Better Bounds for Incremental Medians

Marek Chrobak* Mathilde Hurand†

Abstract

In the incremental version of the well-known *k-median problem* the objective is to compute an incremental sequence of facility sets $F_1 \subseteq F_2 \subseteq .... \subseteq F_n$, where each F_k contains at most k facilities. We say that this incremental medians sequence is *R-competitive* if the cost of each F_k is at most R times the optimum cost of k facilities. The smallest such R is called the *competitive ratio* of the sequence $\{F_k\}$. Mettu and Plaxton [6, 7] presented a polynomial-time algorithm that computes an incremental sequence with competitive ratio ≈ 30. They also showed a lower bound of 2. The upper bound on the ratio was improved to 8 in [5] and [4]. We improve both bounds in this paper. We first show that no incremental sequence can have competitive ratio better than 2.01 and we give a probabilistic construction of a sequence whose competitive ratio is at most $2 + 4\sqrt{2} \approx 7.656$. We also propose a new approach to the problem that for instances that we refer to as *equable* achieves an optimal ratio of 2.

keywords Incremental medians, approximation algorithm, online algorithm, analysis of algorithme

1 Introduction

The *k-median* problem is one of the most studied facility location problems. We are given two sets: a set $\mathcal{C}$ of *customers* and a set $\mathcal{F}$ of n *facilities*, with a metric function d that specifies the distance d_{xy} between any two points $x, y \in \mathcal{C} \cup \mathcal{F}$. The cost of a facility set $F \subseteq \mathcal{F}$, denoted by *cost*(F), is defined as the minimum sum, over all customers $c \in \mathcal{C}$, of d_{cF}, where $d_{cF} = \min_{f \in F} d_{cf}$ is the minimum distance from c to F. Given k, the objective is to compute a set of k facilities with minimum cost.

Not surprisingly, the k-median problem is NP-hard. A number of polynomial-time approximation algorithms have been proposed, with the latest one, by Arya *et al.* [1, 2] achieving the ratio of $3 + \epsilon$, for any $\epsilon > 0$.

Mettu and Plaxton [6, 7] introduced the *incremental medians problem*, where the permitted number k of facilities is not specified in advance. Starting with the empty set, an algorithm receives authorizations for new facilities over time, and after each authorization it is allowed to add another facility to the existing ones. As a result, such an algorithm produces an incremental sequence of facility sets $F_1 \subseteq F_2 \subseteq ... \subseteq F_n$, where $|F_k| \leq k$ for all k. This sequence $\{F_k\}$ is said to be *R-competitive* if *cost*(F_k) is at most R times the optimum cost of k facilities, for each k. The smallest such R is called the *competitive ratio* of $\{F_k\}$.

*Department of Computer Science, University of California, Riverside, CA 92521. Research supported by NSF Grant CCR-0208856.

†Department d'Informatique (LIX), Ecole Polytechnique, Palaiseau, France. hurand@lix.polytechnique.fr

Mettu and Plaxton [6, 7] gave a polynomial-time algorithm that computes such an incremental sequence with competitive ratio ≈ 30. This result is quite remarkable, for there is no apparent reason why an incremental sequence $\{F_k\}$ of facility sets, with each $cost(F_k)$ within a constant factor of the the optimum, would even exist – let alone be computed efficiently.

It is thus natural to address the issue of *existence* separately from *computational complexity*, and this is what we focus on in this paper. As shown by Mettu and Plaxton [6, 7], no ratio better than 2 is possible, that is, for each $\epsilon > 0$ there is a metric space where each incremental facility sequence has competitive ratio at least $2 - \epsilon$. The upper bound on the ratio was improved to 8 by Lin *et al.* [5] and, independently, by Chrobak *et al.* [4]. In [5], the authors also show that a 16-competitive incremental median sequence can be computed in polynomial time.

Our results. We improve both the lower and upper bounds for incremental medians. For the lower bound, we show that, in general, no competitive ratio better than 2.01 is possible. We also prove, via a probabilistic argument, that each instance has an incremental medians sequence with competitive ratio at most $2 + 4\sqrt{2} \approx 7.656$.

In numerical terms, the improvement of the lower bound is mostly symbolic, as it implies that 2 is not the "right" ratio. For the upper bound, our result shows that the doubling method from [5, 4] (see also [3]) is not optimal – even though it gives the optimal ratio of 4 for the closely related "resource augmentation" version of incremental medians [4]. As discussed in Section 6, we believe that our methods can be refined to further improve both the lower and upper bounds.

In addition, we consider a special case of the incremental medians problem where for any fixed value of k, each customer has the same distance to the optimal k-median. We refer to such instances as *equable*. (See Section 5 for a formal definition.) For this case, we show a construction of a 2-competitive incremental medians sequence, matching the lower bound from [6, 7]. Our method for this case is very different from previous constructions and we believe that it will be useful in improving the upper bound for general spaces. In fact, this result implies that if there is a constant $\gamma \geq 1$ such that for each fixed k all customers' optimal costs are within factor γ of each other, then our construction achieves ratio at most 2γ – improving our own bound above if $\gamma < 1 + 2\sqrt{2}$.

2 Preliminaries

Let $(\mathcal{F}, \mathcal{C})$ be an instance of the medians problem, where $\mathcal{F}$ is a set of n facilities, $\mathcal{C}$ is the set of customers, and $\mathcal{F} \cup \mathcal{C}$ forms a metric space. By d_{xy} or $d(x, y)$ we denote the distance between points x, y. If Y is a set, we also write $d_{xY} = \min_{y \in Y} d_{xy}$ for the minimum distance from x to Y. For a facility set $F \subseteq \mathcal{F}$, denote by $cost(F)$ the cost of F, that is $\sum_{x \in \mathcal{C}} d_{xF}$.

For a point x and a set Y, denote by $\Gamma_Y(x)$ the point $y \in Y$ that is closest to x, that is $d_{xy} = d_{xY}$ (if this point is not unique, then break the tie arbitrarily.) If X is a set, we also define $\Gamma_Y(X) = \{\Gamma_Y(x) \mid x \in X\}$. Clearly, $|\Gamma_Y(X)| \leq |X|$. Note that if F is a facility set and X is a set of customers, then $\Gamma_F(X)$ is exactly the set of facilities in F that serve customers in X if F is the facility set under consideration.

By opt_k we denote the optimum cost of k facilities, that is

$$opt_k \;=\; \min\{cost(F) \mid F \subseteq \mathcal{F} \,\&\, |F| = k\}. \tag{1}$$

By $F^*_k \subseteq \mathcal{F}$ we will denote the optimal set of k facilities, that is, the k-*median*. (As before, ties are broken arbitrarily.) Thus $cost(F^*_k) = opt_k$.

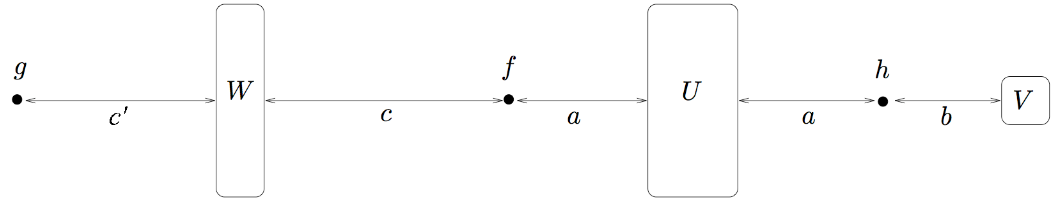

Figure 1: Metric space used in the lower bound.

3 A New Lower Bound

In this section we prove our lower bound of 2.01 on the competitive ratio for incremental medians, improving slightly the previous bound of 2 from [6, 7].

Theorem 1. *There is an instance* $(\mathcal{C}, \mathcal{F})$ *for which no incremental median sequence has competitive ratio smaller than* 2.01.

Proof. The set of customers is $\mathcal{C} = U \cup V \cup W$, where U, V, W are disjoint sets with $|U|+|V|+|W| = n-3$, where n is a large integer. The set of facilities is $\mathcal{F} = \{f, g, h\} \cup \mathcal{C}$. The distances between customers and facilities are shown in Figure 1. For each set U, V, W, all customers in a set have the same distance to each facility. For example, the distance from f to all $u \in U$ is a, the distance from h to all $v \in V$ is b, etc. Other distances are measured along the shortest paths in the graph from Figure 1. This is also true for two customers from a same set (they are *not* at distance 0 from one-another). For example, if $v, v' \in V$ and $v' \neq v$ then the distance from v to v' is $2b$.

Since for $k = n-3$ the optimal cost is 0, the first $n-3$ facilities in any competitive incremental sequence must be chosen from $\mathcal{C}$. In fact, we will only use only three values of k: $k = 1, 2$ and $n-3$.

To prove that there is no incremental median with ratio better than R, we only need to give some values a, b, c, c', $|U|$, $|V|$ and $|W|$ such that:

$$\min\{\mathit{cost}(v), \mathit{cost}(w)\} \geq R \cdot \mathit{cost}(f), \quad \text{and} \tag{2}$$

$$\min\{\mathit{cost}(u,v), \mathit{cost}(u,w)\} \geq R \cdot \mathit{cost}(g,h). \tag{3}$$

These inequalities imply the lower bound of R, for (2) implies that, for $k = 1$, to beat ratio R we must pick some $u \in U$ as the first facility, and (3) implies that, for $k = 2$, it is not possible to add to u another facility and preserve ratio R.

In order to simplify calculations, we slightly modify the way we compute the costs. If a facility at some point $x \in U \cup V \cup W$ serves a customer $z \neq x$ then the cost of z is the length of the shortest path from z to x via one facility f, g, or h, while the cost of $z = x$ is 0. Our first modification is that we will charge this $z = x$ the cost of such a shortest path as well, that is, c cannot serve itself directly at cost 0. For example, if there is a facility at $x \in U$, then we will charge x the cost of $2a$ to get to this facility. Since this increases the cost by a factor of at most $1 + \Theta(1/n)$, by taking n large enough in the proof below, the argument remains valid for the true cost values.

With this convention in mind, we set $a = 5/4$, $b = 1$, $c = 211/100$, $c' = 141/100$, $|U| = 295\lambda$, $|V| = 25\lambda$, and $|W| = 149\lambda$, for some large integer λ. (Thus $n = 469\lambda+3$.) Note that $b \leq a \leq c \leq c'$.

Then, for $k = 1$ we have

$$\begin{aligned} \mathit{cost}(f) &= |U|a + |V|(b + 2a) + |W|c \\ \mathit{cost}(v) &= |U|(a + b) + |V|(2b) + |W|(b + 2a + c) \\ \mathit{cost}(w) &= |U|(a + c) + |V|(b + 2a + c) + |W|(2c') \end{aligned}$$

and for $k = 2$ we have

$$\begin{aligned} \mathit{cost}(g, h) &= |U|a + |V|b + |W|c' \\ \mathit{cost}(u, v) &= |U|(a + b) + |V|(2b) + |W|(a + c) \\ \mathit{cost}(u, w) &= |U|(2a) + |V|(a + b) + |W|(2c') \end{aligned}$$

Then

$$\begin{aligned} \frac{\min\{\mathit{cost}(v), \mathit{cost}(w)\}}{\mathit{cost}(f)} &= \frac{2039}{1014} \geq 2.01, \quad \text{and} \\ \frac{\min\{\mathit{cost}(u, v), \mathit{cost}(u, w)\}}{\mathit{cost}(g, h)} &= \frac{121393}{60384} \geq 2.01. \end{aligned}$$

This implies that inequalities (2), (3) hold with $R = 2.01$, and the lower bound follows. $\square$

4 A New Upper Bound

In this section we construct an incremental medians sequence with competitive ratio $R = 2 + 4\sqrt{2}$. First, we show that, given a facility set H we can find subsets $F \subseteq G \subseteq H$ of specified sizes and of appropriately small cost. We then use this result to construct our incremental medians sequence.

4.1 Choosing Two Nested Facility Sets

Let $1 \leq k \leq l \leq m \leq n$. (Recall that $n = |\mathcal{F}|$ is the number of facilities.) Throughout this section we consider three facility sets: H of cardinality m, U of cardinality k, and V of cardinality l. Intuitively, U and V represent optimal $k-$ and $l-$ medians. We use a probabilistic argument to show that there exist two sets F and G, with $|F| = k$, $|G| = l$ and $F \subseteq G \subseteq H$, such that $\mathit{cost}(F)$ and $\mathit{cost}(G)$ are bounded in terms of $\mathit{cost}(U)$, $\mathit{cost}(V)$ and $\mathit{cost}(H)$.

Lemma 2. *Let $1 \leq k \leq l \leq m \leq n$, and let U, V and H be facility sets with $|H| = m$, $|V| = l$ and $|U| = k$. Then there is a set $T \subseteq V$ with $|T| = k$ such that, denoting $\bar{T} = V - T$, we have*

$$\mathit{cost}(\Gamma_H(T)) + \mathit{cost}(\Gamma_H(U \cup \bar{T})) \leq 2 \cdot \mathit{cost}(H) + 4 \cdot \mathit{cost}(V) + 2 \cdot \mathit{cost}(U). \tag{4}$$

Proof. We use a probabilistic argument, by defining a probability distribution on subsets $T \subseteq V$ and proving that inequality (4) holds in expectation.

Define a random mapping $\Phi : U \to \mathcal{C}$, where $\Phi(u)$ is chosen uniformly from the set $\mathcal{C}_u = \{x \in \mathcal{C} \mid \Gamma_U(x) = u\}$. In other words, $\Phi(u)$ is a random customer of u when U is the facility set. Order arbitrarily the elements of V, and for any given Φ define T_Φ as the subset of V that consists of $\Gamma_V(\Phi(U))$ and $k - |\Gamma_V(\Phi(U))|$ smallest elements of V that are not in $\Gamma_V(\Phi(U))$. Thus $|T_\Phi| = k$.

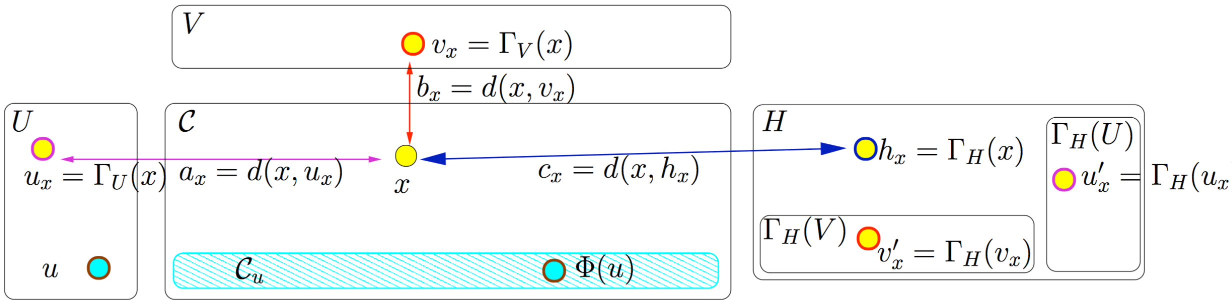

Figure 2: Notations.

For each point x in $\mathcal{C}$, let $u_x = \Gamma_U(x)$, $v_x = \Gamma_V(x)$ and $h_x = \Gamma_H(x)$ be the points serving x respectively in U, V and H. The corresponding distances from x are denoted $a_x = d(x, u_x)$, $b_x = d(x, v_x)$ and $c_x = d(x, h_x)$. Let also $u'_x = \Gamma_H(u_x)$ and $v'_x = \Gamma_H(v_x)$. (See Figure 4.1.)

We now temporarily fix the mapping Φ and a customer $x \in \mathcal{C}$. To simplify notation, we write $T_\Phi = T$ and $u = u_x$. We claim that

$$d(x, \Gamma_H(T)) + d(x, \Gamma_H(U \cup \bar{T})) \;\leq\; a_x + 2b_x + c_x + a_{\Phi(u)} + 2b_{\Phi(u)} + c_{\Phi(u)}. \tag{5}$$

To prove the claim, we consider two cases, for $v_x \in T$ and $v_x \in \bar{T}$.

Case 1: $v_x \in \bar{T}$. This case is illustrated in Figure 3.

Since $v'_{\Phi(u)} \in \Gamma_H(T)$, using the definition of $v'_{\Phi(u)}$ and several applications of the triangle inequality, we have $d(x, \Gamma_H(T)) \leq d(x, v'_{\Phi(u)}) \leq a_x + d(u, v_{\Phi(u)}) + d(v_{\Phi(u)}, v'_{\Phi(u)}) \leq a_x + [a_{\Phi(u)} + b_{\Phi(u)}] + d(v_{\Phi(u)}, h_{\Phi(u)}) \leq a_x + a_{\Phi(u)} + 2b_{\Phi(u)} + c_{\Phi(u)}$.

Since $v'_x \in \Gamma_H(U \cup \bar{T})$, using the definition of v'_x and triangle inequality, $d(x, \Gamma_H(U \cup \bar{T})) \leq d(x, v'_x) \leq b_x + d(v_x, v'_x) \leq b_x + d(v_x, h_x) \leq 2b_x + c_x$.

Combining the two bounds, we get

$$d(x, \Gamma_H(T)) + d(x, \Gamma_H(U \cup \bar{T})) \;\leq\; a_x + 2b_x + c_x + a_{\Phi(u)} + 2b_{\Phi(u)} + c_{\Phi(u)}.$$

Case 2: $v_x \in T$. This case is illustarted in Figure 4.

Since $v'_x \in \Gamma_H(T)$, using the triangle inequality and the definition of v'_x, we have $d(x, \Gamma_H(T)) \leq d(x, v'_x) \leq b_x + d(v_x, v'_x) \leq b_x + d(v_x, h_x) \leq 2b_x + c_x$.

Since $u'_x \in \Gamma_H(U \cup \bar{T})$, using the definition of $u'_x = \Gamma_H(u)$, we have $d(x, \Gamma_H(U \cup \bar{T})) \leq d(x, u'_x) \leq a_x + d(u, u'_x) \leq a_x + d(u, h_{\Phi(u)}) \leq a_x + a_{\Phi(u)} + c_{\Phi(u)}$.

Combining the two bounds we get

$$\begin{aligned} d(x, \Gamma_H(T)) + d(x, \Gamma_H(U \cup \bar{T})) \;&\leq\; a_x + 2b_x + c_x + a_{\Phi(u)} + c_{\Phi(u)} \\ &\leq\; a_x + 2b_x + c_x + a_{\Phi(u)} + 2b_{\Phi(u)} + c_{\Phi(u)}, \end{aligned}$$

completing the proof of inequality (5).

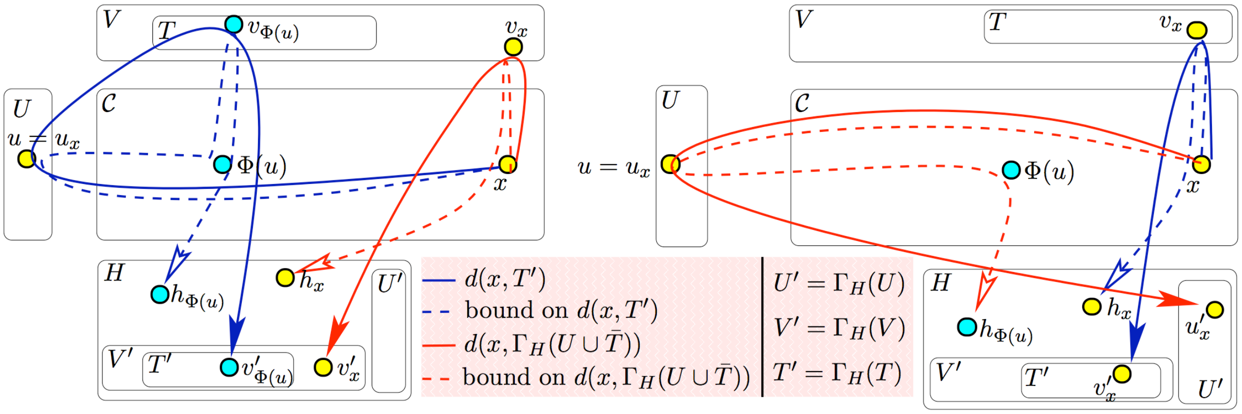

Figure 3: The proof of (5) when $v_x \in \bar{T}$. Figure 4: The proof of (5) when $v_x \in T$.

From (5), for a fixed Φ we have

$$\begin{aligned} cost(\Gamma_H(T_\Phi)) + cost(\Gamma_H(U\cup\bar{T}_\Phi)) &\le \sum_{u\in U}\sum_{x\in\mathcal{C}_u}\left[a_x + 2b_x + c_x + a_{\Phi(u)} + 2b_{\Phi(u)} + c_{\Phi(u)}\right] \\ &\le cost(H) + 2\cdot cost(V) + cost(U) \\ &\qquad + \sum_{u\in U}|\mathcal{C}_u|\cdot[a_{\Phi(u)} + 2b_{\Phi(u)} + c_{\Phi(u)}]. \qquad (6) \end{aligned}$$

For any facility set Z, we have $cost(Z) = \sum_{u\in U}|\mathcal{C}_u|\cdot \mathrm{Exp}_\Phi[d(\Phi(u),Z)]$. Applying it to $Z=U$, V and H, and using the linearity of expectation, inequality (6) yields

$$\begin{aligned} \mathrm{Exp}_\Phi\left[cost(\Gamma_H(T_\Phi)) + cost(\Gamma_H(U\cup\bar{T}_\Phi))\right] &\le cost(H) + 2\cdot cost(V) + cost(U) \\ &\qquad + \sum_{u\in U}|\mathcal{C}_u|\cdot \mathrm{Exp}_\Phi\left[a_{\Phi(u)} + 2b_{\Phi(u)} + c_{\Phi(u)}\right] \\ &= 2\cdot cost(H) + 4\cdot cost(V) + 2\cdot cost(U). \end{aligned}$$

This implies that there is a $T = T_\Phi$ that satisfies the lemma. □

Theorem 3. *Let $1\le k\le l\le m\le n$. For any facility sets H, U and V with $|U|=k$, $|V|=l$, $|H|=m$, there exist $F\subseteq G\subseteq H$ with $|F|=k$, $|G|=l$ such that*

(i) *$cost(F)\le cost(H)+2\cdot cost(U)$ and*

(ii) *$cost(G)\le cost(H)+4\cdot cost(V)$.*

Proof. Let $U'=\Gamma_H(U)$ and $V'=\Gamma_H(V)$ be the facilities in H that are closest to those in U and V, respectively. Using the triangle inequality, it is not difficult to show (see [5, 4], for example) that $cost(U')\le cost(H)+2\cdot cost(U)$ and $cost(V')\le cost(H)+2\cdot cost(V)$.

Let $T\subseteq V$ be the set from Lemma 2. Then either $cost(\Gamma_H(T))\le cost(H)+2\cdot cost(U)$ or $cost(\Gamma_H(U\cup\bar{T}))\le cost(H)+4\cdot cost(V)$. In the first case, we take $F=\Gamma_H(T)$ and $G=V'$, and in the second case we take $F=U'$ and $G=\Gamma_H(U\cup\bar{T})$. (If $|F|<k$ or $|G|<l$, we can increase their cardinalities by adding a sufficient number of elements of H while preserving the inclusion $F\subseteq G$.) The theorem then follows from Lemma 2 and the bounds on $cost(U')$ and $cost(V')$. □

4.2 Competitive Incremental Medians

Recall that n is the number of facilities, F_j^* is the optimal j-median and $opt_j = cost(F_j^*)$, for each $j = 1, 2, ..., n$. Our objective is to construct an incremental medians sequence $F_1 \subseteq F_2 \subseteq ... \subseteq F_n$.

The general approach is similar to that in [5, 4]: we construct the sequence backwards, at each step extracting a smaller set of facilities from among those selected earlier. These sets F_j will be constructed only for values of j in a predefined sequence $\{\kappa(a)\}$ of indices, for which the optimal costs increase exponentially with a. For the intermediate values of j, we simply let F_j to be $F_{\kappa(a)}$, where a is the smallest index for which $\kappa(a) \le j$.

The crucial difference between our method and the previous constructions is in how we extract facilities from $F_{\kappa(a)}$ to form $F_{\kappa(a+1)}$. The algorithms in [5] and [4] select $\kappa(a+1)$ facilities in $F_{\kappa(a)}$ that are closest to those in the optimal set $F^*_{\kappa(a+1)}$. Instead, we use our probabilistic construction from the previous section to simultaneously extract *two* facility sets next in the sequence, namely $F_{\kappa(a+1)}$ and $F_{\kappa(a+2)}$, with Theorem 3 providing an upper bound on their costs.

Construction of incremental medians. Without loss of generality we can assume that $opt_n = 1$, for otherwise we can normalize the instance by dividing all distances by opt_n. (If $opt_n = 0$, instead of n, we can start the process with the largest n' for which $opt_{n'} > 0$.)

We use two parameters $\gamma = 2 + \sqrt{2}/2 \approx 2.71$ and $\lambda = 3\sqrt{2}/2 - 1 \approx 1.16$. We now define a sequence of indices $n = \kappa(0) \ge \kappa(1) \ge ... \ge \kappa(h) = 1$. For $a = 0, 1, ...,$ let

$$\kappa(a) \quad = \quad \begin{cases} \min\{j \mid opt_j \le \gamma^{a/2}\} & \text{if } a \text{ is even} \\ \min\{j \mid opt_j \le \lambda\gamma^{(a-1)/2}\} & \text{if } a \text{ is odd} \end{cases}$$

and choose h to be the smallest a for which $\kappa(a) = 1$. For simplicity, we will assume that h is even. Note that we allow some of the elements in the sequence $\{\kappa(a)\}$ to be equal.

We first define facility sets F_j for $j = \kappa(0), \kappa(1), ..., \kappa(h)$. Initially, $F_{\kappa(0)} = \mathcal{F}$, the set of all facilities. Suppose that $F_{\kappa(a)}$ has been already defined for some even $a \ge 0$. In Theorem 3 let $m = \kappa(a)$, $H = F_{\kappa(a)}$, $l = \kappa(a+1)$, $k = \kappa(a+2)$, $V = F^*_{\kappa(a+1)}$ and $U = F^*_{\kappa(a+2)}$. We then choose $F_{\kappa(a+2)} \subseteq F_{\kappa(a+1)} \subseteq F_{\kappa(a)}$ such that

$$cost(F_{\kappa(a+1)}) \quad \le \quad cost(F_{\kappa(a)}) + 4\,opt_{\kappa(a+1)}, \quad \text{and} \tag{7}$$

$$cost(F_{\kappa(a+2)}) \quad \le \quad cost(F_{\kappa(a)}) + 2\,opt_{\kappa(a+2)}. \tag{8}$$

The existence of such sets is guaranteed by Theorem 3; namely take $F_{\kappa(a+1)} = G$ and $F_{\kappa(a+2)} = F$.

Next, we extend the sequence to other values of j. If $\kappa(a+1) < j < \kappa(a)$, we simply let $F_j = F_{\kappa(a+1)}$. This completes the construction.

Theorem 4. *The incremental sequence $\{F_j\}$ constructed above is R-competitive, where $R = 2 + 4\sqrt{2} \approx 7.656$.*

Proof. For each $j = 1, ..., n$, denote $cost_j = cost(F_j)$. Using the bounds (7), (8), and the definition of the sequence $\{\kappa(a)\}$, each value $cost_{\kappa(a)}$ can be estimated as follows: if a is even, then $cost_{\kappa(a)} \le 2\sum_{b=1}^{a/2} opt_{2b} \le 2\sum_{b=1}^{a/2}\gamma^b$, and if a is odd then $cost_{\kappa(a)} \le 2\sum_{b=1}^{(a-1)/2}\gamma^b + 4\lambda\gamma^{(a-1)/2}$. Summing up the geometric sequences, we thus get

$$cost_{\kappa(a)} \quad \le \quad \begin{cases} \dfrac{2\gamma^{a/2+1}}{\gamma - 1} & \text{if } a \text{ is even} \\ \dfrac{2\gamma^{(a-1)/2+1}}{\gamma - 1} + 4\lambda\gamma^{(a-1)/2} & \text{if } a \text{ is odd} \end{cases}$$

Fix some number of facilities j, and choose a such that $\kappa(a+1) \le j < \kappa(a)$. We want to show that $cost_j \le R \cdot opt_j$. By the construction, $F_j = F_{\kappa(a+1)}$, so $cost_j = cost_{\kappa(a+1)}$. We have two cases.

Suppose first that a is even. By the choice of j and the definition of $\kappa(a)$, we get $opt_j > \gamma^{a/2}$. Since $cost_j = cost_{\kappa(a+1)} \le 2\gamma^{a/2+1}/(\gamma - 1) + 4\lambda\gamma^{a/2}$, the ratio is

$$\frac{cost_j}{opt_j} \;\le\; \frac{2\gamma}{\gamma - 1} + 4\lambda \;=\; R.$$

If a is odd, then by the choice of j and the definition of $\kappa(a)$, we get $opt_j > \lambda\gamma^{(a-1)/2}$. Since $cost_j = cost_{\kappa(a+1)} \le 2\gamma^{(a+1)/2+1}/(\gamma - 1)$, the ratio is

$$\frac{cost_j}{opt_j} \;\le\; \frac{2\gamma^2}{(\gamma - 1)\lambda} \;=\; R,$$

completing the proof.

□

5 2-Competitive Incremental Medians for Equable Instances

In this section we present a construction of a 2-competitive incremental medians sequence for a special case where, for any fixed value of k, each customer has the same distance to the optimal k-median. More formally, suppose $(\mathcal{F}, \mathcal{C})$ is an instance of the medians problem with $n \le |\mathcal{C}|$ such that (i) for each $k = 1, 2, ..., n$ there exist an optimal k-median F^*_k such that all distances $d(x, F^*_k)$ are the same, for all $x \in \mathcal{C}$, and that (ii) for $k = n$ we have $d(x, F^*_n) = 0$ for all $x \in \mathcal{C}$ (or, equivalently, $\mathcal{C} \subseteq \mathcal{F}$.) An instance with this property will be called *equable*.

Our method is different from previous constructions of incremental medians, including the one from Section 4. Unlike in these previous approaches, we construct the sequence $F_1, F_2, ..., F_n$ *forward*, maintaining an invariant ensuring that we not only do well at step k, but also that we make good progress towards obtaining a low-cost l-median for all $l > k$.

Throughout this section, $(\mathcal{C}, \mathcal{F})$ denotes an equable instance of the medians problem. For each $k = 1, 2, ..., n$, let F^*_k be the optimal k-median such that $d(x, F^*_k) = \delta_k$ for all $x \in \mathcal{C}$. Thus $opt_k = |\mathcal{C}|\delta_k$ for all k. Without loss of generality, we can assume that $\delta_1 > \delta_2 > ... > \delta_n = 0$.

Incremental spanners. Suppose that for each $k = 1, 2, ..., n$ we have a family $\mathcal{S}_k \subseteq 2^{\mathcal{C}}$ of k sets that forms a partition of $\mathcal{C}$, that is, all sets in $\mathcal{S}_k$ are disjoint and $\bigcup_{A \in \mathcal{S}_k} A = \mathcal{C}$. (Our proof can be modified to work even if the sets in $\mathcal{S}_k$ are not disjoint.) For a set $X \subseteq \mathcal{C}$, define its k-*span* as

$$Span_k(X) \;\;=\;\; \bigcup \{A \in \mathcal{S}_i \mid i \ge k \;\&\; A \cap X \ne \emptyset\}.$$

A set $X \subseteq \mathcal{C}$ is called a k-*spanner* if $Span_k(X) = \mathcal{C}$. Note that if X is a k-spanner then it is also a j-spanner for any $j < k$. A sequence $X_1 \subseteq X_2 \subseteq ... \subseteq X_n$ is called an *incremental spanner* if for each $k = 1, 2, ..., n$, $|X_k| \le k$ and X_k is a k-spanner.

We now show how to construct an incremental spanner. For $X \subseteq \mathcal{C}$ and any $j = 1, 2, ..., n$, let $setscov_j(X)$ be the collection of sets in $\mathcal{S}_j$ covered by the j-span of X, that is

$$setscov_j(X) \;\;=\;\; \{A \in \mathcal{S}_j \mid A \subseteq Span_j(X)\}.$$

Note that $|setscov_j(X)| = j$ if and only if X is a j-spanner, because $\mathcal{S}_j$ covers $\mathcal{C}$.

We will construct the sets $\emptyset = X_0 \subseteq X_1 \subseteq ... \subseteq X_n$ so that, for each $k = 0, 1, 2, ..., n$, we will have $|X_k| \le k$ and the following invariant will hold:

$$|setscov_j(X_k)| \quad \ge \quad k, \quad \text{for all } j = k, k+1, ..., n. \tag{9}$$

Initially, for $k = 0$, we set $X_0 = \emptyset$, and (9) holds trivially. Suppose we have $X_0, X_1, ..., X_{k'}$, for some $k' < n$ and that (9) holds for $k = 0, 1, ..., k'$. This implies, in particular, that $|setscov_{k'}(X_{k'})| \ge k'$, that is, $X_{k'}$ is a k'-spanner. Thus $X_{k'}$ is also a k-spanner for all $k \le k'$. Let l be the minimum index for which $X_{k'}$ is *not* an l-spanner, that is $\mathcal{C} - Span_l(X_{k'}) \neq \emptyset$. By the choice of l, we have $l > k'$. Pick any $x \in \mathcal{C} - Span_l(X_{k'})$ and take $X_{k'+1} = X_{k'} \cup \{x\}$. Clearly, $|X_{k'+1}| \le k' + 1$.

We now show that (9) holds for $k = k'+1$. By the choice of l, for $j = k'+1, k'+2, ..., l-1$, $X_{k'}$ is a j-spanner. Therefore $X_{k'+1}$ is also a j-spanner, and thus (9) holds. Consider any $j \ge l \ge k' + 1$. Let $A \in \mathcal{S}_j$ be the set for which $x \in A$. By induction, since $x \in \mathcal{C} - Span_j(X_{k'})$, we have $A \notin setscov_j(X_{k'})$. But now $x \in X_{k'+1}$, so $A \in setscov_j(X_{k'+1})$, and we get $|setscov_j(F_{k'+1})| \ge |setscov_j(F_{k'})| + 1 \ge k' + 1$. This completes the proof that our construction preserves invariant (9).

By (9), for each k we have $|setscov_k(X_k)| \ge k$, and thus X_k is a k-spanner. We can conclude then that $X_1, X_2, ..., X_n$ is an incremental spanner.

Incremental medians. We now show how to use incremental spanners to construct incremental medians. For $k = 1, 2, ..., n$, assign each customer $x \in \mathcal{C}$ to its closest facility $f \in F^*_k$ (that is, $d_{xf} = \delta_k$), breaking ties arbitrarily. Define C^f_k to be the set of customers assigned to f, and let $\mathcal{S}_k = \{C^f_k \mid f \in F^*_k\}$. Then each $\mathcal{S}_k$ contains k sets and forms a partition of $\mathcal{C}$. As we showed above, for these partitions $\mathcal{S}_1, \mathcal{S}_2, ..., \mathcal{S}_n$ there exists an incremental spanner $F_1, F_2, ..., F_n$.

We claim that $F_1, F_2, ..., F_n$ is a 2-competitive incremental medians sequence. Consider some fixed k. Since F_k is a k-spanner, for each customer $x \in \mathcal{C}$ there is $i \ge k$, $f \in F^*_i$ and $y \in F_k$ such that both $x, y \in C^f_i$. Thus $d(x, F_k) \le d_{xy} \le d_{xf} + d_{yf} = 2\delta_i \le 2\delta_k$. This implies that $cost(F_k) \le 2m\delta_k = 2\,opt_k$, and the claim follows.

Summarizing, we obtain the following result:

Theorem 5. *Any equable instance $(\mathcal{C}, \mathcal{F})$ of the medians problem has a 2-competitive incremental medians sequence.*

6 Final Comments

We improved both the lower and upper bounds for incremental medians, from 2 to 2.01 and from 8 to $2 + 4\sqrt{2} \approx 7.656$, respectively, thus proving that neither 2 nor 8 are the "right" bounds for this problem. (By optimizing the the parameters in Section 3 it is possible to improve the lower bound slightly, to about 2.01053.) In addition to its own independent interest, closing or significantly reducing the remaining gap would shed more light on the computational hardness of approximating incremental medians, as it would show to what degree the difficulty of the problem can be attributed to non-existence of incremental median sequences with small competitive ratios.

The expected values in the proof of Lemma 2 can be computed in polynomial-time, and thus our probabilistic construction can be derandomized using the method of conditional expectations. However, since our improvement is relatively minor, we did not pursue this direction of research, nor possible implications for upper bounds achievable in polynomial time.

We believe that some of the ideas in the paper can be used to prove even better bounds. In the upper bound proof in Section 4 we construct our sequence backwards, starting with all facilities, and gradually extracting smaller and smaller facility sets, two at a time. By extending the probabilistic construction to more than two steps at a time, we should be able to get a better bound. Even our two-step method still might have room for improvement, as the two choices for F and G considered in the proof of Theorem 3 are not "balanced", that is, the bounds on the cost of F and G in the two cases are not the same. Also, our construction of a 2-competitive incremental medians sequence for equable spaces is very different from previous constructions and we believe that its basic idea will be useful in improving the upper bound for general spaces.

Our lower bound argument uses only three steps, for $k = 1, 2, n$. It should be possible to improve our bound by using either $k > 2$ as the intermediate number of facilities or more (perhaps an unbounded number of) steps. Both ideas lead to difficulties that we were not able to overcome at this time. In a three-step strategy using $k = 1, k', n$ with $k' > 2$, an algorithm can place facilities $2, .., k'$ optimally (given the choice of the first facility), and thus increasing k' seems only to help the algorithm. A strategy that uses additional steps leads to a different problem. Average costs for the customers must decrease with k, and thus introducing additional steps creates shortcuts via optimal k'-medians for large k', reducing the algorithm's cost for small values of k.

The result from Section 5 may also be useful for lower bound proofs, as it shows that in "hard" instances, for a fixed k, the optimal customers' costs should be significantly different.

References

[1] Vijay Arya, Naveen Garg, Rohit Khandekar, Adam Meyerson, Kamesh Munagala, and Vinayaka Pandit. Local search heuristic for k-median and facility location problems. In *Proc. 33rd Symp. Theory of Computing (STOC)*, pages 21–29. ACM, 2001.

[2] Vijay Arya, Naveen Garg, Rohit Khandekar, Adam Meyerson, Kamesh Munagala, and Vinayaka Pandit. Local search heuristics for k-median and facility location problems. *SIAM Journal on Computing*, 33(3):544–562, 2004.

[3] Marek Chrobak and Claire Kenyon. Competitiveness via doubling. *SIGACT News*, pages 115–126, 2006.

[4] Marek Chrobak, Claire Kenyon, John Noga, and Neal Young. Online medians via online bidding. In *Proc. 7th Latin American Theoretical Informatics Symp. (LATIN)*, volume 3887 of *Lecture Notes in Comput. Sci.*, pages 311–322, 2006.

[5] Guolong Lin, Chandrashekha Nagarajan, Rajmohan Rajamaran, and David P. Williamson. A general approach for incremental approximation and hierarchical clustering. In *Proc. 17th Symp. on Discrete Algorithms (SODA)*, pages 1147–1156, 2006.

[6] Ramgopal R Mettu and C. Greg Plaxton. The online median problem. In *Proc. 41st Symp. Foundations of Computer Science (FOCS)*, pages 339–348. IEEE, 2000.

[7] Ramgopal R. Mettu and C. Greg Plaxton. The online median problem. *SIAM J. Comput.*, 32:816–832, 2003.