by Charikar and Guha [6].

8.2 Facility location game

An important consideration, in cooperative game theory, while distributing the cost of a shared utility, is that the cost shares should satisfy the *coalition participation constraint*, i.e., the total cost share of any subset of the users shall not be larger than their stand-alone cost of receiving the service, so as to prevent this subset from seceding. In general, this turns out to be a stringent condition to satisfy. For the facility location problem, Goemans and Skutella [13] showed that such a cost allocation is only possible for a very special case. Furthermore, intractability sets in as well, for instance, in the case of the facility location problem, computing the optimal cost of serving a set of users is **NP-hard**.

In [24] Jain and Vazirani relax this notion: for a constant k , ensure that the cost share of any subset is no more than k times its stand-alone cost. They also observe that LP-based approximation algorithms directly yield a cost sharing method compatible with this relaxed notion. However, this involves solving an LP, as in the case of LP-rounding. We observe that our facility location algorithms automatically yield such a cost sharing method, with $k = 1.861$ and $k = 1.61$ respectively, by defining the cost share of city j to be α_j .

8.3 Arbitrary demands

In this version, for each city j , a non-negative integer demand d_j , is specified. An open facility i can serve this demand at the cost of $c_{ij}d_j$. The best way to look at this modification is to reduce it to unit demand case by making d_j copies of city j . This reduction suggests that we need to change our algorithms, so that each city j raises its contribution α_j at rate d_j . Note that the modified algorithms still have the same running time in more general cases, where d_j is fractional or exponentially large, and achieve the same approximation ratio.

8.4 Fault tolerant facility location with uniform connectivity requirements

We are given a connectivity requirement r_j for each city j , which specifies the number of open facilities that city j should be connected to. We can see that this problem is closely related to the

set multi-cover problem, in the case that every set can be picked at most once [38]. The greedy algorithm for set-cover can be adapted for this variant of the multi-cover problem achieving the same approximation factor. We can use the same approach to deal with the fault tolerant facility location: The mechanism of raising dual variables and opening facilities is the same as in our initial algorithms. The only difference is that city j stops raising its dual variable and withdraws its contribution from other facilities, when it is connected to r_j open facilities. We can show that when all r_j 's are equal, our algorithms can still achieve the approximation factor of 1.861 and 1.61.

8.5 Facility location with penalties

In this version we are not required to connect every city to an open facility; however, for each city j , there is a specified penalty, p_j , which we have to pay, if it is not connected to any open facility. We can modify our algorithms for this problem as follows: If α_j reaches p_j before j is connected to any open facility, the city j stops raising its dual variable and keeps its contribution equal to its penalty until it is either connected to an open facility or all remaining cities stop raising their dual variables. At this point, the algorithm terminates and unconnected cities remain unconnected. Using the linear programming formulation introduced in Charikar et al. ([8] inequalities (4.6)-(4.10)), we can show that the approximation ratio and running time of our modified algorithms have not changed.

8.6 Robust facility location

In this variant, we are given a number l and we are only required to connect $n_c - l$ cities to open facilities. This problem can be reduced to the previous one via Lagrangian relaxation. Very recently, Charikar et al. [8] proposed a primal-dual algorithm, based on JV algorithm, which achieves an approximation ratio of 3. As they showed, the linear programming formulation of this variant has an unbounded integrality gap. In order to fix this problem, they use the technique of parametric pruning, in which they guess the most expensive facility in the optimal solution. After that, they run JV algorithm on the pruned instance, where the only allowable facilities are those that are not more expensive than the guessed facility. Here we can use the same idea, using Algorithm 1 rather than the JV algorithm. Using a proof similar to the proof of the Theorem 3.2 in [8], we can prove that this algorithm solves the robust facility location problem with an approximation factor of 2.

8.7 Dealing with capacities

In real applications, it is not usually the case that the cost of opening a facility is independent of the number of cities it will serve. But we can assume that we have *economy of scales*, i.e., the cost of serving each city decreases when the number of cities increases (since publication of the first draft of this paper, this problem has also been studied in [19]). In order to capture this property, we define the following variant of the capacitated metric facility location problem. For each facility i , there is an initial opening cost f_i . After facility i is opened, it will cost s_i to serve each city. This variant can be solved using metric uncapacitated facility location problem: We just have to change the metric such that for each city j and facility i , $c'_{ij} = c_{ij} + s_i$. Clearly, c' is also a metric and the solution of the metric uncapacitated version to this problem can be interpreted as a solution to the original problem with the same cost.

We can reduce the variant of the capacitated facility location problem in which each facility can be opened many times [22] to this problem by defining $s_i = f_i/u_i$. If in the solution to this problem k cities are connected to facility i , we open this facility $\lceil k/u_i \rceil$ times. The cost of the solution will be at most two times the original cost so any α -approximation for the uncapacitated facility location problem can be turned into a 2α -approximation for this variant of the capacitated version. We can also use the same technique as in [22] to give a factor 3-approximation algorithm for this problem based on the LMP 2-approximation algorithm for uncapacitated facility location problem.

9 Discussion

The method of dual fitting can be seen as an implementation of the primal-dual schema in which, instead of relaxing complementary slackness conditions (which is the most common way of implementing the schema), we relax feasibility of the dual. However, we prefer to reserve the term primal-dual for algorithms that produce feasible primal and dual solutions.

Let us show how the combination of dual fitting with factor-revealing LP applies to the set cover problem. The duality-based restatement of the greedy algorithm (see [44]) is: All elements in the universal set U increase their dual variables uniformly. Each element contributes its dual towards paying for the cost of each of the sets it is contained in. When the total contribution offered to a set equals its cost, the set is picked. At this point, the newly covered elements freeze their dual variables and withdraw their contributions from all other sets. As stated in the introduction, the latter (important) step ensures that the primal is fully paid for by the dual. However, we might not get a feasible dual solution. To make the dual solution feasible we look for the smallest positive number Z , so that when the dual solution is shrunk by a factor of Z , it becomes feasible. An upper bound on the approximation factor of the algorithm is obtained by maximizing Z over all possible instances.

Clearly Z is also the maximum factor by which any set is over-tight. Consider any set S . We want to see what is the worst factor, over all sets and over all possible instances of the problem, by which a set S is over-tight. Let the elements in S be $1, 2, \dots, k$. Let x_i be the dual variable corresponding to the element i at the end of the algorithm. Without loss of generality we may assume that $x_1 \leq x_2 \leq \dots \leq x_k$. It is easy to see that at time $t = x_i^-$, total duals offered to S is at least $(k - i + 1)x_i$. Therefore, this value cannot be greater than the cost of the set S (denoted by c_S). So, the optimum solution of the following mathematical program gives an upper bound on the value of Z . (Note that c_S is a variable not a constant).

$$\begin{aligned}
& \text{maximize} && \frac{\sum_{i=1}^k x_i}{c_S} && (25) \\
& \text{subject to} && \forall 1 \leq i < k : x_i \leq x_{i+1} \\
& && \forall 1 \leq i \leq k : (k - i + 1)x_i \leq c_S \\
& && \forall 1 \leq i \leq k : x_i \geq 0 \\
& && c_S \geq 1
\end{aligned}$$

The above optimization program can be turned into a linear program by adding the constraint $c_S = 1$ and changing the objective function to $\sum_{i=1}^k x_i$. We call this linear program the *factor-revealing LP*. Notice that the factor-revealing LP has nothing to do with the LP formulation of

the set cover problem; it is only used in order to analyze this particular algorithm. This is the important distinction between the factor-revealing LP technique, and other LP-based techniques in approximation algorithms.

One advantage of reducing the analysis of the approximation guarantee of an algorithm to obtaining an upper bound on the optimal solution to a factor-revealing LP is that one can introduce empirical experimentation into the latter task. This can also help decide which aspects of the execution of the algorithm to introduce into the factor-revealing LP to obtain the best possible bound on the performance of the algorithm, e.g., we needed to introduce the variables $r_{j,i}$ in Section 5.1 in order to get a good bound on the approximation ratio of Algorithm 2.

In general, this technique is not guaranteed to yield a tight analysis of the algorithm, since the algorithm may be performing well not because of local reasons but for some global reasons that are difficult to capture in a factor-revealing LP. In the case of set cover, this method not only produces a tight analysis, but the factor-revealing LP also helps produce a tight example for the algorithm. From any feasible solution x of factor-revealing LP 25, one can construct the following instance: There are k elements $1, \dots, k$, a set $S = \{1, \dots, k\}$ of cost $1 + \epsilon$ which is the optimal solution, and sets $S_i = \{i\}$ of cost x_i for $i = 1, \dots, k$. It is easy to verify that the greedy algorithm gives a solution that is $\sum x_i$ times worse than the optimal on this instance. Picking x to be the optimal solution, we get a tight example, and also show that the approximation ratio of the greedy algorithm is precisely equal H_n , the optimal solution of the factor-revealing LP.

Finally, in terms of practical impact, what is the significance of improving the approximation guarantee for facility location from 3 to 1.81 or 1.61 when practitioners are seeking algorithms that come within 2% to 5% of the optimal? The superior experimental results of our algorithms, as compared with the JV algorithm, seem to provide the answer and to support the argument made in [44] (Preface, page IX) that the approximation factor should be viewed as a “measure that forces us to explore deeper into the combinatorial structure of the problem and discover more powerful tools for exploiting this structure” and the observation that “sophisticated algorithms do have the error bounds of the desired magnitude, 2% to 5%, on typical instances, even though their worst case error bounds are much higher”.

Acknowledgments. We would like to thank Michel Goemans, Mohammad Ghodsi, Nicole Immorlica, Nisheet K. Vishnoi, Milena Mihail, and Christos Gkantsidis for their helpful comments and discussions.

References

- [1] M. Andrews and L. Zhang. The access network design problem. In *Proceedings of the 39th Annual IEEE Symposium on Foundations of Computer Science*, 1998.
- [2] V. Arya, N. Garg, R. Khandekar, A. Meyerson, K. Munagala, and V. Pandit. Local search heuristics for k-median and facility location problems. In *Proceedings of 33rd ACM Symposium on Theory of Computing*, 2001.
- [3] M.L. Balinski. On finding integer solutions to linear programs. In *Proc. IBM Scientific Computing Symposium on Combinatorial Problems*, pages 225–248, 1966.
- [4] J. E. Beasley. Operations research library, 2002.

- [5] C. W. Cameron, S. H. Low, and D. X. Wei. High-density model for server allocation and placement. unpublished manuscript, 2002.
- [6] M. Charikar and S. Guha. Improved combinatorial algorithms for facility location and k-median problems. In *Proceedings of the 40th Annual IEEE Symposium on Foundations of Computer Science*, pages 378–388, October 1999.
- [7] M. Charikar, S. Guha, E. Tardos, and D.B. Shmoys. A constant-factor approximation algorithm for the k-median problem. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing*, pages 1–10, May 1999.
- [8] M. Charikar, S. Khuller, D. Mount, and G. Narasimhan. Facility location with outliers. In *12th Annual ACM-SIAM Symposium on Discrete Algorithms*, Washington DC, January 2001.
- [9] F.A. Chudak and D. Shmoys. Improved approximation algorithms for the uncapacitated facility location problem. unpublished manuscript, 1998.
- [10] V. Chvatal. A greedy heuristic for the set covering problem. *Math. Oper. Res.*, 4:233–235, 1979.
- [11] G. Cornuejols, G.L. Nemhauser, and L.A. Wolsey. The uncapacitated facility location problem. In P. Mirchandani and R. Francis, editors, *Discrete Location Theory*, pages 119–171. John Wiley and Sons Inc., 1990.
- [12] M. Goemans and J. Kleinberg. An improved approximation ratio for the minimum latency problem. *Mathematical Programming*, 82:111–124, 1998.
- [13] M.X. Goemans and M. Skutella. Cooperative facility location games. In *Symposium on Discrete Algorithms*, pages 76–85, 2000.
- [14] S. Guha. *Approximation algorithms for facility location problems*. PhD thesis, Stanford University, 2000.
- [15] S. Guha and S. Khuller. Greedy strikes back: Improved facility location algorithms. *Journal of Algorithms*, 31:228–248, 1999.
- [16] S. Guha, A. Meyerson, and K. Munagala. Hierarchical placement and network design problems. In *Proceedings of the 41th Annual IEEE Symposium on Foundations of Computer Science*, 2000.
- [17] S. Guha, A. Meyerson, and K. Munagala. Improved combinatorial algorithms for single sink edge installation problems. Technical Report STAN-CS-TN00 -96, Stanford University, 2000.
- [18] S. Guha, A. Meyerson, and K. Munagala. Improved algorithms for fault tolerant facility location. In *Symposium on Discrete Algorithms*, pages 636–641, 2001.
- [19] M. Hajiaghayi, M. Mahdian, and M. Mirrokni. Facility location problem with concave cost functions. unpublished manuscript, 2002.
- [20] D. S. Hochbaum. Heuristics for the fixed cost median problem. *Mathematical Programming*, 22(2):148–162, 1982.

- [21] K. Jain, M. Mahdian, and A. Saberi. A new greedy approach for facility location problems. In *ACM Symposium on Theory of Computing*, 2002.
- [22] K. Jain and V.V. Vazirani. Primal-dual approximation algorithms for metric facility location and k-median problems. In *Proceedings of the 40th Annual IEEE Symposium on Foundations of Computer Science*, pages 2–13, October 1999.
- [23] K. Jain and V.V. Vazirani. An approximation algorithm for the fault tolerant metric facility location problem. In K. Jansen and S. Khuller, editors, *Approximation Algorithms for Combinatorial Optimization, Proceedings of APPROX 2000*, volume 1913 of *Lecture Notes in Computer Science*, pages 177–183, September 2000.
- [24] K. Jain and V.V. Vazirani. Applications of approximation algorithms to cooperative games. In *ACM Symposium on Theory of Computing*, pages 364–372, 2001.
- [25] S. Jamin, C. Jin, Y. Jin, D. Raz, Y. Shavitt, and L. Zhang. On the placement of internet instrumentations. In *Proceedings of IEEE INFOCOM'00*, pages 26–30, March 2000.
- [26] L. Kaufman, M. vanden Eede, and P. Hansen. A plant and warehouse location problem. *Operational Research Quarterly*, 28:547–557, 1977.
- [27] M.R. Korupolu, C.G. Plaxton, and R. Rajaraman. Analysis of a local search heuristic for facility location problems. In *Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1–10, January 1998.
- [28] A.A. Kuehn and M.J. Hamburger. A heuristic program for locating warehouses. *Management Science*, 9:643–666, 1963.
- [29] B. Li, M. Golin, G. Italiano, Deng X., and K. Sohraby. On the optimal placement of web proxies in the internet. In *Proceedings of IEEE INFOCOM'99*, pages 1282–1290, 1999.
- [30] L. Lovasz. On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13:383–390, 1975.
- [31] 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*, volume 2129 of *Lecture Notes in Computer Science*, pages 127–137. Springer-Verlag, 2001.
- [32] M. Mahdian, Y. Ye, and J. Zhang. Improved approximation algorithms for metric facility location problems. Unpublished manuscript, 2002.
- [33] R.J. McEliece, E.R. Rodemich, H. Rumsey Jr., and L.R. Welch. New upper bounds on the rate of a code via the delarte-macwilliams inequalities. *IEEE Transactions on Information Theory*, 23:157–166, 1977.
- [34] R.R. Mettu and G. Plaxton. The online median problem. In *IEEE Symposium on Foundations of Computer Science*, pages 339–348, 2000.
- [35] G.L. Nemhauser and L.A. Wolsey. *Integer and Combinatorial Optimization*. John Wiley and Sons, 1990.

- [36] L. Qiu, V.N. Padmanabhan, and G. Voelker. On the placement of web server replicas. In *Proceedings of IEEE INFOCOM'01*, Anchorage, AK, USA, April 2001.
- [37] Lili Qiu, Venkata N. Padmanabhan, and Geoffrey M. Voelker. On the placement of web server replicas. In *INFOCOM*, pages 1587–1596, 2001.
- [38] S. Rajagopalan and Vazirani V.V. Primal-dual RNC approximation algorithms for set cover and covering integer programs. *SIAM Journal on Computing*, 28(2):525–540, 1999.
- [39] D.B. Shmoys, E. Tardos, and K.I. Aardal. Approximation algorithms for facility location problems. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, pages 265–274, 1997.
- [40] J.F. Stollsteimer. *The effect of technical change and output expansion on the optimum number, size and location of pear marketing facilities in a California pear producing region*. PhD thesis, University of California at Berkeley, 1961.
- [41] J.F. Stollsteimer. A working model for plant numbers and locations. *J. Farm Econom.*, 45:631–645, 1963.
- [42] M. Sviridenko. An 1.582-approximation algorithm for the metric uncapacitated facility location problem. to appear in the Ninth Conference on Integer Programming and Combinatorial Optimization, 2002.
- [43] M. Thorup. Quick k -median, k -center, and facility location for sparse graphs. In *Automata, Languages and Programming, 28th International Colloquium, Crete, Greece*, volume 2076 of *Lecture Notes in Computer Science*, pages 249–260, 2001.
- [44] V. V. Vazirani. *Approximation Algorithms*. Springer-Verlag, Berlin, 2001.
- [45] Ellen W. Zegura, Kenneth L. Calvert, and Samrat Bhattacharjee. How to model an internet-work. In *IEEE Infocom*, volume 2, pages 594–602, San Francisco, CA, March 1996. IEEE.