On the Competitive Ratio for Online Facility Location*

Dimitris Fotakis

Max-Planck-Institut für Informatik
Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany
Email: fotakis@mpi-sb.mpg.de

Abstract. We consider the problem of Online Facility Location, where demands arrive online and must be irrevocably assigned to an open facility upon arrival. The objective is to minimize the sum of facility and assignment costs. We prove that the competitive ratio for Online Facility Location is $\Theta(\frac{\log n}{\log\log n})$. On the negative side, we show that no randomized algorithm can achieve a competitive ratio better than $\Omega(\frac{\log n}{\log\log n})$ against an oblivious adversary even if the demands lie on a line segment. On the positive side, we present a deterministic algorithm achieving a competitive ratio of $\mathrm{O}(\frac{\log n}{\log\log n})$. The analysis is based on a hierarchical decomposition of the optimal facility locations such that each component either is relatively well-separated or has a relatively large diameter, and a potential function argument which distinguishes between the two kinds of components.

1 Introduction

The (metric uncapacitated) Facility Location problem is, given a metric space along with a facility cost for each point and a (multi)set of demand points, to find a set of facility locations which minimize the sum of facility and assignment costs. The assignment cost of a demand point is the distance to the nearest facility. Facility Location provides a simple and natural model for network design and clustering problems and has been the subject of intensive research over the last decade (e.g., see [17] for a survey and [9] for approximation algorithms and applications).

The definition of *Online Facility Location* [16] is motivated by practical applications where either the demand set is not known in advance or the solution must be constructed incrementally using limited information about future demands. In Online Facility Location, the demands arrive one at a time and must be irrevocably assigned to an open facility without any knowledge about future demands. The objective is to minimize the sum of facility and assignment costs, where each demand's assignment cost is the distance to the facility it is assigned to.

We evaluate the performance of online algorithms using *competitive analysis* (e.g., [5]). An online algorithm is c-competitive if for all instances, the cost incurred by the algorithm is at most c times the cost incurred by an optimal offline algorithm, which has full knowledge of the demand sequence, on the same instance. We always use n to denote the number of demands.

Previous Work. In the offline case, where the demand set is fully known in advance, there are constant factor approximation algorithms based on Linear Programming rounding (e.g., [18]), local search (e.g., [10]), and the primal-dual schema (e.g., [12]). The best known polynomial-time algorithm achieves an approximation ratio of 1.52 [14], while no polynomial-time algorithm can achieve an approximation ratio less than 1.463 unless $\mathrm{NP} = \mathrm{DTIME}(n^{\mathrm{O}(\log\log n)})$ [10].

Online Facility Location was first defined and studied in [16], where a simple randomized algorithm is shown to achieve a constant performance ratio if the demands, which are adversarially selected, arrive in random order. In the standard framework of competitive analysis, where not

* This work was partially supported by the Future and Emerging Technologies programme of the EU under contract number IST-1999-14186 (ALCOM–FT).

only the demand set but also the demand order is selected by an oblivious adversary, the same algorithm achieves a competitive ratio of $O(\frac{\log n}{\log\log n})^1$. It is also shown a lower bound of $\Omega(\log^* n)$ on the competitive ratio of any online algorithm, where $\log^*$ is the inverse Ackerman function.

Online Facility Location should not be confused with the problem of Online Median [15]. In Online Median, the demand set is fully known in advance and the number of facilities increases online. An O(1)-competitive algorithm is known for Online Median [15].

Online Facility Location bears a resemblance to the extensively studied problem of Online File Replication (e.g., [4, 2, 1, 13, 8]). In Online File Replication, we are given a metric space, a point initially holding the file, and a replication cost factor. Read requests are generated by points in an online fashion. Each request accesses the nearest file copy at a cost equal to the corresponding distance. In between requests, the file may be replicated to a set of points at a cost equal to the replication cost factor times the total length of the minimum Steiner tree connecting the set of points receiving the file to at least one point already holding the file. Similarly to Facility Location, File Replication asks for a set of file locations which minimize the sum of replication and access costs. The important difference is that the cost of each facility only depends on the location, while the cost of each replication depends on the set of points which hold the file and the set of points which receive the file.

Online File Replication is a generalization of Online Steiner Tree [11]. Hence, there are metric spaces in which no randomized online algorithm can achieve a competitive ratio better than $\Omega(\log n)$ against an oblivious adversary. They are known both a randomized [4] and a deterministic [2] algorithm achieving a competitive ratio of $O(\log n)$ for the more general problem of Online File Allocation. For trees and rings, algorithms of constant competitive ratio are known [1, 13, 8].

Contribution. We prove that the competitive ratio for Online Facility Location is $\Theta(\frac{\log n}{\log\log n})$. On the negative side, we show that no randomized algorithm can achieve a competitive ratio better than $\Omega(\frac{\log n}{\log\log n})$ against an oblivious adversary even if the metric space is a line segment. The only previously known lower bound was $\Omega(\log^* n)$ [16]. On the positive side, we present a deterministic algorithm achieving a competitive ratio of $O(\frac{\log n}{\log\log n})$ in every metric space. To the best of our knowledge, this is the first deterministic upper bound on the competitive ratio for Online Facility Location.

As for the analysis, the technique of [2], which is based on a hierarchical decomposition/cover of the optimal file locations such that each component's diameter is not too large, cannot be adapted to yield a sub-logarithmic competitive ratio for Online Facility Location. On the other hand, it is not difficult to show that our algorithm achieves a competitive ratio of $O(\frac{\log n}{\log\log n})$ for instances whose optimal solution consists of a single facility. To establish a tight bound for general instances, we show that any metric space has a hierarchical cover with the additional property that any component either is relatively well-separated or has a relatively large diameter. Then, we prove that the sub-instances corresponding to well-separated components can be treated as essentially independent instances whose optimal solutions consist of a single facility, and we bound the additional cost incurred by the algorithm because of the sub-instances corresponding to large diameter components.

Problem Definition. The problem of Online Facility Location is formally defined as follows. We are given a metric space $\mathcal{M} = (C, d)$, where C denotes the set of points and $d : C \times C \mapsto \mathbb{R}^+$ denotes the distance function which is symmetric and satisfies the triangle inequality. For each point $v \in C$, we are also given the cost f_v of opening a facility at v. The demand sequence consists of (not necessarily distinct) points $w \in C$. When a demand w arrives, the algorithm can open some new facilities. Once opened, a facility cannot be closed. Then, w must be irrevocably assigned to

[1] Only a logarithmic competitive ratio is claimed in [16]. However, a competitive ratio of $O(\frac{\log n}{\log\log n})$ follows from a simple modification of the same argument.

the nearest facility. If w is assigned to a facility at v, w's assignment cost is $d(w,v)$. The objective is to minimize the sum of facility and assignment costs.

Throughout this paper, we only consider unit demands by allowing multiple demands to be located at the same point. We always use n to denote the total number of demands. We distinguish between the case of uniform facility costs, where the cost of opening a facility, denoted by f, is the same for all points, and the general case of non-uniform facility costs, where the cost of opening a facility depends on the point.

Notation. A metric space $\mathcal{M} = (C, d)$ is usually identified by its point set C. For a subspace $C' \subseteq C$, $D(C') = \max_{u,v \in C'} \{d(u,v)\}$ denotes the diameter of C'. For a point $u \in C$ and a subspace $C' \subseteq C$, $d(C', u) = \min_{v \in C'} \{d(v,u)\}$ denotes the distance from u to the nearest point in C'. We use the convention that $d(u, \emptyset) = \infty$. For subspaces $C', C'' \subseteq C$, $d(C', C'') = \min_{u \in C''} \{d(C', u)\}$ denotes the minimum distance between a point in C' and a point in C''. For a point $u \in C$ and a non-negative number r, $B(u,r)$ denotes the ball of center u and radius r, $B(u,r) = \{v \in C : d(u,v) \leq r\}$.

2 A Lower Bound on the Competitive Ratio

In this section, we restrict our attention to uniform facility costs and instances whose optimal solution consists of a single facility. These assumptions can only strengthen the proven lower bound.

Theorem 1. *No randomized algorithm for Online Facility Location can achieve a competitive ratio better than $\Omega(\frac{\log n}{\log \log n})$ against an oblivious adversary even if the metric space is a line segment.*

Proof Sketch. We first prove that the lower bound holds if the metric space is a complete binary Hierarchically Well-Separated Tree (HST) [3]. Let T be a complete binary rooted tree of height h such that (i) the distance from the root to each of its children is D, and (ii) on every path from the root to a leaf, the edge length drops by a factor exactly m on every step. The height of a vertex is the number of edges on the path to the root. Every non-leaf vertex has exactly two children and every leaf has height exactly h. The distance from a vertex of height i to each of its children is exactly $\frac{D}{m^i}$. Let f be the cost of opening a new facility, which is the same for every vertex of T.

For a vertex v, let T_v denote the subtree rooted at v. The lower bound is based on the following property of T: The distance from a vertex v of height i to any vertex in T_v is at most $\frac{m}{m-1}\frac{D}{m^i}$, while the distance from v to any vertex not in T_v is at least $\frac{D}{m^{i-1}}$.

By Yao's principle (e.g., [5, Chapter 8]), it suffices to show that there is a probability distribution over demand sequences for which the ratio of the expected cost of any deterministic online algorithm to the expected optimal cost is $\Omega(\frac{\log n}{\log \log n})$.

We define an appropriate probability distribution by considering demand sequences divided into $h+1$ phases. Phase 0 consists of a single demand at the root v_0. After the end of phase i, if v_i is not a leaf, the adversary proceeds to the next phase by selecting v_{i+1} uniformly at random and independently (u.i.r.) among the two children of v_i. Phase $i+1$ consists of m^{i+1} consecutive demands at v_{i+1}.

The total number of demands is at most $m^h \frac{m}{m-1}$, which must not exceed n. The optimal solution opens a single facility at v_h and, for each phase i, incurs an assignment cost no greater than $D\frac{m}{m-1}$. Therefore, the optimal cost is at most $f + hD\frac{m}{m-1}$.

Let Alg be any deterministic online algorithm. We fix the adversary's random choices $v_0, \ldots, v_i$ up to phase i, $0 \leq i \leq h-1$, (equivalently, we fix T_{v_i}), and we consider the expected cost (conditional on T_{v_i}) incurred by Alg for demands and facilities not in $T_{v_{i+1}}$. If Alg has no facilities in T_{v_i} when the first demand at v_{i+1} arrives, the assignment cost of demands at $v_i \in T_{v_i} \setminus T_{v_{i+1}}$ is at least

$A \leftarrow \emptyset$; $L \leftarrow \emptyset$; `/* Initialization */`
For each demand w:
  $r_w \leftarrow \frac{d(A,w)}{x}$; $B_w \leftarrow \{u \in L \cup \{w\} : d(w,u) \leq r_w\}$; $\mathsf{Pot}(B_w) \leftarrow \sum_{u \in B_w} d(A,u)$;
  if $\mathsf{Pot}(B_w) \geq f$ then `/* A new facility is opened */`
    if $d(A,w) < f$ then
      Let ν be the smallest integer:
        either there exists exactly one $u \in B_w$ such that
          $\mathsf{Pot}\big(B_w \cap B\big(u, \frac{r_w}{2^\nu}\big)\big) > \frac{\mathsf{Pot}(B_w)}{2}$,
        or, for any $u \in B_w$, $\mathsf{Pot}\big(B_w \cap B\big(u, \frac{r_w}{2^{\nu+1}}\big)\big) \leq \frac{\mathsf{Pot}(B_w)}{2}$.
      Let $\hat{w}$ be any demand in B_w: $\mathsf{Pot}\big(B_w \cap B\big(\hat{w}, \frac{r_w}{2^\nu}\big)\big) > \frac{\mathsf{Pot}(B_w)}{2}$.
    else $\hat{w} \leftarrow w$;
    $A \leftarrow A \cup \{\hat{w}\}$; $L \leftarrow L \setminus B_w$;
  else $L \leftarrow L \cup \{w\}$; `/* w is marked unsatisfied */`
  Assign w to the nearest facility in A.

Fig. 1. The algorithm Deterministic Facility Location – DFL.

mD. Otherwise, since v_{i+1} is selected u.i.r. among v_i's children, with probability at least $\frac{1}{2}$, there is at least one facility in $T_{v_i} \setminus T_{v_{i+1}}$. Therefore, for any fixed T_{v_i}, the (conditional) expected cost incurred by Alg for demands and facilities not in $T_{v_{i+1}}$ is at least $\min\{mD, \frac{f}{2}\}$ plus the cost for demands and facilities not in T_{v_i}. Since this holds for any fixed choice of $v_0, \ldots, v_i$ (equivalently, for any fixed T_{v_i}), the (unconditional) expected cost incurred by Alg for demands and facilities not in $T_{v_{i+1}}$ is at least $\min\{mD, \frac{f}{2}\}$ plus the (unconditional) expected cost for demands and facilities not in T_{v_i}. Hence, at the beginning of phase i, $0 \leq i \leq h$, the expected cost incurred by Alg for demands and facilities not in T_{v_i} is at least $i \min\{mD, \frac{f}{2}\}$. For the last phase, Alg incurs a cost no less than $\min\{mD, f\}$ inside T_{v_h}.

For $m = h$ and $D = \frac{f}{h}$, the total expected cost of Alg is at least $\frac{h+2}{2}hD$, while the optimal cost is at most $\frac{2h-1}{h-1}hD$. For the chosen value of h, the quantity $\frac{h^{h+1}}{h-1}$ must not exceed n. Setting $h = \left\lfloor \frac{\log n}{\log\log n} \right\rfloor$ yields the claimed lower bound.

To conclude the proof, we consider the following embedding of T in a line segment. The root is mapped to 0 (i.e., the center of the segment). Let v be a vertex of height i mapped to $\tilde{v}$. Then, v's left child is mapped to $\tilde{v} - \frac{D}{m^i}$ and v's right child is mapped to $\tilde{v} + \frac{D}{m^i}$. It can be shown that, for any $m \geq 4$, this embedding results in a hierarchically well-separated metric space. □

3 A Deterministic Algorithm for Uniform Facility Costs

In this section, we present the algorithm Deterministic Facility Location – DFL (Fig. 1) and prove that its competitive ratio is $O(\frac{\log n}{\log\log n})$.

Outline. The algorithm maintains its *facility configuration* A and the set L of *unsatisfied demands*, which are the demands not having contributed towards opening a new facility so far. A new demand w is marked unsatisfied and added to L only if no new facilities are opened when w arrives. Each unsatisfied demand $u \in L$ can contribute an amount of $d(A,u)$ to the cost of opening a new facility in its neighborhood. We refer to the quantity $d(A,u)$ as the *potential of* u. Only unsatisfied demands and the demand currently being processed have non-zero potential. For a set S consisting of demands of non-zero potential, let $\mathsf{Pot}(S) = \sum_{u \in S} d(A,u)$ be the potential of S.

The high level idea is to keep a balance between the algorithm's assignment and facility costs. For each demand w, the algorithm computes the set B_w consisting of w and the unsatisfied demands at distance no greater than $\frac{d(A,w)}{x}$ from w, where x is a sufficiently large constant. If B_w's potential is less than f, w is assigned to the nearest facility, marked unsatisfied and added to L. Otherwise, the algorithm opens a new facility at an appropriate location $\hat{w} \in B_w$ and assigns w to it. In this case, the demands in B_w are marked satisfied and removed from L. The location $\hat{w}$ is chosen as the center of a smallest radius ball/subset of B_w contributing more than half of B_w's potential.

An Overview of the Analysis. For an arbitrary sequence of n demands, we compare the algorithm's cost with the cost of a fixed offline optimal solution. The optimal solution is determined by k facility locations $c_1^*, c_2^*, \ldots, c_k^*$. The set of optimal facilities is denoted by C^*. Each demand u is assigned to the nearest facility in C^*. Hence, C^* defines a partition of the demand sequence into optimal clusters $C_1, C_2, \ldots, C_k$. Let $d_u^* = d(C^*, u)$ denote the assignment cost of u in the optimal solution, let $\mathsf{S}^* = \sum_u d_u^*$ be the total optimal assignment cost, let $\mathsf{F}^* = kf$ be the total optimal facility cost, and let $\sigma^* = \frac{\mathsf{S}^*}{n}$ be the average optimal assignment cost.

Let ρ, ψ denote a fixed pair of integers such that $\rho^\psi > n$. For any integer j, $0 \leq j \leq \psi$, let $r(j) = \rho^j \sigma^*$. We also define $r(-1) = 0$ and $r(\psi + 1) = \infty$. We observe that, for any demand u, $d_u^* < r(\psi)$. Let λ be some appropriately large constant, and, for any integer j, $-1 \leq j \leq \psi + 1$, let $R(j) = \lambda\, r(j)$. Throughout the analysis of DFL, we use $\lambda = 3x + 2$.

The Case of a Single Optimal Cluster. We first restrict our attention to instances whose optimal solution consists of a single facility c^*. The convergence of A to c^* is divided into $\psi + 2$ phases, where the current phase ℓ, $-1 \leq \ell \leq \psi$, starts just after the first facility within a distance of $R(\ell + 1)$ from c^* is opened and ends when the first facility within a distance of $R(\ell)$ from c^* is opened. In other words, the current phase ℓ lasts as long as $d(A, c^*) \in (R(\ell), R(\ell + 1)]$.

The demands arriving in the current phase ℓ and the demands remaining in L from the previous phase are partitioned into *inner* demands, whose optimal assignment cost is less than $r(\ell)$, and *outer* demands. The last phase ($\ell = -1$) never ends and only consists of outer demands.

For any outer demand u, $d(A, u)$ is at most $\lambda\sigma^* + (\lambda\rho + 1)d_u^*$ (Ineq. (3)). Hence, the assignment cost of an outer demand arriving in phase ℓ can be charged to its optimal assignment cost. We charge the total assignment cost of inner demands arriving in phase ℓ and the total facility cost incurred by the algorithm in phase ℓ to the optimal facility cost and the optimal assignment cost of the outer demands marked satisfied in phase ℓ.

The set of inner demands is included in a ball of center c^* and radius $r(\ell)$. If $R(\ell)$ is large enough compared to $r(\ell)$ (namely, if λ is chosen sufficiently large), we can think of the inner demands as being essentially located at c^*, because they are much closer to each other than to the current facility configuration A. Hence, we refer to the total potential accumulated by unsatisfied inner demands as the potential accumulated by c^* or simply, the potential of c^*. For any inner demand w, B_w includes the entire set of unsatisfied inner demands. Therefore, the potential accumulated by c^* is always less than f (Lemma 3).

However, a new facility may decrease the potential of c^*, because (i) it may be closer to c^*, and (ii) some unsatisfied inner demands may contribute their potential towards opening the new facility, in which case they are marked satisfied and removed from L. As a result, the upper bound of f on the potential accumulated by c^* cannot be directly translated into an upper bound on the total assignment cost of the inner demands arriving in phase ℓ as in [16].

Each time a new facility is opened, the algorithm incurs a facility cost of f and an assignment cost no greater than $\frac{f}{x}$. The algorithm must also be charged with an additional cost accounting for the decrease in the potential accumulated by c^*, which cannot exceed f. Hence, for each new facility, the algorithm is charged with a cost no greater than $\frac{2x+1}{x} f$.

Using the fact that $R(\ell)$ is much larger than $r(\ell)$, we show that if the inner demands included in B_w contribute more than half of B_w's potential, the new facility at $\hat{w}$ is within a distance of $R(\ell)$ from c^* (Lemma 4). In this case (Lemma 8, Case Isolated.B), the current phase ends and the algorithm's cost is charged to the optimal facility cost. Otherwise (Lemma 8, Case Isolated.A), the algorithm's cost is charged to the potential of the outer demands included in B_w, which is at least $f/2$. The optimal facility cost is charged $O(\psi)$ times and the optimal assignment cost is charged $O(\lambda\rho)$ times. Hence, setting $\psi = \rho = O(\frac{\log n}{\log \log n})$ yields the desired competitive ratio.

The General Case. If the optimal solution consists of $k > 1$ facilities $c_1^*, \ldots, c_k^*$, the demands are partitioned into the optimal clusters $C_1, \ldots, C_k$. The convergence of A to an optimal facility c_i^* is divided into $\psi + 2$ phases, where the current phase ℓ_i, $-1 \leq \ell_i \leq \psi$, lasts as long as $d(A, c_i^*) \in (R(\ell_i), R(\ell_i + 1)]$. For the current phase ℓ_i, the demands of C_i are again partitioned into inner and outer demands, and the inner demands of C_i can be thought of as being essentially located at c_i^*.

As before, the potential accumulated by an optimal facility c_i^* cannot exceed f. However, a single new facility can decrease the potential accumulated by many optimal facilities. Therefore, if we bound the decrease in the potential of each optimal facility separately and charge the algorithm with the total additional cost, we can only guarantee a logarithmic upper bound on the competitive ratio. To establish a tight bound, we show that the average (per new facility) decrease in the total potential accumulated by optimal facilities is $O(f)$.

We first observe that as long as the distance from the algorithm's facility configuration A to a set of optimal facilities K is large enough compared to the diameter of K, the inner demands assigned to facilities in K are much closer to each other than to A. Consequently, we can think of the inner demands assigned to K as being located at some optimal facility $c_K^* \in K$. Therefore, the total potential accumulated by optimal facilities in K is always less than f (Lemma 3). This observation naturally leads to the definition of an (optimal facility) *coalition* (Definition 2).

Our potential function argument is based on a *hierarchical cover* (Definition 1) of the subspace C^* comprising the optimal facility locations. Given a facility configuration A, the hierarchical cover determines a minimal collection of *active* coalitions which form a partition of C^* (Definition 3).

A coalition is *isolated* if it is well-separated from any other disjoint coalition, and *typical* otherwise. A new facility can decrease the potential accumulated by at most one isolated active coalition. Therefore, for each new facility, the decrease in the total potential accumulated by isolated active coalitions is at most f (Lemma 8, Case Isolated).

On the other hand, a new facility can decrease the potential accumulated by several typical active coalitions. We prove that any metric space has a hierarchical cover such that each component either is relatively well-separated or has a relatively large diameter (i.e., its diameter is within a constant factor from its parent's diameter (Lemma 1). Typical active coalitions correspond to the latter kind of components. Hence, we obtain a bound on the relative length of the interval for which an active coalition remains typical (Lemma 2), which can be translated into a bound of $O(f)$ on the total decrease in the potential accumulated by an active coalition, while the coalition remains typical (potential function component $\Xi_K^{(2)}$ and Lemma 7).

In the remaining paragraphs, we prove the following theorem by turning the aforementioned intuition into a formal potential function argument.

Theorem 2. *For any constant $x \geq 10$, the competitive ratio of Deterministic Facility Location is $O(\frac{\log n}{\log \log n})$.*

Hierarchical Covers and Optimal Facility Coalitions. We start by showing that any metric space has a hierarchical cover with the desired properties.

Definition 1. A hierarchical cover of a metric space C *is a collection* $\mathcal{K} = \{K_1, \ldots, K_m\}$ *of non-empty subsets of* C *which can be represented by a rooted tree* $T_{\mathcal{K}}$ *in the following sense:*

(A) C belongs to $\mathcal{K}$ and corresponds to the root of $T_{\mathcal{K}}$.
(B) For any $K \in \mathcal{K}$, $|K| > 1$, $\mathcal{K}$ contains sets $K_1, \ldots, K_\mu$, each of diameter less than $D(K)$, which form a partition of K. The sets $K_1, \ldots, K_\mu$ correspond to the children of K in $T_{\mathcal{K}}$.

We use $\mathcal{K}$ and its tree representation $T_{\mathcal{K}}$ interchangeably. By definition, every non-leaf set has at least two children. Therefore, $T_{\mathcal{K}}$ has at most $2|C|-1$ nodes. For a set K different from the root, we use P_K to denote the immediate ancestor/parent of K in $T_{\mathcal{K}}$. Our potential function argument is based on the following property of metric spaces.

Lemma 1. *For any metric space C and any $\gamma \geq 16$, there exists a hierarchical cover $T_{\mathcal{K}}$ of C such that for any set K different from the root, either $D(K) > \frac{D(P_K)}{\gamma^2}$ or $d(K, C \setminus K) > \frac{D(P_K)}{4\gamma}$.*

Proof Sketch. Let C be any metric space, and let $D = D(C)$. We first show that, for any integer $i \geq 0$, C can be partitioned into a collection of *level i groups* $G_1^i, \ldots, G_m^i$ such that (i) for any $j_1 \neq j_2$, $d(G_{j_1}^i, G_{j_2}^i) > \frac{D}{4\gamma^i}$, and (ii) if $D(G_j^i) > \frac{D}{\gamma^i}$, then G_j^i does not contain any subset $G \subseteq G_j^i$ such that both $D(G) \leq \frac{D}{\gamma^{i+1}}$ and $d(G, G_j^i \setminus G) > \frac{D}{4\gamma^i}$. Since the collection of level i groups is a partition of C, for any G_j^i, $d(G_j^i, C \setminus G_j^i) > \frac{D}{4\gamma^i}$.

Level i groups are further partitioned into *level i components* $K_1^i, \ldots, K_{m'}^i$ such that (i) $D(K_j^i) \leq \frac{D}{\gamma^i}$, and (ii) either $D(K_j^i) > \frac{D}{\gamma^{i+1}}$ or $d(K_j^i, C \setminus K_j^i) > \frac{D}{4\gamma^i}$. To ensure a hierarchical structure, we proceed inductively in a bottom-up fashion. We create a single level i component for each level i group G_j^i of diameter no greater than $\frac{D}{\gamma^i}$. We recall that $d(G_j^i, C \setminus G_j^i) > \frac{D}{4\gamma^i}$. If $D(G_j^i) > \frac{D}{\gamma^i}$, G_j^i is partitioned into level i components of diameter in the interval $(\frac{D}{\gamma^{i+1}}, \frac{D}{\gamma^i}]$. For $\gamma \geq 16$, such a partition exists, because G_j^i does not contain any well-separated subsets of small diameter. Finally, we eliminate multiple occurrences of the same component at different levels. □

Definition 2. *A set of optimal facilities $K \subseteq C^*$ with representative $c_K^* \in K$ is* a coalition with respect to the facility configuration A *if $d(A, c_K^*) \geq \lambda D(K)$. A coalition K is called* isolated *if $d(K, C^* \setminus K) \geq 2\, d(A, c_K^*)$, and* typical *otherwise. A coalition K* becomes broken *as soon as $d(A, c_K^*) < \lambda D(K)$.*

Given a hierarchical cover $T_{\mathcal{K}}$ of the subspace C^* comprising the optimal facility locations, we choose an arbitrary optimal facility as the representative of each set K. The representative of K always remains the same and is denoted by c_K^*. Then, $T_{\mathcal{K}}$ can be regarded as a system of optimal facility coalitions which hierarchically covers C^*. The current facility configuration A defines a minimal collection of active coalitions which form a partition of C^*.

Definition 3. *Given a hierarchical cover $T_{\mathcal{K}}$ of C^*, a coalition $K \in T_{\mathcal{K}}$ is* an active coalition with respect to A *if $d(A, c_K^*) \geq \lambda D(K)$ and for any other coalition K' on the path from K to the root of $T_{\mathcal{K}}$, $d(A, c_{K'}^*) < \lambda D(K')$.*

Lemma 2. *For any $\gamma \geq 8\lambda$, there is a hierarchical cover $T_{\mathcal{K}}$ of C^* such that if K is a typical active coalition with respect to the facility configuration A, then $\lambda \frac{D(P_K)}{\gamma^2} < d(A, c_K^*) < (\lambda + 1) D(P_K)$.*

Proof. For some $\gamma \geq 8\lambda$, let $T_{\mathcal{K}}$ be the hierarchical cover of C^* implied by Lemma 1. We show that $T_{\mathcal{K}}$ has the claimed property. The root of $T_{\mathcal{K}}$ is an isolated coalition by definition. Hence, we can restrict our attention to coalitions $K \in T_{\mathcal{K}}$ different from the root for which the parent function P_K is well-defined.

Since K is an active coalition, its parent coalition P_K must have become broken. The upper bound on $d(A, c_K^*)$ follows from the triangle inequality and the fact that c_K^* also belongs to P_K.

For the lower bound, we consider two cases. If K has a relatively large diameter ($D(K) > \frac{D(P_K)}{\gamma^2}$), the lower bound on $d(A, c_K^*)$ holds as long as K remains a coalition. If K is relatively well-separated ($d(K, C^* \setminus K) > \frac{D(P_K)}{4\gamma}$) and the lower bound on $d(A, c_K^*)$ does not hold, we conclude that $2\, d(A, c_K^*) < d(K, C^* \setminus K)$ (K is an isolated coalition), which is a contradiction. □

Notation. The set of active coalitions with respect to the current facility configuration A is denoted by $\mathsf{Act}(A)$. For a coalition K, ℓ_K denotes the index of the current phase. Namely, ℓ_K is equal to the integer j, $-1 \leq j \leq \psi$, such that $d(A, c_K^*) \in (R(j), R(j+1)]$. If $d(A, c_K^*) > R(\psi)$, $\ell_K = \psi$ (the first phase), while if $d(A, c_K^*) \leq R(0)$, $\ell_K = -1$ (the last phase). Let $C_K = \bigcup_{c_i^* \in K} C_i$ be the optimal cluster corresponding to K. Since $\mathsf{Act}(A)$ is always a partition of C^*, the collection $\{C_K : K \in \mathsf{Act}(A)\}$ is a partition of the demand sequence. For the current phase ℓ_K, the demands of C_K are partitioned into inner demands $\mathsf{In}(K) = \{u \in C_K : d_u^* < r(\ell_K)\}$ and outer demands $\mathsf{Out}(K) = C_K \setminus \mathsf{In}(K)$. Let also $\Lambda_K = L \cap \mathsf{In}(K)$ be the set of unsatisfied inner demands assigned to K.

We should emphasize that ℓ_K, $\mathsf{In}(K)$, $\mathsf{Out}(K)$, and Λ_K depend on the current facility configuration A. In addition, Λ_K depends on the current set of unsatisfied demands L. For simplicity of notation, we omit the explicit dependence on A and L by assuming that while a demand w is being processed, ℓ_K, $\mathsf{In}(K)$, $\mathsf{Out}(K)$, and Λ_K keep the values they had when w arrived.

Properties. Let K be a coalition with respect to the current facility configuration A. Then, $d(A, c_K^*) \geq \lambda \max\{D(K), r(\ell_K)\}$. The diameter of the subspace comprising the inner demands of K is $D(\mathsf{In}(K)) < 3\max\{D(K), r(\ell_K)\}$. We repeatedly use the following inequalities. Let u be any demand in C_K and let $c_u^* \in K$ be the optimal facility to which u is assigned. Then,

$$d(A, u) \leq d(A, c_K^*) + d(c_K^*, c_u^*) + d(c_u^*, u) \leq d(A, c_K^*) + D(K) + d_u^* \leq \tfrac{\lambda+1}{\lambda} d(A, c_K^*) + d_u^* \qquad (1)$$

If u is an inner demand of K ($u \in \mathsf{In}(K)$),

$$d(u, c_K^*) \leq d(u, c_u^*) + d(c_u^*, c_K^*) < r(\ell_K) + D(K) \leq 2\max\{D(K), r(\ell_K)\} \leq \tfrac{2}{\lambda} d(A, c_K^*) \qquad (2)$$

If u is an outer demand of K ($u \in \mathsf{Out}(K)$),

$$d(A, u) \leq (\lambda+1)\sigma^* + ((\lambda+1)\rho + 1)d_u^* \qquad (3)$$

Proof of Ineq. (3). Since u is an outer demand, it must be the case that $d_u^* \geq r(\ell_K)$. In addition, by Ineq. (1), $d(A, u) \leq \frac{\lambda+1}{\lambda} d(A, c_K^*) + d_u^*$. If the current phase is the last one ($\ell_K = -1$), then $d(A, c_K^*) \leq \lambda\sigma^*$, and the inequality follows. Otherwise, the current phase cannot be the first one (i.e., it must be $\ell_K < \psi$), because $d_u^* < r(\psi)$ and u could not be an outer demand. Therefore, $d(A, u) \leq R(\ell_K + 1) = \lambda\,\rho\, r(\ell_K) \leq \lambda\,\rho\, d_u^*$, and the inequality follows. □

Lemma 3 and Lemma 4 establish the main properties of DFL.

Lemma 3. *For any coalition K, $\mathsf{Pot}(\Lambda_K) = \sum_{u \in \Lambda_K} d(A, u) < f$.*

Proof. In the last phase ($\ell_K = -1$), $\mathsf{Pot}(\Lambda_K) = 0$, because there are no inner demands ($\mathsf{In}(K) = \emptyset$). If $\ell_K \geq 0$, for any inner demand u of K ($u \in \mathsf{In}(K)$),

$$d(A, u) \geq d(A, c_K^*) - d(c_K^*, u) > 3x \max\{D(K), r(\ell_K)\}\ ,$$

where the last inequality follows from (i) $d(A, c_K^*) \geq \lambda \max\{D(K), r(\ell_K)\}$, because K is a coalition, (ii) $d(u, c_K^*) < 2\max\{D(K), r(\ell_K)\}$, because of Ineq. (2), and (iii) $\lambda = 3x + 2$.

Let w be the demand in Λ_K which has arrived last, and let A_w be the facility configuration when w arrived. The last time $\mathsf{Pot}(\Lambda_K)$ increased was when w was added to L (and hence, to Λ_K). Since $D(\mathsf{In}(K)) < 3\max\{D(K), r(\ell_K)\} < \frac{d(A,w)}{x} \leq \frac{d(A_w,w)}{x}$, B_w must have contained the entire set Λ_K (including w). $\mathsf{Pot}(B_w)$ must have been less than f, because w was added to L. Therefore, $\mathsf{Pot}(\Lambda_K) \leq \mathsf{Pot}(B_w) < f$. □

Lemma 4. *Let w be any demand such that $\mathsf{Pot}(B_w) \geq f$, and, for a coalition K, let $\Lambda_K^w = B_w \cap \mathsf{In}(K)$. If there exists an active coalition K such that $\mathsf{Pot}(\Lambda_K^w) > \frac{\mathsf{Pot}(B_w)}{2}$, then $d(\hat{w}, c_K^*) < 8\max\{D(K), r(\ell_K)\}$.*

Proof. We first consider the case that $d(A,w) \geq f$ and $\hat{w}$ coincides with w. If there exists an active coalition K such that $w \in \mathsf{In}(K)$, the conclusion of the lemma follows from Ineq. (2). For any active coalition K' such that $w \notin \mathsf{In}(K')$, Lemma 3 implies that $\mathsf{Pot}(\Lambda^w_{K'}) < \frac{\mathsf{Pot}(B_w)}{2}$, because $\mathsf{Pot}(B_w \setminus \Lambda^w_{K'}) \geq d(A,w) \geq f$.

We have also to consider the case that $d(A,w) < f$. We observe that any subset of B_w including a potential greater than $\frac{\mathsf{Pot}(B_w)}{2}$ must have a non-empty intersection with Λ^w_K. If $\frac{r_w}{2^\nu} < 6\max\{D(K), r(\ell_K)\}$, let u be any demand in $\Lambda^w_K \cap B(\hat{w}, \frac{r_w}{2^\nu})$. Since u is an inner demand of K, using Ineq. (2), we show that

$$d(\hat{w}, c^*_K) \leq d(\hat{w}, u) + d(u, c^*_K) < 6\max\{D(K), r(\ell_K)\} + 2\max\{D(K), r(\ell_K)\} \ .$$

Otherwise, it must be $\frac{r_w}{2^{\nu+1}} \geq 3\max\{D(K), r(\ell_K)\} > D(\mathsf{In}(K))$. Therefore, for any $u \in \Lambda^w_K$, $B_w \cap B\big(u, \frac{r_w}{2^{\nu+1}}\big)$ includes the entire set Λ^w_K and hence, a potential greater than $\frac{\mathsf{Pot}(B_w)}{2}$. Consequently, there must be a single demand $u \in B_w$ such that $\mathsf{Pot}\big(B_w \cap B\big(u, \frac{r_w}{2^\nu}\big)\big) > \frac{\mathsf{Pot}(B_w)}{2}$. Since the previous inequality is satisfied by any demand $u \in \Lambda^w_K$, there must be only one demand in Λ^w_K, and $\hat{w}$ must coincide with it. The lemma follows from Ineq. (2), because $\hat{w}$ is an inner demand of K. □

Potential Function Argument. We use the potential function Φ to bound the total algorithm's cost. Let $T_{\mathcal{K}}$ be the hierarchical cover of C^* implied by Lemma 2.

$$\Phi = \sum_{K \in T_{\mathcal{K}}} \Phi_K \ , \text{ where } \Phi_K = \frac{(2x+1)(\lambda+1)}{x(\lambda-2)}\, \Xi_K - \frac{\lambda+1}{\lambda}\, \Upsilon_K \ .$$

The function Ξ_K is the sum of three components, $\Xi_K = \Xi^{(1)}_K + \Xi^{(2)}_K + \Xi^{(3)}_K$, where

$$\Xi^{(1)}_K = \sum_{j=0}^{\psi} \xi^{(1)}(K,j) \ , \quad \xi^{(1)}(K,j) = \begin{cases} f & \text{if } d(A, c^*_K) > R(j). \\ 0 & \text{if } d(A, c^*_K) \leq R(j). \end{cases}$$

$$\Xi^{(2)}_K = \begin{cases} 0 & \text{if } K \text{ is the root of } T_{\mathcal{K}}. \\ f \max\left\{ \ln\left(\dfrac{\min\{d(A,c^*_K), (\lambda+1)D(P_K)\}}{\lambda \frac{D(P_K)}{\gamma^2}} \right), 0 \right\} & \text{otherwise.} \end{cases}$$

$$\Xi^{(3)}_K = \begin{cases} 2f & \text{if } K \text{ is a typical coalition.} \\ f & \text{if } K \text{ is an isolated coalition.} \\ 0 & \text{if } K \text{ has become broken.} \end{cases}$$

The function Υ_K is defined as $\Upsilon_K = \begin{cases} \sum_{u \in \Lambda_K} d(A, c^*_K) & \text{if } K \in \mathsf{Act}(A). \\ 0 & \text{otherwise.} \end{cases}$

Let K be an active coalition. The function $\Xi^{(1)}_K$ compensates for the cost of opening the facility concluding the current phase of K. $\Xi^{(2)}_K$ compensates for the additional cost charged to the algorithm while K is typical active coalition (Lemma 7). $\Xi^{(3)}_K$ compensates for the cost of opening a facility which changes the status of K either from typical to isolated or from isolated to broken. The function Ξ_K never increases and can decrease only if a new facility closer to c^*_K is opened. The function Υ_K is equal to the potential accumulated by c^*_K. Υ_K increases when an inner demand of K is added to L and decreases when a new facility closer to c^*_K is opened.

In the following, $\Delta\Phi$ denotes the change in the potential function because of a demand w. More specifically, let Φ be the value of the potential function just before the arrival of w, and let Φ' be the value of the potential function just after the algorithm has finished processing w. Then, $\Delta\Phi = \Phi' - \Phi$. The same notation is used with any of the potential function components above.

We first prove that Φ_K remains non-negative (Lemma 5). If a demand w is added to L (i.e., no new facilities are opened), the algorithm incurs an assignment cost of $d(A,w)$, while if w

is not added to L (i.e., a new facility at $\hat{w}$ is opened), the algorithm incurs a facility cost of f and an assignment cost of $d(\hat{w}, w) < \frac{f}{x}$. In the former case, we show that $d(A, w) + \Delta\Phi \leq (\lambda+1)\sigma^* + ((\lambda+1)\rho+1)d^*_w$ (Lemma 6). In the latter case, we show that $f + d(\hat{w}, w) + \Delta\Phi \leq \frac{4(\lambda+1)}{\lambda-2}\left[(\lambda+1)\sigma^*|B_w| + ((\lambda+1)\rho+1)\sum_{u\in B_w} d^*_u\right]$ (Lemma 8).

Lemma 5. *For any coalition K, if $\ell_K \geq 0$, then $\Upsilon_K < \frac{\lambda}{\lambda-2}f$, while if $\ell_K = -1$, then $\Upsilon_K = 0$.*

Proof. In the last phase ($\ell_K = -1$), $\Upsilon_K = 0$ because there are no inner demands ($\mathsf{In}(K) = \emptyset$). Otherwise, DFL maintains the invariant that $\mathsf{Pot}(\Lambda_K) < f$ (Lemma 3). In addition, for any $u \in \Lambda_K$, $d(A, u) > \frac{\lambda-2}{\lambda} d(A, c^*_K)$, because of Ineq. (2). Therefore, $\Upsilon_K < \frac{\lambda}{\lambda-2}\mathsf{Pot}(\Lambda_K) < \frac{\lambda}{\lambda-2}f$. □

Lemma 5 implies that Φ_K is non-negative, because if K is an active coalition and $\ell_K \geq 0$, then $\frac{\lambda+1}{\lambda}\Upsilon_K < \frac{\lambda+1}{\lambda-2}f \leq \frac{(2x+1)(\lambda+1)}{x(\lambda-2)}\Xi^{(1)}_K$. On the other hand, if either K is not an active coalition or $\ell_K = -1$, then $\Upsilon_K = 0$.

Lemma 6. *If the demand w is added to L, then $d(A, w) + \Delta\Phi \leq (\lambda+1)\sigma^* + ((\lambda+1)\rho+1)d^*_w$.*

Proof. Let K be the unique active coalition such that $w \in C_K$. If w is an inner demand of K, w is added to Λ_K, and $\Delta\Phi = -\frac{\lambda+1}{\lambda}\Delta\Upsilon_K = -\frac{\lambda+1}{\lambda}d(A, c^*_K)$. Using Ineq. (1), we conclude that $d(A, w) + \Delta\Phi \leq d^*_w$. If w is an outer demand of K, then $\Delta\Phi = 0$. Using Ineq. (3), we conclude that $d(A, w) + \Delta\Phi \leq (\lambda+1)\sigma^* + ((\lambda+1)\rho+1)d^*_w$. □

We have also to consider demands w which are not added to L (i.e., a new facility at $\hat{w}$ is opened). Let A be the facility configuration just before the arrival of w, and let $A' = A \cup \{\hat{w}\}$. We observe that if either K is not an active coalition or $\ell_K = -1$, $\Upsilon_K = 0$ and Φ_K cannot increase due to the new facility at $\hat{w}$. Therefore, we focus on active coalitions K such that $\ell_K \geq 0$.

Lemma 7. *Let $\hat{w}$ be the facility opened when the demand w arrives. Then, for any typical active coalition K, the quantity $\frac{(2x+1)(\lambda+1)}{x(\lambda-2)}\Xi_K - \frac{(2x+1)(\lambda+1)}{x\lambda}\Upsilon_K$ cannot increase due to $\hat{w}$.*

Proof. If either the current phase ends ($d(\hat{w}, c^*_K) \leq R(\ell_K)$) or K stops being a typical active coalition due to $\hat{w}$, then $\Delta\Xi_K \leq -f$, and the lemma follows from $-\Delta\Upsilon_K < \frac{\lambda}{\lambda-2}f$.

If K remains a typical active coalition with respect to A' and the current phase does not end ($d(\hat{w}, c^*_K) > R(\ell_K)$), let $\tau^w_K = \frac{d(A, c^*_K)}{d(A', c^*_K)} \geq 1$ be factor by which $d(A, c^*_K)$ decreases because of the new facility at $\hat{w}$. K cannot be the root of $T_{\mathcal{K}}$, which is an isolated coalition by definition. Moreover, since K is a typical active coalition with respect to both A and A', Lemma 2 implies that $(\lambda+1)D(P_K) > d(A, c^*_K) \geq d(A', c^*_K) > \lambda\frac{D(P_K)}{\gamma^2}$. Therefore,

$$\Delta\Xi^{(2)}_K = \left[\ln\left(\frac{d(A', c^*_K)}{\lambda\frac{D(P_K)}{\gamma^2}}\right) - \ln\left(\frac{d(A, c^*_K)}{\lambda\frac{D(P_K)}{\gamma^2}}\right)\right] f = \ln\left(\frac{d(A', c^*_K)}{d(A, c^*_K)}\right) f = -\ln(\tau^w_K)\, f$$

If $B_w \cap \mathsf{In}(K) = \emptyset$, no demands are removed from Λ_K, and $-\Delta\Upsilon_K \leq (1 - \frac{1}{\tau^w_K})\,\Upsilon_K \leq \ln(\tau^w_K)\,\Upsilon_K$. Otherwise, we can show that $\tau^w_K > \frac{x}{3} > 3$, and $-\Delta\Upsilon_K \leq \Upsilon_K < \ln(\tau^w_K)\,\Upsilon_K$. In both cases, the lemma follows from $\Upsilon_K < \frac{\lambda}{\lambda-2}f$. □

Lemma 8. *Let $\hat{w}$ be the facility opened when the demand w arrives. Then,*

$$f + d(\hat{w}, w) + \Delta\Phi \leq \tfrac{4(\lambda+1)}{\lambda-2}\left[(\lambda+1)\sigma^*|B_w| + ((\lambda+1)\rho+1)\textstyle\sum_{u\in B_w} d^*_u\right] .$$

Proof Sketch. Let Λ_w be the set of inner demands in B_w, and let $M_w = B_w \setminus \Lambda_w$ be the set of outer demands in B_w. We recall that $f + d(\hat{w}, w) \leq \frac{x+1}{x}f$.

Case Isolated. There exists an isolated active coalition K such that $d(\hat{w}, c^*_K) < d(A, c^*_K)$. Lemma 7 implies that for any typical active coalition K', $\Delta\Phi_{K'} \leq 0$. In addition, for $x \geq 10$, we can prove that (i) for any isolated active coalition K' different from K, $d(\hat{w}, c^*_{K'}) \geq d(A, c^*_{K'})$, and (ii) for any active coalition K' different from K, $B_w \cap \mathsf{In}(K') = \emptyset$. As a result, for any isolated active coalition K' different from K, $\Delta\Phi_{K'} = 0$. In addition, only inner demands of K are included in B_w ($\Lambda_w \subseteq \mathsf{In}(K)$).

We have also to bound $\frac{x+1}{x} f + \Delta\Phi_K$. Since $-\Delta\Upsilon_K < \frac{\lambda}{\lambda-2} f$ and $\lambda = 3x+2$, $\frac{x+1}{x} f + \Delta\Phi_K < \frac{2(\lambda+1)}{\lambda-2} f + \Delta\Xi_K$. We distinguish between two cases depending on the potential contributed by Λ_w.

Case Isolated.A. $\mathsf{Pot}(\Lambda_w) \leq \frac{\mathsf{Pot}(B_w)}{2}$. Then, $\frac{2(\lambda+1)}{\lambda-2} f$ cannot exceed $\frac{4(\lambda+1)}{\lambda-2}\,\mathsf{Pot}(M_w)$. We also recall than $\Delta\Xi_K \leq 0$. Hence, both the algorithm's cost and the increase in the potential function can be charged to the potential of the outer demands in B_w. Using Ineq. (3), we conclude that

$$\tfrac{x+1}{x} f + \Delta\Phi_K < \tfrac{4(\lambda+1)}{\lambda-2}\mathsf{Pot}(M_w) \leq \tfrac{4(\lambda+1)}{\lambda-2}\left[(\lambda+1)\sigma^*|B_w| + ((\lambda+1)\rho+1)\textstyle\sum_{u\in B_w} d^*_u\right] \ .$$

Case Isolated.B. $\mathsf{Pot}(\Lambda_w) > \frac{\mathsf{Pot}(B_w)}{2}$. Since $\Lambda_w \subseteq \mathsf{In}(K)$, Lemma 4 implies that $d(\hat{w}, c^*_K) < 8\max\{D(K), r(\ell_K)\}$. Hence, either the current phase ends or the coalition K becomes broken. In both cases, $\Delta\Xi_K \leq -f$ and the decrease in Ξ_K compensates for both the algorithm's cost and the decrease in Υ_K.

Case Typical. For any isolated active coalition K, $d(\hat{w}, c^*_K) \geq d(A, c^*_K)$. Therefore, no inner demands of K are included in B_w, because it would be $d(\hat{w}, c^*_K) < \frac{x}{3} d(A, c^*_K)$ otherwise. As a result, $\Delta\Phi_K = \Delta\Upsilon_K = 0$.

If w is an inner demand, let K_w be the unique typical active coalition such that $w \in \mathsf{In}(K_w)$. Similarly to the proof of Lemma 7, we can show that $\frac{x+1}{x} f + \Delta\Phi_{K_w} \leq 0$. In addition, for any typical active coalition K' different from K_w, Lemma 7 implies that $\Delta\Phi_{K'} \leq 0$.

If w is an outer demand, using the following upper bound on $\mathsf{Pot}(B_w)$, we can charge the algorithm's cost to the potential of B_w.

$$\tfrac{x+1}{x} f \leq \mathsf{Pot}(B_w) \leq \tfrac{x+1}{x}\left[(\lambda+1)\sigma^*|B_w| + ((\lambda+1)\rho+1)\sum_{u\in B_w} d^*_u\right] - \tfrac{x+1}{x}\tfrac{\lambda+1}{\lambda}\sum_{K\in\mathsf{Act}(K)} \Delta\Upsilon_K$$

We conclude the proof by applying Lemma 7 for each typical active coalition. □

In addition to the initial credit provided by the potential function Φ, a demand's optimal assignment cost is considered at most once by Lemma 6 (i.e., when the demand is added to L) and at most once by Lemma 8 (i.e., when the demand is removed from L). Therefore, the algorithm's total cost cannot exceed $\frac{2(2x+1)(\lambda+1)}{x(\lambda-2)}\left[\psi + 3 + \ln\left(\frac{\lambda+1}{\lambda}\gamma^2\right)\right]\mathsf{F}^* + \frac{5\lambda+2}{\lambda-2}\left[(\lambda+1)\rho + \lambda + 2\right]\mathsf{S}^*$. Setting $\gamma = 8\lambda$ and $\psi = \rho = \mathrm{O}(\frac{\log n}{\log\log n})$ yields the claimed competitive ratio. □

4 The Algorithm for Non-Uniform Facility Costs

In this section, we outline the algorithm Non-Uniform Deterministic Facility Location – NDFL, which is a generalization of DFL and can handle non-uniform facility costs.

The algorithm first rounds down the facility costs to the nearest integral power of two. For each demand w, the algorithm computes r_w, B_w, $\mathsf{Pot}(B_w)$, and $\hat{w}$ as in Fig. 1. If $|B_w| > 1$, NDFL opens the cheapest facility in $B(w, r_w) \cup B(\hat{w}, r_w)$ if its cost does not exceed $\mathsf{Pot}(B_w)$. Ties are always broken in favour of $\hat{w}$. Namely, if there are many facilities of the same (cheapest) cost, the one nearest to $\hat{w}$ is opened. If a new facility is opened, the demands of B_w are removed from L. Otherwise, w is added to L. If $|B_w| = 1$, NDFL keeps opening the cheapest facility in $B(w, r_w)$ while there is a facility of cost no greater than $\mathsf{Pot}(B_w)$. In this case, $\hat{w}$ coincides with w and ties are broken in favour of w. After opening a new facility, the algorithm updates r_w and $\mathsf{Pot}(B_w)$ according to the new facility configuration and iterates. After the last iteration, w is added to L. As in Fig. 1, the algorithm finally assigns w to the nearest facility.

The following theorem can be proven by generalizing the techniques described in Section 3.

Theorem 3. *For any constant* $x \geq 12$*, the competitive ratio of* NDFL *is* $\mathrm{O}(\frac{\log n}{\log \log n})$.

5 An Open Problem

In the framework of incremental clustering (e.g., [6, 7]), an algorithm is also allowed to merge some of the existing clusters. On the other hand, the lower bound of Theorem 1 on the competitive ratio for Online Facility Location crucially depends on the restriction that facilities cannot be closed. A natural open question is how much the competitive ratio can be improved if the algorithm is also allowed to close a facility by re-assigning the demands to another facility (i.e., merge some of the existing clusters). This research direction is related to an open problem of [7] concerning the existence of an incremental algorithm for k-Median which achieves a constant performance ratio using $\mathrm{O}(k)$ medians.

References

1. S. Albers and H. Koga. New online algorithms for the page replication problem. *J. of Algorithms*, 27(1):75–96, 1998.
2. B. Awerbuch, Y. Bartal, and A. Fiat. Competitive distributed file allocation. *Proc. of STOC '93*, pp. 164–173, 1993.
3. Y. Bartal. Probabilistic approximations of metric spaces and its algorithmic applications. *Proc. of FOCS '96*, pp. 184–193, 1996.
4. Y. Bartal, A. Fiat, and Y. Rabani. Competitive algorithms for distributed data management. *J. of Computer and System Sciences*, 51(3):341–358, 1995.
5. A. Borodin and R. El-Yaniv. *Online Computation and Competitive Analysis.* Cambridge University Press, 1998.
6. M. Charicar, C. Chekuri, T. Feder, and R. Motwani. Incremental clustering and dynamic information retrieval. *Proc. of STOC '97*, pages 626–635, 1997.
7. M. Charicar and R. Panigrahy. Clustering to minimize the sum of cluster diameters. *Proc. of STOC '01*, pages 1–10, 2001.
8. R. Fleischer and S. Seiden. New results for online page replication. *Proc. of APPROX '00, LNCS 1913*, pp. 144–154, 2000.
9. S. Guha. *Approximation Algorithms for Facility Location Problems.* PhD Thesis, Stanford University, 2000.
10. S. Guha and S. Khuller. Greedy strikes back: Improved facility location algorithms. *Proc. of SODA '98*, pp. 649–657, 1998.
11. M. Imase and B.M. Waxman. Dynamic Steiner tree problem. *SIAM J. on Discrete Mathematics*, 4(3):369–384, 1991.
12. K. Jain and V. Vazirani. Approximation algorithms for metric facility location and k-median problems using the primal-dual schema and Lagrangian relaxation. *J. of the ACM*, 48(2):274–296, 2001.
13. C. Lund, N. Reingold, J. Westbrook, and D.C.K. Yan. Competitive online algorithms for distributed data management. *SIAM J. on Computing*, 28(3):1086–1111, 1999.
14. M. Mahdian, Y. Ye, and J. Zhang. Improved approximation algorithms for metric facility location problems. *Proc. of APPROX '02, LNCS 2462*, pp. 229–242, 2002.
15. R.R. Mettu and C.G. Plaxton. The online median problem. *Proc. of FOCS '00*, pp. 339–348, 2000.
16. A. Meyerson. Online facility location. *Proc. of FOCS '01*, pp. 426–431, 2001.
17. D. Shmoys. Approximation algorithms for facility location problems. *Proc. of APPROX '00, LNCS 1913*, pp. 27–33, 2000.
18. D. Shmoys, E. Tardos, and K. Aardal. Approximation algorithms for facility location problems. *Proc. of STOC '97*, pp. 265–274, 1997.