On the Competitive Ratio for Online Facility Location*

Dimitris Fotakis

Department of Mathematical, Physical and Computational Sciences,
Aristotle University of Thessaloniki, 54006 Thessaloniki, Greece.
Email: fotakis@gen.auth.gr

Abstract. We consider the problem of Online Facility Location, where the demand points 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 which achieves a competitive ratio of $\mathrm{O}(\frac{\log n}{\log\log n})$ in every metric space.

Keywords: Online Algorithms, Competitive Analysis, Metric Facility Location.

1 Introduction

In the (metric uncapacitated) Facility Location problem, we are given a metric space along with a facility cost for each point and a (multi)set of demand points, and we seek for a set of facility locations which minimize the sum of facility and assignment costs. Facility Location provides a simple and natural model for network design and data clustering problems and has been the subject of intensive research over the last decade (e.g., see [26] for a survey and [15] for approximation algorithms and applications).

In addition to the offline setting, there are many practical applications where the demand points either are not known in advance or the solution must be constructed incrementally using limited (if any) information about future demands (e.g., see [24] for examples from the areas of network design and data clustering). For instance, telecommunication networks are usually upgraded incrementally to accommodate new clients because it is expensive and often infeasible to reconfigure the complete network every time a few new clients are added (e.g., [25]).

The definition of *Online Facility Location* [24] is motivated by similar considerations. In Online Facility Location, the demand points arrive one at a time and must be irrevocably assigned to either an existing or a new facility upon arrival. The objective is to minimize the sum of facility and assignment costs, where the assignment cost of a demand is its distance from the facility it is assigned to.

Related Work. In the offline case, where the demand points are fully known in advance, there are constant factor approximation algorithms based on Linear Programming rounding (e.g., [27, 28]), local search (e.g., [20, 9, 3]), and the primal-dual method (e.g., [19, 18]). The best known polynomial-time algorithm achieves an approximation ratio of 1.52 [22], while no polynomial-time algorithm can achieve an approximation ratio less than 1.463 unless $\mathrm{NP} = \mathrm{DTIME}(n^{\mathrm{O}(\log\log n)})$ [16].

* A preliminary version of this work appeared in the Proceedings of the 30th International Colloquium on Automata, Languages and Programming (ICALP 2003), Lecture Notes in Computer Science 2719. This work was done while the author was at the Max-Planck-Institut für Informatik, Saarbrücken, Germany, and was partially supported by the Future and Emerging Technologies programme of the EU under contract number IST-1999-14186 (ALCOM–FT).

Meyerson [24] introduced the online version of Facility Location and presented a simple randomized algorithm which achieves a constant performance ratio if the demands are examined in random order. In the standard framework of competitive analysis (e.g., [7]), where both the set of demands and their arrival order are selected by an oblivious adversary, the same algorithm achieves a competitive ratio of $\mathrm{O}(\frac{\log n}{\log\log n})$[1]. Meyerson also presented a lower bound of $\Omega(\log^* n)$ on the competitive ratio of any online algorithm, where $\log^*$ denotes the inverse Ackerman function.

After a preliminary version of this work had appeared [13], Anagnostopoulos et al. [2] presented a $\Theta(2^d \log n)$-competitive deterministic algorithm for Online Facility Location in d-dimensional Euclidean spaces. Their algorithm works by partitioning a d-dimensional hypercube with sides of length μ into 2^d hypercubes with sides of length $\mu/2$ as soon as the assignment cost for the demands in the hypercube has exceeded the cost of opening a new facility. For Euclidean spaces of small dimension, this approach results in a deterministic algorithm which is simpler and faster than the algorithm presented in this paper. Nevertheless, the approach of [2] cannot guarantee an optimal (sub-logarithmic) competitive ratio (even on the line) and cannot be applied to general metric spaces. In addition to the theoretical analysis, an experimental comparison of several online algorithms for Facility Location in the Euclidean plane is presented in [2].

Online Facility Location bears a resemblance to the problem of Online File Replication which has been extensively and productively studied (e.g., [6, 4, 1, 21, 12]). 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 points receiving the file with some points already holding the file. Similarly to Facility Location, File Replication asks for a set of locations which minimize the sum of replication/facility and access/assignment 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 distance between the points already holding the file and the points which receive the file.

Online File Replication is a generalization of Online Steiner Tree [17]. 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 [6] and a deterministic [4] algorithm achieving a competitive ratio of $\mathrm{O}(\log n)$ for the more general problem of Online File Allocation. For trees and rings, algorithms of constant competitive ratio are known [1, 21, 12].

Very recently, the incremental version of Facility Location has also been studied [14] (see also [8] for an introduction to the framework of incremental algorithms for data clustering problems). Similarly to the online version, the demand points arrive one at a time and must be assigned to either an existing or a new facility upon arrival. In the incremental version, the algorithm can also merge some of the existing facilities at any point in time. Namely, the algorithm can close a facility by re-assigning all the demands assigned to it to another open facility. In [14], the first incremental algorithm for Facility Location which achieves a constant competitive ratio is presented[2].

[1] A competitive ratio of $\mathrm{O}(\log n)$ is stated in [24]. Nevertheless, a competitive ratio of $\mathrm{O}(\frac{\log n}{\log\log n})$ follows from a simple modification of the same argument.

[2] We should remark that the lower bound of $\mathrm{O}(\frac{\log n}{\log\log n})$ only applies if the decision of opening a facility at a particular location is irrevocable. Hence, the lower bound of this paper does not apply to the incremental version of Facility Location.

We also remark that Online Facility Location should not be confused with the problem of Online Median [23]. In Online Median, the demand points are fully known in advance and the number of facilities increases online. A $\mathrm{O}(1)$-competitive algorithm for Online Median is presented in [23].

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 $\Omega(\log^* n)$ [24].

On the positive side, we present a deterministic algorithm which achieves a competitive ratio of $\mathrm{O}(\frac{\log n}{\log\log n})$ in every metric space. We formulate and analyze the algorithm for both uniform and non-uniform facility costs. To the best of our knowledge, this is the only known deterministic algorithm which achieves an (asymptotically) optimal competitive ratio. In addition, this is the only known deterministic algorithm which works in every metric space and has been formulated and analyzed for the general case of non-uniform facility costs.

For the analysis, we introduce the notion of *isolated active coalitions* of optimal facilities and establish that the algorithm's competitive ratio is essentially determined by the special case that the optimal solution consists of a single facility. The basic idea is that each isolated active coalition can be analyzed similarly to the special case of a single optimal facility. Furthermore, we show that no coalition can remain both *non-isolated* and *active* for too long. To establish this claim, we show that every metric space admits a hierarchical decomposition such that each component is either well-separated or has a large diameter. To the best of our knowledge, this fundamental property of metric spaces has not been noticed and proven before.

Our approach gives new insight into the behaviour of a broad and natural class of algorithms for Online Facility Location and related problems. A non-trivial generalization of this approach has already been applied to the incremental version of Facility Location [14]. We believe that our approach is of independent interest and will find applications to other online problems as well.

Problem Definition and Notation. We evaluate the performance of online algorithms using *competitive analysis* (e.g., [7]). 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.

In Online Facility Location, we are given a metric space $\mathcal{M} = (M, d)$, where M denotes the set of points and $d : M \times M \mapsto \mathbb{R}_+$ denotes the distance function which is non-negative, symmetric and satisfies the triangle inequality. For each point $v \in M$, we are also given the cost f_v of opening a facility at v. We abuse the notation by letting the same symbol denote both a demand/facility and the corresponding point of the metric space.

The demand sequence consists of (not necessarily distinct) points $w \in M$. When a demand w arrives, the algorithm can open some new facilities, in which case it incurs the corresponding cost. The decision of opening a facility at a particular location is irrevocable. Then, w must be irrevocably assigned to 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.

By the definition above, it should be clear that we only consider unit demands and allow multiple demands to be located at the same point. We 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.

A metric space $\mathcal{M} = (M, d)$ is usually identified by its point set M. For a subspace $M' \subseteq M$, $D(M') = \max_{u,v\in M'}\{d(u, v)\}$ denotes the diameter of M'. For a point $u \in M$ and a subspace $M' \subseteq M$, $d(M', u) = \min_{v\in M'}\{d(v, u)\}$ denotes the distance between u and M'. It is $d(\emptyset, u) = \infty$. For subspaces $M', M'' \subseteq M$, $d(M', M'') = \min_{u\in M''}\{d(M', u)\}$ denotes the minimum distance between M' and M''. For a subspace $M' \subseteq M$, $\mathrm{sep}_M(M') = d(M', M \setminus M')$ denotes the distance separating the points in M' from the remaining points in M. It is $\mathrm{sep}_M(\emptyset) = \mathrm{sep}_M(M) = \infty$.

For a point $u \in M$ and a non-negative number r, we let $\mathrm{Ball}(u, r)$ denote the set of points within distance r from u, $\mathrm{Ball}(u, r) = \{v \in M : d(u, v) \leq r\}$. We also let $\ln\alpha$ denote the natural logarithm and $\log\alpha$ the logarithm to base 2 of a positive number α.

Organization. In Section 2, we establish a lower bound of $\Omega(\frac{\log n}{\log\log n})$ on the competitive ratio of any algorithm for Online Facility Location.

Section 3 is devoted to the presentation and the analysis of the deterministic algorithm for uniform facility costs. In Section 3.1, we give a sketch of the analysis for the special case that the optimal solution consists of a single facility and introduce the basic notions required to apply the same ideas to the general case. The properties of the algorithm along with several useful propositions are established in sections 3.2 and 3.3. In Section 3.4, we develop a potential function argument and show that the algorithm's competitive ratio is $\mathrm{O}(\frac{\log n}{\log\log n})$.

In Section 4, we generalize the deterministic algorithm to deal with non-uniform facility costs. In Section 4.1, we introduce the notation used in the analysis of the algorithm, and in sections 4.2 and 4.3, we establish its basic properties. In Section 4.4, we generalize the potential function argument to deal with non-uniform facility costs and prove that algorithm's competitive ratio remains $\mathrm{O}(\frac{\log n}{\log\log n})$.

2 A Lower Bound on the Competitive Ratio

The lower bound of this section holds for uniform facility costs and instances whose optimal solution consists of a single facility.

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. We first establish the lower bound in a complete binary Hierarchically Well-Separated Tree [5]. 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 $\frac{D}{m^i}$ (Fig. 1). Also, let f be the cost of opening a new facility, which is the same for every vertex of T (uniform facility costs).

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 $\frac{D}{m^{i-1}}$.

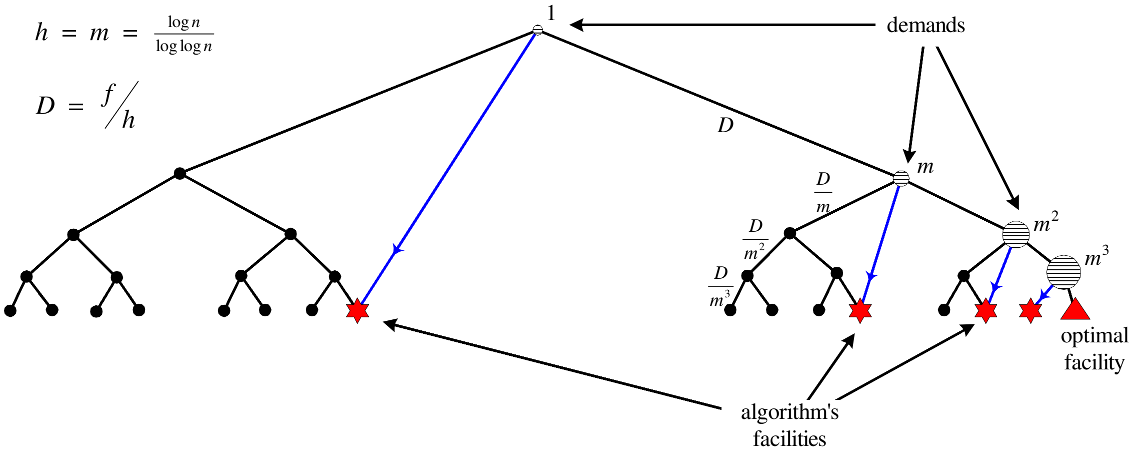

Fig. 1. The metric space and the demand sequence used in the proof of Theorem 1.

By Yao's principle (e.g., [7, Chapter 8]), it suffices to show that there is a probability distribution over the demand sequences for which the ratio of the expected cost of any deterministic online algorithm to the expected optimal cost is $\Omega(\frac{\log n}{\log\log n})$.

To define an appropriate probability distribution, we divide the demand sequence into $h+1$ phases. Phase 0 consists of a single demand located at the root v_0. After the end of phase i, $0 \le i \le h$, if v_i is not a leaf, the adversary selects v_{i+1} uniformly at random and independently among the two children of v_i. Phase $i+1$ consists of m^{i+1} unit demands located 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 except for the last one, 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, and let i, $0 \le i \le h-1$, be any phase except for the last one. We fix the adversary's random choices $v_0, \dots, v_i$ up to the end of phase i (hence, we also fix T_{v_i}) and consider the (conditional) expected cost incurred by Alg for demands and facilities not in $T_{v_{i+1}}$.

If Alg has no facilities in T_{v_i} at the moment the first demand at v_{i+1} arrives, the assignment cost for the demands located at $v_i \in T_{v_i} \setminus T_{v_{i+1}}$ is at least $m^i \frac{D}{m^{i-1}} = mD$. Otherwise, since v_{i+1} is selected uniformly at random and independently 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 every fixed choice of $v_0, \dots, v_i$ (and every 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 every fixed choice of $v_0, \dots, v_i$, it also holds regardless of the conditioning. Therefore, 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, $0 \le i \le h$, 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 an additional cost no less than $\min\{mD, f\}$ for demands and facilities in 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 $\frac{2h-1}{h-1}hD$. We have also the constraint that $\frac{h^{h+1}}{h-1}$ 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 observe that if m is large enough, T is not really different from a line segment (as far as the lower bound construction is concerned). More specifically, we consider the following embedding of T in a line segment L. The root is mapped to the center of the segment, namely, point 0. 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 $\tilde{v} + \frac{D}{m^i}$. In the following, we use $d(u, v)$ to denote the distance between the vertices u and v in T and $\tilde{d}(u, v)$ to denote the distance between the images of u and v in L.

We show that for every $m > 2$, this embedding preserves the main property of T used in the proof of the lower bound. In particular, we prove that the distance (in the line segment L) between the image of any vertex v of height i and the image of any vertex in T_v is at most $\frac{m}{m-1}\frac{D}{m^i}$, while the distance (in L) between the image of v and the image of any vertex not in T_v is at least $\frac{D}{m^{i-1}}$.

We first observe that this embedding does not increase the distance between any pair of vertices. Namely, for every $u, v \in T$, $\tilde{d}(u, v) \leq d(u, v)$. Hence, for every vertex v of height i and every vertex $u \in T_v$, $\tilde{d}(u, v) \leq \frac{m}{m-1}\frac{D}{m^i}$.

In addition, if u and v are either adjacent vertices (u is a child of v or vice versa) or siblings (u and v have the same immediate ancestor/parent) in T, the embedding preserves the distance between them and $\tilde{d}(u, v) = d(u, v)$.

Next, we show that for every vertex w, the image of any vertex in the left (right) subtree of w must lie on the left (resp. right) side of the image of w in L. Formally, let w be any vertex of height i, and let u' and v' be the two children of w. By the definition of the embedding, u' and v' do not lie on the same side of w. More specifically, $\tilde{d}(w, v') = \tilde{d}(w, u') = \frac{D}{m^i}$ and $\tilde{d}(u', v') = \frac{2D}{m^i}$. Let now u be any successor of u', and let v be any successor of v' in T. Since u' and v' are vertices of height $i+1$ in T and $u \in T_{u'}/v \in T_{v'}$, we obtain that $\tilde{d}(v, v') \leq \frac{D}{m^{i+1}}\frac{m}{m-1} < \tilde{d}(w, v')$ and $\tilde{d}(u, u') \leq \frac{D}{m^{i+1}}\frac{m}{m-1} < \tilde{d}(w, u')$. In other words, the image of u must lie on the same side of w with the image of u' and the image of v must lie on the same side of w with the image of v'. Therefore, the images of v and u cannot lie on the same side of the image of w in L.

Finally, we use induction on i and show that the distance (in L) between the image of a vertex v of height i and the image of any vertex not in T_v is at least $\frac{D}{m^{i-1}}$. Let v be any vertex of height i, $i \geq 1$, let u be any vertex not in T_v, and let w be the parent of v in T. If $u \in T_w \setminus T_v$, then either $u = w$ or u and v lie on different sides of w. In both cases, $\tilde{d}(v, u) \geq \tilde{d}(v, w) = \frac{D}{m^{i-1}}$. This argument also establishes the base case (namely, $i = 1$) because if v is a child of the root, every vertex not in T_v must belong to $T_w \setminus T_v$. If v is not a child of the root and $u \notin T_w$, then

$$\tilde{d}(v, u) \geq \tilde{d}(u, w) - \tilde{d}(v, w) \geq \frac{D}{m^{i-2}} - \frac{D}{m^{i-1}} \geq \frac{D}{m^{i-1}}$$

The second inequality follows from the inductive hypothesis, since w is a vertex of height $i-1$ and $u \notin T_w$, and $\tilde{d}(v, w) = \frac{D}{m^{i-1}}$, since v is a child of w. The last inequality holds for every $m \geq 2$. □

3 A Deterministic Algorithm for Uniform Facility Costs

The algorithm Deterministic Facility Location – DFL (Fig. 2) maintains its *facility configuration* F and the set L of *unsatisfied demands*. A demand becomes unsatisfied and is added to L upon arrival. Each unsatisfied demand holds a *potential* which is always equal to its distance from the nearest facility. If the unsatisfied neighborhood B_w of a new demand w has accumulated a potential no less than f, the algorithm opens a new facility at an appropriately selected location $\hat{w} \in B_w$. The location of the new

Let x be an appropriately chosen constant.
$F \leftarrow \emptyset;\ L \leftarrow \emptyset;$
For each new demand w:
  $L \leftarrow L \cup \{w\};\ \ r_w \leftarrow d(F,w)/x;$
  $B_w \leftarrow \mathrm{Ball}(w, r_w) \cap L;\ \mathrm{Pot}(B_w) \leftarrow \sum_{u \in B_w} d(F,u);$
  if $\mathrm{Pot}(B_w) \geq f$ then `/* A new facility opens */`
    if $d(F,w) < f$ then
      Let $\nu \geq 0$ be the smallest integer:
        either there exists exactly one point $u \in B_w$ such that
          $\mathrm{Pot}(B_w \cap \mathrm{Ball}(u, r_w/2^\nu)) > \mathrm{Pot}(B_w)/2$,
        or for every $u \in B_w$, $\mathrm{Pot}(B_w \cap \mathrm{Ball}(u, r_w/2^{\nu+1})) \leq \mathrm{Pot}(B_w)/2$.
      Let $\hat{w}$ be any point in B_w: $\mathrm{Pot}(B_w \cap \mathrm{Ball}(\hat{w}, r_w/2^\nu)) > \mathrm{Pot}(B_w)/2$.
    else $\hat{w} \leftarrow w$;
    $F \leftarrow F \cup \{\hat{w}\}$; $L \leftarrow L \setminus B_w$;
  Assign w to the nearest facility in F.

Fig. 2. The algorithm Deterministic Facility Location – DFL.

facility $\hat{w}$ is the center of the smallest radius ball/subset of B_w contributing more than half of B_w's potential[3]. Then, the unsatisfied demands in B_w lose their potential, become satisfied and are removed from L. Finally, the new demand is assigned to the nearest facility currently in F.

The remaining of this section is devoted to the proof of the following theorem.

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

3.1 Notation and Techniques

The notion of unsatisfied demands (the set L) ensures that each demand contributes to the facility cost at most once. More specifically, each unsatisfied demand $u \in L$ can contribute an amount of $d(F,u)$ to the cost of opening a new facility in its neighborhood. Then, it becomes satisfied and cannot contribute to the algorithm's cost anymore. For each unsatisfied demand u, we usually refer to $d(F,u)$ as the *potential of* u. Only unsatisfied demands have non-zero potential, which is non-increasing with time. For a set S of unsatisfied demands $\mathrm{Pot}(S) = \sum_{u \in S} d(F,u)$ denotes the total potential of S.

For an arbitrary fixed sequence of n demands, we compare the algorithm's cost with the cost of the offline optimal solution. We denote the set of optimal facilities by F^*. To avoid confusing the algorithm's facilities in F with the optimal facilities in F^*, we use the term *optimal center*, or simply *center*, to refer to an optimal facility in F^* and the term *facility* to refer to an algorithm's facility in F.

The optimal solution F^* consists of k centers $c_1, c_2, \ldots, c_k$. In the optimal solution, each demand is assigned to the nearest center in F^*. Hence, F^* defines an optimal partitioning of the demand sequence into clusters $C_1, C_2, \ldots, C_k$. For each demand u, c_u denotes the optimal center and C_u denotes the optimal cluster u is assigned to. Let $d^*_u = d(c_u, u) = d(F^*, u)$ denote the assignment cost of u in the

[3] To achieve a sub-logarithmic competitive ratio, the algorithm must carefully select the location of the new facility. In the Appendix, Section A.1, we show that the variant of DFL which opens a new facility at the location of each new demand w with $\mathrm{Pot}(B_w) \geq f$ has a competitive ratio of $\Omega(\log n)$.

optimal solution, let $\mathrm{Asg}^* = \sum_u d_u^*$ be the total optimal assignment cost, let $\mathrm{Fac}^* = kf$ be the total optimal facility cost, and let $\delta^* = \mathrm{Asg}^*/n$ be the average optimal assignment cost per demand. For the analysis, we use δ^* as the basic distance unit.

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

We break down the analysis of the cost charged to each optimal center c_i into $\psi + 2$ disjoint *phases* according to the distance between the algorithm's facility configuration F and c_i. For each optimal center c_i, phase j starts just after the first facility within a distance of $R(j+1)$ from c_i opens and ends when the first facility within a distance of $R(j)$ from c_i opens. In other words, the j-th phase for the optimal center c_i lasts as long as $d(F, c_i) \in (R(j), R(j+1)]$, $-1 \leq j \leq \psi + 1$. We use the convention that the last phase (i.e., $j = -1$) never ends.

In the following, a small neighborhood/ball should be interpreted as a ball of radius approximately λ times smaller than the distance of its demands/points from the nearest algorithm's facility. Hence, the unsatisfied neighborhood B_w of each new demand w includes every other unsatisfied demand located in a small ball around w (cf. Fig. 5). Intuitively, the demands of a small ball could be regarded as being located at the center of the ball.

An Overview of the Analysis. We start by informally stating the main properties of DFL.

The main invariant maintained by the algorithm is that the potential accumulated by any small neighborhood does not exceed the cost of opening a facility in this neighborhood (cf. Fig. 5 and Lemma 3). Otherwise, a new facility would have opened and the neighborhood's potential would have dropped to 0 when the last demand in the neighborhood arrived.

The second basic property has to do with the location of the new facility. Each new facility $\hat{w}$ is located in the smallest radius ball/subset of B_w which contributes more than half of B_w's potential. Therefore, if there is a single small neighborhood which contributes more than half of B_w's potential, the new facility $\hat{w}$ will be located very close to it (cf. Lemma 4).

The Case of a Single Optimal Center. To provide some intuition about the reason that the properties above imply a competitive ratio of $\mathrm{O}(\frac{\log n}{\log \log n})$, we sketch the analysis of the algorithm in the special case that the optimal solution consists of a single center c. In the following, we let ℓ denote the current phase for the optimal center c.

The demands which either arrive in the current phase or remain unsatisfied from the previous phases are partitioned into *inner* and *outer* demands (see also Fig. 3 for the definition of inner and outer demands).

Outer demands are at distance no less than $r(\ell)$ from c. Therefore, both the assignment cost (of each new outer demand) and the potential (of each unsatisfied outer demand) are within a constant factor from the demand's optimal assignment cost. More specifically, if u is an outer demand in the current phase, its distance from the nearest facility is $d(F, u) \leq \lambda\delta^* + (\lambda\rho + 1)d_u^*$ (cf. Ineq. (3)).

Inner demands are included in a ball of radius $r(\ell)$ around the optimal center c. We have chosen λ so as $R(\ell)$ to be sufficiently larger than $r(\ell)$ and each new inner demand w to include every unsatisfied inner demand in its neighborhood B_w (Fig. 5). Hence, the potential accumulated by the set of unsatisfied inner demands can never exceed f (Lemma 3). Intuitively, inner demands can also be thought of as being

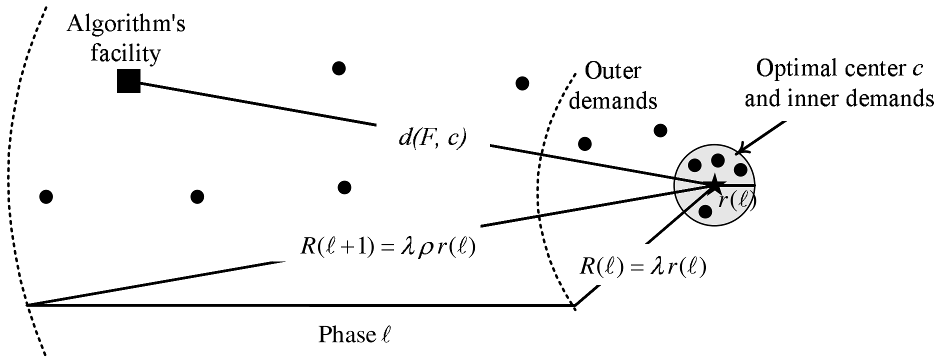

Fig. 3. The current phase, the set of inner demands, and the set of outer demands.

essentially located at c because they are much closer to c (and to each other) than to any facility in F (cf. Ineq (2)).

We have already noticed that the assignment cost of new outer demands can be directly charged to their optimal assignment cost. We use a typical potential function argument to bound the algorithm's facility cost and the total assignment cost of the inner demands which remain unsatisfied at their assignment time (in other words, the inner demands which do not cause any new facilities to open when they arrive).

For simplicity, we refer to the potential of unsatisfied inner demands as the potential of c (recall that the inner demands can be regarded as being essentially located at c). We also assume that the set of unsatisfied inner demands becomes empty at the beginning of each new phase. Hence, at the beginning of the current phase, the potential of c is 0. Each time an inner demand arrives and remains unsatisfied at its assignment time, the potential of c increases by (approximately) the demand's assignment cost. Since the potential of c drops to 0 at the beginning of the next phase, the total assignment cost of the inner demands which arrive in the current phase and remain unsatisfied at their assignment time does not exceed the total decrease in the potential of c in the current phase (see also Fig. 4).

The potential of c can decrease only if a new facility opens. This can happen because either the new facility is even closer to c or some unsatisfied inner demands contribute their potential to the cost of the new facility (or both). We have already mentioned that the potential of c is always less than f (cf. Fig. 5 and Lemma 3). Hence, we charge the algorithm with a cost of f for the decrease in the potential of c due to a new facility.

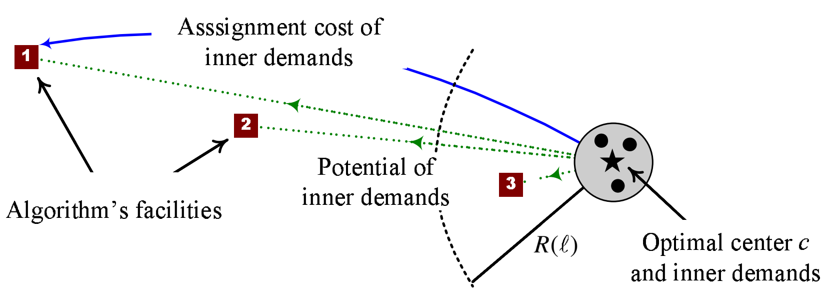

Fig. 4. The potential of an inner demand is initially equal to its assignment cost and drops to 0 by the end of the current phase.

In addition, the algorithm incurs a cost of f for each new facility and an assignment cost no greater than $\frac{f}{x}$ for the demand opening the new facility (this demand has become satisfied by its assignment time). Therefore, there is a total cost of $\frac{2x+1}{x} f$ associated with each new facility.

If the inner demands in B_w contribute more than half of B_w's potential, the new facility is located within a distance of $R(\ell)$ from c. More specifically, the new facility is located in the smallest radius subset of B_w which contributes more than half of B_w's potential and inner demands are included in the small neighborhood $\mathrm{Ball}(c, r(\ell))$. Hence, the new facility must be located in a ball of radius at most $4\,r(\ell)$ which intersects the set of inner demands (cf. Lemma 4). Therefore, the current phase ends and the cost associated with the new facility is charged to the optimal facility cost. Otherwise, the potential of the outer demands in B_w must be at least $f/2$. Thus, the cost associated with the new facility is charged to the potential of the outer demands in B_w, which is in turn charged to their optimal assignment cost (cf. Ineq (3)). Furthermore, these demands become satisfied and cannot be charged with any additional cost in the future.

The analysis implies that the total algorithm's cost is $\mathrm{O}(\psi)$ times the optimal facility cost plus $\mathrm{O}(\lambda\rho)$ times the optimal assignment cost. Setting $\psi = \rho = \mathrm{O}(\frac{\log n}{\log \log n})$ yields the desired competitive ratio.

The General Case. If there are $k > 1$ optimal centers $c_1, \ldots, c_k$, the demands are partitioned into k clusters $C_1, \ldots, C_k$ according to the center they are assigned to in the optimal solution. We let ℓ_i denote the current phase of each optimal center c_i. The current phase lasts as long as $d(F, c_i) \in (R(\ell_i), R(\ell_i + 1)]$. As before, the demands of each optimal cluster C_i which either arrive in the current phase or remain unsatisfied from the previous phases are partitioned into inner and outer demands. In addition, the inner demands of each optimal cluster C_i can be thought of as being essentially located at c_i.

Therefore, the potential of each optimal center does not exceed f. However, a single new facility can significantly decrease the potential of many optimal centers. As a result, if we bound the decrease in the potential of each optimal center separately, we can only prove an upper bound of $\mathrm{O}(\log n)$ on the algorithm's competitive ratio. To establish a competitive ratio of $\mathrm{O}(\frac{\log n}{\log \log n})$, we use the notions of *isolated* and *non-isolated active coalitions* of optimal centers and show that the amortized (per new facility) decrease in the total potential of optimal centers remains $\mathrm{O}(f)$.

Coalitions of Optimal Centers. A set of optimal centers $K \subseteq F^*$ with representative $c_K \in K$ forms a *coalition* as long as $d(F, c_K) > \lambda D(K)$. A coalition K becomes *broken* as soon as $d(F, c_K) \leq \lambda D(K)$. A coalition K is *isolated* if $d(F, c_K) \leq \frac{1}{2}\,\mathrm{sep}_{F^*}(K)$ and *non-isolated* otherwise. We recall that $\mathrm{sep}_{F^*}(K) = d(K, F^* \setminus K)$ denotes the distance separating the optimal centers in K from the remaining optimal centers. For simplicity, we usually omit the subscript F^* and write $\mathrm{sep}(K)$ instead of $\mathrm{sep}_{F^*}(K)$.

Intuitively, we can think of a coalition K as a single optimal center located at c_K. More specifically, as long as K's diameter is much smaller than its distance from the algorithm's facilities (K is a coalition), the total potential of the optimal centers in K (i.e., the potential accumulated by the corresponding unsatisfied inner demands) does not exceed f. In the following, we sometimes refer to the total potential of the optimal centers in a coalition K as the potential of c_K (i.e., the potential of K's representative) or even the potential of K.

For each isolated coalition K, there is a facility in F which is significantly closer to K than any optimal center not in K. At an intuitive level, we can think of an isolated coalition K as if there were no optimal centers outside K. More specifically, given a collection of disjoint isolated coalitions, a new facility can be closer than the existing facilities in F to at most one of them. Hence, a new facility can decrease the potential of at most one isolated coalition and this decrease cannot exceed f. In other words, a collection of disjoint isolated coalitions can be analyzed independently of each other and similarly to the special case that the optimal solution consists of a single facility.

We use a hierarchical decomposition to maintain a partition of F^* into coalitions. A *hierarchical decomposition* $\mathcal{K}$ of F^* is a complete laminar set system[4] on F^*. Every hierarchical decomposition of F^* contains at most $2|F^*| - 1$ distinct sets.

Given a hierarchical decomposition $\mathcal{K}$ of F^*, we fix an arbitrary representative c_K for each $K \in \mathcal{K}$ and interpret $\mathcal{K}$ as a system of coalitions which hierarchically covers F^*. Formally, given a hierarchical decomposition $\mathcal{K}$ of F^* and the current algorithm's facility configuration F, a set $K \in \mathcal{K}$ is *an active coalition* if K is still a coalition (i.e., $d(F, c_K) > \lambda D(K)$), while every superset of K in $\mathcal{K}$ has become broken (i.e., for every $K' \in \mathcal{K}, K \subset K'$, $d(F, c_{K'}) \leq \lambda D(K')$). The current algorithm's facility configuration induces a collection of active coalitions which form a partitioning of F^*. Furthermore, since the distance between the algorithm's facility configuration and any point in the metric space is non-increasing with time, no coalition which has become broken (isolated) can become active (resp. non-isolated) again.

Let $D_N(K) = \max\{D(K), \frac{1}{2\lambda}\,\mathrm{sep}(K)\}$. By definition, K becomes either isolated or broken as soon as $d(F, c_K) \leq \lambda D_N(K)$. Next, we show that there is a hierarchical decomposition of F^* such that no coalition K becomes active before $d(F, c_K) < (\lambda + 1)\gamma^2 D_N(K)$.

Lemma 1. *For every $\gamma \geq 8\lambda$, there is a hierarchical decomposition $\mathcal{K}$ of F^* such that for every non-isolated active coalition $K \in \mathcal{K}$, $\lambda D_N(K) < d(F, c_K) < (\lambda + 1)\gamma^2 D_N(K)$.*

Proof. A hierarchical decomposition $\mathcal{K}$ of F^* can be represented by the *decomposition tree* $T_{\mathcal{K}}$, where the nodes of the tree correspond to the sets in $\mathcal{K}$ and there are edges connecting each set with its maximal subsets. The root of $T_{\mathcal{K}}$ corresponds to F^* and there is a leaf for each singleton set $\{c\}$, $c \in F^*$. For every $K \in \mathcal{K}$ different from the root, let p_K denote the immediate ancestor/parent of K in $T_{\mathcal{K}}$.

For the lower bound, we observe that the coalition K becomes either isolated or broken as soon as $d(F, c_K) \leq \lambda D_N(K)$. For the upper bound, we first observe that the root of $T_{\mathcal{K}}$, which corresponds to F^*, is an isolated coalition because $\mathrm{sep}(F^*) = \infty$. A coalition K different from the root cannot become active before its parent-coalition p_K becomes broken. Therefore, at the moment K becomes active, it must be $d(F, c_K) \leq d(c_K, c_{\mathrm{p}_K}) + d(F, c_{\mathrm{p}_K}) \leq (\lambda + 1)D(\mathrm{p}_K)$, The following lemma implies that there exists a hierarchical decomposition of F^* such that $\gamma^2 D_N(K) > D(\mathrm{p}_K)$.

Lemma 2. *For every $\gamma \geq 16$, every metric space M admits a hierarchical decomposition $\mathcal{K}$ such that for any set $K \in \mathcal{K}$ different from M, either $D(K) > \frac{D(\mathrm{p}_K)}{\gamma^2}$ or $\mathrm{sep}_M(K) > \frac{D(\mathrm{p}_K)}{4\gamma}$.*

The proof of Lemma 2 can be found in the Appendix, Section A.2. Intuitively, Lemma 2 states that every metric space admits a hierarchical decomposition such that each component either has a large diameter or is well-separated from the remaining metric space. Well-separated components become isolated coalitions soon after they have become active, while large diameter components become broken soon after they have become active coalitions. Therefore, no coalition can have become active long before it becomes either isolated or broken.

[4] A set system is *laminar* if it contains no intersecting pair of sets. The sets K, K' form an *intersecting pair* if neither of $K \setminus K'$, $K' \setminus K$ and $K \cap K'$ are empty. A laminar set system on F^* is *complete* if it contains F^* and every singleton set $\{c\}, c \in F^*$.

Let $\mathcal{K}$ be the hierarchical decomposition of F^* implied by Lemma 2. We conclude the proof of the lemma by establishing that for every $K \in \mathcal{K}$ different from the root F^*, $\gamma^2 D_N(K) = \gamma^2 \max\{D(K), \frac{1}{2\lambda}\operatorname{sep}(K)\} > D(\mathrm{p}_K)$. If $D(K) > \frac{D(\mathrm{p}_K)}{\gamma^2}$, then $\gamma^2 D_N(K) \geq \gamma^2 D(K) > D(\mathrm{p}_K)$. Otherwise, it must be $\operatorname{sep}(K) > \frac{D(\mathrm{p}_K)}{4\gamma}$. Hence,

$$\gamma^2 D_N(K) \geq \frac{\gamma^2}{2\lambda}\operatorname{sep}(K) > \frac{\gamma^2}{2\lambda}\frac{D(\mathrm{p}_K)}{4\gamma} \geq D(\mathrm{p}_K)\,,$$

for every $\gamma \geq 8\lambda$. □

From now on, we always assume that the set of active coalitions is determined by a fixed hierarchical decomposition $\mathcal{K}$ of F^* such that for every non-isolated active coalition K, $\lambda D_N(K) < d(F, c_K) < (\lambda+1)\gamma^2 D_N(K)$.

We use the notions of isolated and non-isolated active coalitions to establish a competitive ratio of $\mathrm{O}(\frac{\log n}{\log\log n})$ for the general case that the optimal solution consists of k centers. More specifically, each isolated active coalition is analyzed independently and similarly to the special case that the optimal solution consists of a single center. On the other hand, we show that non-isolated active coalitions can only increase the competitive ratio by a constant additive term. The analysis of isolated active coalitions is based on the ideas sketched at the beginning of this section. For non-isolated active coalitions, we use the property that every coalition becomes either isolated or broken soon after it has become active.

Inner and Outer Demands. We always use c_K to denote the representative of a coalition K. For an active coalition K, ℓ_K denotes the current phase for K. Namely, ℓ_K is equal to the integer j, $-1 \leq j \leq \psi$, such that $d(F, c_K) \in (R(j), R(j+1)]$. If $d(F, c_K) > R(\psi)$, then $\ell_K = \psi$ (first phase), and if $d(F, c_K) \leq R(0)$, then $\ell_K = -1$ (last phase). We use the convention that the last phase never ends.

For an active coalition K, let $C_K = \bigcup_{c_i \in K} C_i$ denote the optimal cluster corresponding to K. C_K consists of the demands assigned to the optimal centers of K in the optimal solution. Since the set of active coalitions always forms a partitioning of F^*, when a new demand w arrives, there is a unique active coalition such that $w \in C_K$. We sometimes abuse the notation and say that a demand $w \in C_K$ is assigned to the coalition K.

For the current phase ℓ_K, the demands of C_K are partitioned into *inner* demands $\mathrm{in}(K) = \{u \in C_K : d^*_u < r(\ell_K)\}$ and *outer* demands $\mathrm{out}(K) = C_K \setminus \mathrm{in}(K)$. In addition, $\Lambda_K = \mathrm{in}(K) \cap L$ denotes the set of unsatisfied inner demands of the active coalition K. Since $r(\psi) > n$, there are no outer demands in the first phase, and since $r(-1) = 0$, there are no inner demands in the last phase.

We should emphasize that ℓ_K, $\mathrm{in}(K)$, $\mathrm{out}(K)$, and Λ_K depend on the algorithm's configuration consisting of the facility configuration F and the set of unsatisfied demands L. Hence, they may change each time the algorithm considers a new demand. We use the convention that the algorithm updates its configuration just before it performs the assignment of the new demand. Thus, we distinguish between the arrival and the assignment time of a demand because the algorithm's configuration may have changed in between. We use plain symbols to refer to the algorithm's configuration (and any function of it) at the arrival time of a new demand and primed symbols to refer to the updated algorithm's configuration (and any function of it) at the demand's assignment time. For example, F/F' denotes the facility configuration at the arrival/assignment time of a new demand. In particular, when a new facility $\hat{w}$ opens, F denotes the facility configuration just before $\hat{w}$'s opening and $F' = F \cup \{\hat{w}\}$ denotes the facility configuration after $\hat{w}$'s opening. Similarly, Λ_K/Λ'_K denotes the set of unsatisfied inner demands of the active coalition K at the arrival/assignment time of a new demand, etc.

3.2 Preliminaries

In this section, we establish some inequalities and propositions repeatedly used in the analysis of DFL. Throughout this section, K denotes an active coalition in phase ℓ_K and c_K its representative.

As long as K is an active coalition, it must be $d(F, c_K) > \lambda \max\{D(K), r(\ell_K)\}$. More specifically, it is $d(F, c_K) > \lambda D(K)$ because K is a coalition with representative c_K and $d(F, c_K) > R(\ell_K) = \lambda\, r(\ell_K)$ since ℓ_K denotes the current phase of K. Furthermore, as long as K is a non-isolated active coalition, it must be $d(F, c_K) > \lambda \max\{D_N(K), r(\ell_K)\}$ because of Lemma 1.

The main property of a coalition K is that its inner demands are included in a small ball around c_K. In particular, let u, v be inner demands of K, and let $c_u, c_v \in K$ be the optimal centers they are assigned to. By the definition of inner demands, $d(u, c_u) = d_u^* < r(\ell_K)$ and $d(v, c_v) = d_v^* < r(\ell_K)$. Therefore,

$$d(u, v) \leq d(u, c_u) + d(c_u, c_v) + d(c_v, v) < r(\ell_K) + D(K) + r(\ell_K) \leq 3 \max\{D(K), r(\ell_K)\}$$

In other words, the diameter of the set of inner demands of an active coalition K, denoted by $D(\mathrm{in}(K))$, is less than $3 \max\{D(K), r(\ell_K)\}$.

Let $u \in C_K$ be any demand of the active coalition K. Then,

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

In addition, if u is an inner demand of K (i.e., $u \in \mathrm{in}(K)$), then

$$d(u, c_K) \leq d(u, c_u) + d(c_u, c_K) < r(\ell_K) + D(K) \leq 2 \max\{D(K), r(\ell_K)\} < \tfrac{2}{\lambda} d(F, c_K) \tag{2}$$

On the other hand, if u is an outer demand of K (i.e., $u \in \mathrm{out}(K)$), then

$$d(F, u) \leq (\lambda + 1)\delta^* + ((\lambda + 1)\rho + 1) d_u^* \tag{3}$$

Proof of Ineq. (3). Since u is an outer demand of K, it must be $d_u^* \geq r(\ell_K)$. We first observe that the current phase cannot be the first one (i.e., it must be $\ell_K < \psi$) because there are no outer demands in the first phase.

If the current phase is the last one (namely, if $\ell_K = -1$), then $d(F, c_K) \leq R(0) = \lambda \delta^*$. Otherwise, $d(F, c_K) \leq R(\ell_K + 1) = \lambda\, \rho\, r(\ell_K) \leq \lambda\, \rho\, d_u^*$. In both cases, the inequality follows from Ineq. (1). □

The following proposition establishes that every time some inner demands of K become satisfied (and hence the potential of c_K decreases), the distance between c_K and the nearest facility in F decreases by a constant factor.

Proposition 1. *Let K be an active coalition, and let w be a demand opening a new facility $\hat{w}$. If B_w includes some inner demands of K (i.e., $B_w \cap \mathrm{in}(K) \neq \emptyset$), then $d(\hat{w}, c_K) < \frac{1}{3} d(F, c_K)$.*

Proof. Let u be an inner demand of K which also belongs to B_w. Thus, $d(u, c_K) < \frac{2}{\lambda} d(F, c_K)$ by Ineq. (2), and $d(u, w) \leq \frac{1}{x} d(F, w)$ because $u \in B_w$. We observe that

$$d(F, w) \leq d(F, c_K) + d(u, c_K) + d(u, w) < (1 + \tfrac{2}{\lambda}) d(F, c_K) + \tfrac{1}{x} d(F, w) \tag{4}$$

Therefore, $d(F, w) < \frac{(\lambda+2)x}{\lambda(x-1)} d(F, c_K)$. Combining the fact that $\hat{w}$ also belongs to B_w with this upper bound on $d(F, w)$, we obtain that

$$d(\hat{w}, c_K) \leq d(\hat{w}, w) + d(u, w) + d(u, c_K) < \tfrac{2}{x} d(F, w) + \tfrac{2}{\lambda} d(F, c_K) < \tfrac{2}{\lambda}(\tfrac{\lambda+2}{x-1} + 1) d(F, c_K)$$

Using $x \geq 10$ and $\lambda = 3x + 2$, we conclude that $d(\hat{w}, c_K) < \frac{8x+6}{(3x+2)(x-1)} < \frac{3}{x} d(F, c_K) < \frac{1}{3} d(F, c_K)$. □

The following propositions essentially state that each isolated active coalition can be analyzed independently of the remaining active coalitions and similarly to the special case that the optimal solution consists of a single facility.

Proposition 2. *Let $\hat{w}$ be a new facility, and let K be an isolated active coalition such that $d(\hat{w}, c_K) < d(F, c_K)$. For every isolated active coalition K' different from K, $d(\hat{w}, c_{K'}) > d(F, c_{K'})$.*

Proof. The proposition is an immediate consequence of the definition of isolated active coalitions.

We first observe that $d(c_K, c_{K'}) \geq \max\{\mathrm{sep}(K), \mathrm{sep}(K')\}$ because K and K' are disjoint sets of optimal centers, since both of them are isolated active coalitions when $\hat{w}$ opens (recall that $\mathrm{sep}(K) = d(K, F^* \setminus K)$). Assuming that not only $d(\hat{w}, c_K) < d(F, c_K)$ but also $d(\hat{w}, c_{K'}) \leq d(F, c_{K'})$, we reach a contradiction:

$$d(c_K, c_{K'}) \leq d(\hat{w}, c_K) + d(\hat{w}, c_{K'}) < d(F, c_K) + d(F, c_{K'}) \leq \tfrac{1}{2}\,\mathrm{sep}(K) + \tfrac{1}{2}\,\mathrm{sep}(K') \leq d(c_K, c_{K'})$$

The third inequality follows from the definition of isolated coalitions. ☐

Proposition 3. *Let w be a demand opening a new facility $\hat{w}$, and let K be an isolated active coalition such that $d(\hat{w}, c_K) < d(F, c_K)$. For every active coalition K' different from K, B_w does not include any inner demands of K' (i.e., $B_w \cap \mathrm{in}(K') = \emptyset$).*

Proof. The proof follows from the definition of isolated active coalitions and Proposition 1.

To reach a contradiction, we assume that there exists an active coalition K' such that B_w includes some inner demands of K'. Then,

$$\begin{aligned} d(c_K, c_{K'}) &\leq d(\hat{w}, c_K) + d(\hat{w}, c_{K'}) < d(F, c_K) + \tfrac{1}{3} d(F, c_{K'}) \\ &\leq d(F, c_K) + \tfrac{1}{3} d(F, c_K) + \tfrac{1}{3} d(c_K, c_{K'}) \end{aligned}$$

The second inequality follows from (i) the hypothesis that $d(\hat{w}, c_K) < d(F, c_K)$, and (ii) Proposition 1, since we have assumed that $B_w \cap \mathrm{in}(K') \neq \emptyset$.

By the inequality above and the definition of isolated coalitions, we obtain that

$$d(c_K, c_{K'}) < 2\, d(F, c_K) \leq \mathrm{sep}(K)$$

which is a contradiction. More specifically, since both K and K' are active coalitions at w's arrival time, they are disjoint sets of optimal centers. Therefore, $\mathrm{sep}(K) = d(K, F^* \setminus K) \leq d(c_K, c_{K'})$. ☐

3.3 Basic Properties

The algorithm maintains the invariant that for every active coalition K, the potential of c_K (namely, the potential accumulated by the set of unsatisfied inner demands of K) does not exceed f.

Lemma 3. *Let K be an active coalition, and let Λ_K be the set of unsatisfied inner demands of K just before a new demand arrives. Then,* $\mathrm{Pot}(\Lambda_K) = \sum_{u \in \Lambda_K} d(F, u) < f$.

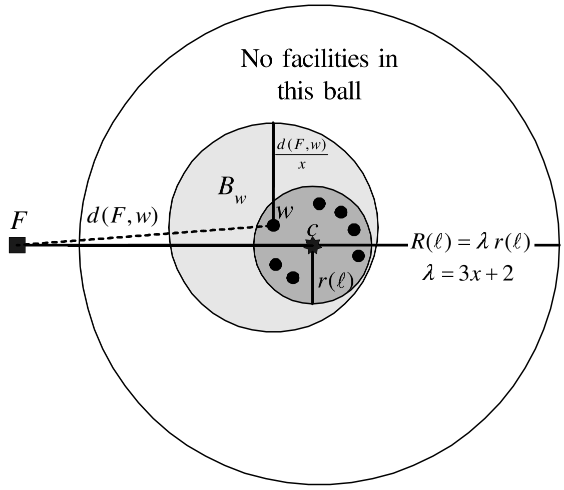

Fig. 5. For each new inner demand w, B_w includes every unsatisfied inner demand assigned to the same optimal center.

Proof. In the last phase ($\ell_K = -1$), there are no inner demands and $\mathrm{Pot}(\Lambda_K) = 0$.

Otherwise (i.e., if $\ell_K \geq 0$), let w be the demand in Λ_K arriving last, and let F_w be the facility configuration at w's arrival time. It suffices to show that the radius of B_w exceeds the diameter of $\mathrm{in}(K)$ and hence, B_w included every unsatisfied inner demand of K currently in Λ_K. Since w is an inner demand of K,

$$d(F,w) \geq d(F,c_K) - d(c_K,w) > 3x \max\{D(K), r(\ell_K)\} \tag{5}$$

The last inequality follows from (i) $d(F,c_K) > \lambda \max\{D(K), r(\ell_K)\}$ because K is an active coalition, (ii) the choice of $\lambda = 3x + 2$, and (iii) $d(w,c_K) < 2\max\{D(K), r(\ell_K)\}$ by Ineq. (2).

The radius of B_w was $\frac{d(F_w,w)}{x} \geq \frac{d(F,w)}{x} > 3\max\{D(K), r(\ell_K)\}$ (recall that $F_w \subseteq F$). On the other hand, the diameter of Λ_K is less than $3\max\{D(K), r(\ell_K)\}$ because it entirely consists of inner demands of K. Therefore, B_w must have included every unsatisfied inner demand currently in Λ_K (including w, see also Fig. 5). In addition, $\mathrm{Pot}(B_w)$ must have been less than f because w remains unsatisfied instead of opening a new facility.

Since the potential of each demand is non-increasing with time, the last time the potential of Λ_K increased was when w was added to Λ_K. Hence, we conclude that $\mathrm{Pot}(\Lambda_K) \leq \mathrm{Pot}(B_w) < f$. □

Next, we prove that if the unsatisfied inner demands of a single active coalition K contribute more than half of the cost for opening a new facility, the new facility is located so close to c_K that either K becomes broken or the current phase for K ends.

Lemma 4. *Let w be a demand opening a new facility $\hat{w}$. If there exists an active coalition K such that $\mathrm{Pot}(B_w \cap \mathrm{in}(K)) > \mathrm{Pot}(B_w)/2$, then $d(\hat{w}, c_K) < 8\max\{D(K), r(\ell_K)\}$.*

Proof. We start by considering the case that $d(F,w) \geq f$ and the location of the new facility $\hat{w}$ coincides with the location of the demand w. If there exists an active coalition K such that $w \in \mathrm{in}(K)$, the lemma follows from $d(w,c_K) < 2\max\{D(K), r(\ell_K)\}$ by Ineq. (2). For any active coalition K' such that $w \notin \mathrm{in}(K')$, $\mathrm{Pot}(B_w \cap \mathrm{in}(K')) \leq \mathrm{Pot}(\Lambda_{K'}) < \mathrm{Pot}(B_w)/2$ because $\mathrm{Pot}(B_w \setminus \Lambda_{K'}) \geq d(F,w) \geq f$, while $\mathrm{Pot}(\Lambda_{K'}) < f$ by Lemma 3.

We have also to consider the case that $d(F, w) < f$. By the formulation of the algorithm (Fig. 2), ν is a non-negative integer such that there exists at least one point $u \in B_w$ with $\mathrm{Pot}(B_w \cap \mathrm{Ball}(u, r_w/2^\nu)) > \mathrm{Pot}(B_w)/2$. Furthermore, either there exists exactly one such point in B_w or ν is the largest integer with this property, namely for every $u \in B_w$, $\mathrm{Pot}(B_w \cap \mathrm{Ball}(u, r_w/2^{\nu+1})) \leq \mathrm{Pot}(B_w)/2$ (recall that $r_w = d(F, w)/x$). In both cases, the location of $\hat{w}$ is selected so as to $\mathrm{Pot}(B_w \cap \mathrm{Ball}(\hat{w}, r_w/2^\nu)) > \mathrm{Pot}(B_w)/2$.

Every subset of B_w holding a potential greater than $\mathrm{Pot}(B_w)/2$ must have a non-empty intersection with the set of K's inner demands. Formally, $\mathrm{Ball}(\hat{w}, r_w/2^\nu) \cap \mathrm{in}(K) \neq \emptyset$.

If $r_w/2^\nu < 6 \max\{D(K), r(\ell_K)\}$, let u be any inner demand of K belonging to $\mathrm{Ball}(\hat{w}, r_w/2^\nu)$. The lemma follows from the observation that u is close to both c_K and $\hat{w}$. Formally, we can use Ineq. (2) and prove 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 $r_w/2^{\nu+1} \geq 3 \max\{D(K), r(\ell_K)\} > D(\mathrm{in}(K))$. In other words, $r_w/2^{\nu+1}$ exceeds the diameter of $\mathrm{in}(K)$. Hence, for any $u \in \mathrm{in}(K)$, $\mathrm{Ball}(u, r_w/2^{\nu+1})$ includes every inner demand of K. Therefore, for any demand u which is an unsatisfied inner demand of K,

$$\mathrm{Pot}(B_w \cap \mathrm{Ball}(u, r_w/2^{\nu+1})) \geq \mathrm{Pot}(B_w \cap \mathrm{in}(K)) > \mathrm{Pot}(B_w)/2$$

By the definition of ν, there must be a single unsatisfied inner demand of K which corresponds to the single point $u \in B_w$ such that $\mathrm{Pot}(B_w \cap \mathrm{Ball}(u, r_w/2^\nu)) > \mathrm{Pot}(B_w)/2$. Thus, the location of $\hat{w}$ coincides with the location of an inner demand of K and the lemma follows from Ineq. (2). □

3.4 The Potential Function Argument

We proceed to develop a potential function argument and establish that the algorithm's competitive ratio is $\mathrm{O}(\frac{\log n}{\log \log n})$. Before we define the potential function and bound its increase due to the actions taken by the algorithm, we give a brief sketch of the argument.

Outline. We first consider the demands which remain unsatisfied at their assignment time (equivalently, the demands which do not open any new facilities when they arrive). We use Ineq. (3) and directly charge the assignment cost of outer demands to the corresponding optimal assignment cost. The assignment cost of each inner demand which remains unsatisfied at its assignment time is charged to the potential of the active coalition the demand is assigned to in the optimal solution. More specifically, for each active coalition K, the potential of c_K is $\Upsilon_K = |\Lambda_K| \cdot d(F, c_K)$. In simple words, the potential of c_K is equal to the number of unsatisfied inner demands of K times the distance of c_K from the nearest algorithm's facility. By Ineq. (1), $d(F, c_K)$ is a good approximation to the assignment cost of each new inner demand of K. Therefore, the increase in the potential of c_K compensates for the assignment cost of each inner demand of K which remains unsatisfied at its assignment time (cf. Lemma 6). For each active coalition K, the total assignment cost of the inner demands which remain unsatisfied at their assignment time is bounded by considering the decrease in the potential of c_K every time a new facility opens.

When a new facility opens, the algorithm is charged with a facility cost of f plus an assignment cost not exceeding $\frac{f}{x}$ for the demand opening the new facility. Hence, the total algorithm's cost is $\frac{x+1}{x} f$. In addition, the algorithm is charged with the decrease in the potential of each active coalition. Lemma 3 and Ineq. (2) imply that for each active coalition K, the potential of c_K is less than $\frac{\lambda}{\lambda-2}f$ (Lemma 5). As a consequence, the decrease in the potential of c_K cannot exceed $\frac{\lambda}{\lambda-2}f$.

Next, we show how to define a potential function which in combination with the optimal cost, compensates for the cost charged to the algorithm when a new facility opens. Let w be a demand opening a new facility, let $\hat{w}$ be the new facility itself, and let K be an active coalition at w's arrival time. We first observe that the potential of c_K can decrease only if the new facility $\hat{w}$ is closer to c_K than any existing facility (see Proposition 1 for the case that $B_w \cap \mathrm{in}(K) \neq \emptyset$). Thus, we charge the active coalition K with some part of the algorithm's cost only if $d(\hat{w}, c_K) < d(F, c_K)$.

We first consider isolated active coalitions. If there exists an isolated active coalition K such that $d(\hat{w}, c_K) < d(F, c_K)$, then K is charged with the total algorithm's cost (i.e., $\frac{x+1}{x}f$) plus the decrease in the potential of c_K (i.e., $\frac{\lambda}{\lambda-2}f$). Hence a total cost of $\frac{2(\lambda+1)}{\lambda-2}f$ is charged to K (recall that $\lambda = 3x + 2$). Since K is an isolated active coalition and $d(\hat{w}, c_K) < d(F, c_K)$, every inner demand in B_w must be an inner demand of K (Proposition 3). Therefore, if the inner demands in B_w contribute more than half of B_w's potential, Lemma 4 implies that either K becomes broken or the current phase of K ends when the new facility $\hat{w}$ opens. In both cases, the credit allocated to K by the potential function must compensate for the cost charged to K. Hence, we allocate a credit of $\mathrm{O}((\psi + 3)f)$ to each active coalition K (cf. the definition of the potential function components $\Xi_K^{(1)}$ and $\Xi_K^{(3)}$ and the proof of Lemma 8, Case Isolated).

Otherwise, the potential of the outer demands in B_w is at least $f/2$. Thus, the cost of $\frac{2(\lambda+1)}{\lambda-2}f$ can be charged to the potential of the outer demands in B_w, which in turn is within a constant factor from their optimal assignment cost (Ineq. (3)). Propositions 2 and 3 imply that for any isolated active coalition K' different from K, the potential of $c_{K'}$ does not change when $\hat{w}$ opens. Therefore, the outer demands in B_w are charged by a single isolated active coalition (cf. the proof of Lemma 8, Case Isolated).

As for non-isolated active coalitions, if there exists an isolated active coalition K such that $d(\hat{w}, c_K) < d(F, c_K)$, K is charged with the total algorithm's cost and each non-isolated active coalition is only charged with the decrease in the potential of its representative. Otherwise, in addition to the decrease in the potential of its representative, each non-isolated active coalition is also charged with its share in the algorithm's cost. Since both cases are treated similarly, we only sketch the latter case. Moreover, we restrict our attention to the most interesting case that w is an outer demand (see the proof of Lemma 8, Case Non-Isolated, for the case that w is an inner demand).

The basic idea is that the part of the algorithm's cost charged to each active coalition is proportional to the coalition's share in the potential of B_w (recall that the potential of B_w is at least f, while the algorithm's cost is $\frac{x+1}{x}f$). By Ineq. (3), we can directly charge the potential of the outer demands in B_w to the corresponding optimal assignment cost. Each inner demand in B_w must be an unsatisfied inner demand of a non-isolated active coalition because we have assumed that w is an outer demand and for any isolated active coalition K, $d(\hat{w}, c_K) \geq d(F, c_K)$. Hence, the potential of the inner demands in B_w cannot exceed the total decrease in the potential of non-isolated active coalitions (cf. Ineq. (11)). Thus, we charge each non-isolated active coalition K with (approximately) twice the decrease in the potential of c_K (cf. the proof of Lemma 8, Case Non-Isolated).

We conclude the sketch of the argument by bounding the total decrease in the potential of c_K while K is a non-isolated active coalition. We first recall that for any active coalition K, the potential of c_K is always less than $\frac{\lambda}{\lambda-2}f$ (Lemma 5). If some inner demands of K are included in B_w and become satisfied, the distance of c_K from the nearest algorithm's facility decreases by a factor of 3 (Proposition 1) and the decrease in the potential of c_K is $\mathrm{O}(\ln(\frac{d(F,c_K)}{d(F\cup\{\hat{w}\},c_K)})f)$. If no unsatisfied inner demands of K are included in B_w, the cardinality of Λ_K does not change and the decrease in the potential of c_K is also $\mathrm{O}(\ln(\frac{d(F,c_K)}{d(F\cup\{\hat{w}\},c_K)})f)$. A non-isolated active coalition K becomes either isolated or broken as soon as $d(F, c_K)$ has decreased by a factor of $(1+\frac{1}{\lambda})\gamma^2$ (Lemma 1). Therefore, the total decrease in the potential of c_K while K is a non-isolated active coalition is $\mathrm{O}(\ln(\gamma^2)f)$. We allocate a credit of $\mathrm{O}(\ln(\gamma^2)f)$ to each coalition which compensates for the decrease in the potential of c_K as long as K remains a non-isolated active coalition (cf. the definition of the potential function component $\Xi_K^{(2)}$ and the proof of Lemma 7).

The Potential Function. In the following, $\mathcal{K}$ denotes the hierarchical decomposition of F^* implied by Lemma 1. We also recall that $D_N(K) = \max\{D(K), \frac{1}{2\lambda}\,\mathrm{sep}(K)\}$. We use the potential function Φ to formalize the argument sketched above.

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

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

$$\Xi_K^{(1)} = \sum_{j=0}^{\psi} \xi^{(1)}(K,j)\ ,\quad \xi^{(1)}(K,j) = \begin{cases} f & \text{if } d(F,c_K) > R(j). \\ 0 & \text{if } d(F,c_K) \le R(j). \end{cases}$$

$$\Xi_K^{(2)} = \begin{cases} 0 & \text{if } K \text{ is the root of } \mathcal{K}\ (K = F^*). \\ f\max\left\{\ln\left(\dfrac{\min\{d(F,c_K),(\lambda+1)\gamma^2 D_N(K)\}}{\lambda D_N(K)}\right),0\right\} & \text{otherwise.} \end{cases}$$

$$\Xi_K^{(3)} = \begin{cases} 2f & \text{if } K \text{ is a non-isolated coalition } (\lambda D_N(K) < d(F,c_K)). \\ f & \text{if } K \text{ is an isolated coalition } (\lambda D(K) < d(F,c_K) \le \frac{1}{2}\,\mathrm{sep}(K)). \\ 0 & \text{if } K \text{ has become broken } (d(F,c_K) \le \lambda D(K)). \end{cases}$$

The function Υ_K is defined as $\Upsilon_K = \begin{cases} |\Lambda_K|\cdot d(F,c_K) & \text{if } K \text{ is an active coalition.} \\ 0 & \text{otherwise.} \end{cases}$

For each active coalition K, the function Ξ_K is non-increasing and can decrease only if a new facility closer to c_K opens. On the other hand, the function Υ_K increases each time a new inner demand of K remains unsatisfied at its assignment time and decreases when a new facility closer to c_K opens.

In the following, $\Delta\Phi$ denotes the change in the value of the potential function when a new demand w is considered. More specifically, Φ denotes the value of the potential function at w's arrival time and Φ' denotes the value of the potential function at w's assignment time. Then, $\Delta\Phi = \Phi' - \Phi$. The same notation is used with any of the potential function components above.

We start by proving that the potential function Φ is non-negative (Lemma 5). If a new demand w remains unsatisfied at its assignment time (in other words, if w does not open any new facilities), the algorithm incurs an assignment cost of $d(F, w)$. Lemma 6 establishes that

$$d(F,w) + \Delta\Phi \le (\lambda+1)\delta^* + ((\lambda+1)\rho+1)d_w^*$$

If w opens a new facility $\hat{w}$, the algorithm incurs a facility cost of f and an assignment cost of $d(\hat{w}, w)$. Lemma 8 establishes that

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

Lemma 5. *For any active 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 and $|\Lambda_K| = 0$.

Otherwise, the algorithm maintains the invariant that $\mathrm{Pot}(\Lambda_K) < f$ (Lemma 3). In addition, for any $u \in \Lambda_K$,

$$d(F, u) \geq d(F, c_K) - d(c_K, u) > \tfrac{\lambda-2}{\lambda} d(F, c_K) \tag{6}$$

where the last inequality follows from Ineq. (2). Therefore,

$$\Upsilon_K = |\Lambda_K| \cdot d(F, c_K) < \tfrac{\lambda}{\lambda-2} \sum_{u \in \Lambda_K} d(F, u) = \tfrac{\lambda}{\lambda-2} \mathrm{Pot}(\Lambda_K) < \tfrac{\lambda}{\lambda-2} f$$

□

Lemma 5 implies that the potential function Φ is non-negative. More specifically, if K is an active coalition and $\ell_K \geq 0$, then $\frac{\lambda+1}{\lambda} \Upsilon_K < \frac{\lambda+1}{\lambda-2} f \leq \frac{(2x+1)(\lambda+1)}{x(\lambda-2)} \Xi_K^{(1)}$. On the other hand, if either K is not active or $\ell_K = -1$, then $\Upsilon_K = 0$. In both cases, $\Phi_K \geq 0$. Consequently, $\Phi \geq 0$.

The following lemma considers the case that a new demand w remains unsatisfied at its assignment time. The algorithm is charged with w's assignment cost, which is equal to $d(F, w)$.

Lemma 6. *Let w be a new demand which remains unsatisfied at its assignment time. Then,* $d(F, w) + \Delta\Phi \leq (\lambda + 1)\delta^* + ((\lambda + 1)\rho + 1)d^*_w$.

Proof. Let K be the unique active coalition (at w's arrival time) to which w is assigned (i.e., $w \in C_K$). Since no new facilities open and no demands become satisfied, $\Delta\Xi_K = 0$ and for any active coalition K' different from K, $\Delta\Phi_{K'} = 0$.

If w is an inner demand of K, it remains an unsatisfied inner demand of K at its assignment time. Therefore, the cardinality of Λ_K increases by 1 and $\Delta\Phi = -\frac{\lambda+1}{\lambda}\Delta\Upsilon_K = -\frac{\lambda+1}{\lambda} d(F, c_K)$. Using Ineq. (1), we conclude that

$$d(F, w) + \Delta\Phi \leq \tfrac{\lambda+1}{\lambda} d(F, c_K) + d^*_w - \tfrac{\lambda+1}{\lambda} d(F, c_K) = d^*_w$$

If w is an outer demand of K, then $\Delta\Phi_K = 0$ and $\Delta\Phi = 0$. Using Ineq. (3), we conclude that $d(F, w) + \Delta\Phi \leq (\lambda + 1)\delta^* + ((\lambda + 1)\rho + 1)d^*_w$. □

We have also to consider the case that a new demand w opens a new facility $\hat{w}$ and becomes satisfied before its assignment time. In the following, let F be the algorithm's facility configuration at w's arrival time, and let $F' = F \cup \{\hat{w}\}$ be the facility configuration at w's assignment time.

The algorithm is charged with a facility cost of f and an assignment cost of $d(\hat{w}, w)$. This cost is amortized among the coalitions which are active at w's arrival time. In addition, for an active coalition K which is not in the last phase (i.e., $\ell_K \geq 0$), the value of Φ_K may increase because the potential of c_K may decrease due to the new facility. In the following, we bound K's share in the algorithm's cost plus the decrease in the potential of c_K for each coalition K being active and not in the last phase at w's arrival time. If K is either not an active coalition or in the last phase (i.e., $\ell_K = -1$) when w arrives, the potential of c_K is 0 and Φ_K cannot increase. Bounding the algorithm's cost charged to an active coalition in the last phase can be done using exactly the same arguments.

Lemma 7. *Let w be a new demand which opens a new facility $\hat{w}$. For every non-isolated active coalition K, $\frac{(2x+1)(\lambda+1)}{x(\lambda-2)}\Delta\Xi_K - \frac{(2x+1)(\lambda+1)}{x\lambda}\Delta\Upsilon_K \leq 0$.*

Proof. If either the current phase ends (i.e., $d(\hat{w}, c_K) \leq R(\ell_K)$) or K stops being a non-isolated active coalition (i.e., $d(\hat{w}, c_K) \leq \lambda D_N(K)$) when $\hat{w}$ opens, then $\Delta\Xi_K \leq -f$. Furthermore, by Lemma 5, $-\Delta\Upsilon_K < \frac{\lambda}{\lambda-2} f$ and the lemma follows.

Otherwise, let $\alpha = \frac{d(F,c_K)}{d(F',c_K)} \geq 1$ be the factor by which the distance of c_K from the nearest algorithm's facility decreases when $\hat{w}$ opens. We first prove that $\Delta\Xi_K^{(2)} = -\ln(\alpha)f$. We observe that K cannot be the root of $\mathcal{K}$, because F^* is isolated as long as it remains an active coalition. In addition, since K is a non-isolated active coalition with respect to both F and F', Lemma 1 implies that $(\lambda+1)\gamma^2 D_N(K) > d(F, c_K) \geq d(F', c_K) > \lambda D_N(K)$. Therefore,

$$\Delta\Xi_K^{(2)} = [\ln(\tfrac{d(F',c_K)}{\lambda D_N(K)}) - \ln(\tfrac{d(F,c_K)}{\lambda D_N(K)})]f = \ln(\tfrac{d(F',c_K)}{d(F,c_K)})f = -\ln(\alpha)\, f \tag{7}$$

We conclude the proof of the lemma by establishing that $-\Delta\Upsilon_K \leq \frac{\lambda}{\lambda-2}\ln(\alpha)f$.

If no inner demands of K are included in B_w, the set of unsatisfied inner demands of K does not change when $\hat{w}$ opens (i.e., $\Lambda'_K = \Lambda_K$). Therefore,

$$\Upsilon'_K = |\Lambda'_K| \cdot d(F', c_K) = |\Lambda_K| \cdot d(F', c_K) = \tfrac{1}{\alpha}\, |\Lambda_K| \cdot d(F, c_K) = \tfrac{1}{\alpha}\Upsilon_K$$

Using the equality above, we obtain that

$$-\Delta\Upsilon_K = \Upsilon_K - \Upsilon'_K = (1 - \tfrac{1}{\alpha})\Upsilon_K \leq \ln(\alpha)\Upsilon_K < \tfrac{\lambda}{\lambda-2}\ln(\alpha)f \tag{8}$$

The third inequality follows from $\ln(\alpha) \geq (1 - \frac{1}{\alpha})$, for all $\alpha \geq 1$, and the fourth inequality from Lemma 5.

If some inner demands of K are included in B_w, it must be $d(\hat{w}, c_K) < \frac{1}{3}d(F, c_K)$ by Proposition 1. Hence, $d(F', c_K) < \frac{1}{3}d(F, c_K)$ and $\alpha > 3$. Therefore,

$$-\Delta\Upsilon_K < \tfrac{\lambda}{\lambda-2} f < \tfrac{\lambda}{\lambda-2}\ln(\alpha)f$$

The first inequality follows from Lemma 5 and the second inequality from $\alpha > 3$. □

We are now ready to prove the main technical lemma of this section. This lemma bounds the algorithm's cost plus the total increase in the potential function when a new facility opens and concludes the potential function argument.

Lemma 8. *Let w be a new demand which opens a new facility $\hat{w}$. Then,*

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

Proof. Throughout the proof of this lemma, let Λ_w denote the set of inner demands in B_w, and let $M_w = B_w \setminus \Lambda_w$ denote the set of outer demands in B_w.

We first observe that the algorithm's cost is at most $\frac{x+1}{x}f$. More specifically, if $d(F, w) \geq f$, then the locations of $\hat{w}$ and w coincide and $d(\hat{w}, w) = 0$. Otherwise, $\hat{w}$ is included in B_w and $d(\hat{w}, w) \leq \frac{d(F,w)}{x} < \frac{f}{x}$. In both cases, $d(\hat{w}, w) < \frac{f}{x}$ and the claim follows.

We distinguish between the case that there exists an isolated active coalition K such that $d(\hat{w}, c_K) < d(F, c_K)$ (*Case Isolated*) and the case that for every isolated active coalition K, $d(\hat{w}, c_K) \geq d(F, c_K)$ (*Case Non-Isolated*).

Case Isolated. Let K be an isolated active coalition such that $d(\hat{w}, c_K) < d(F, c_K)$. We first show that either the optimal assignment cost of the outer demands in B_w or the function Ξ_K compensates for both the total algorithm's cost and the decrease in the potential of c_K.

We start with the case that $\mathrm{Pot}(\Lambda_w) \leq \mathrm{Pot}(B_w)/2$. Namely, the inner demands in B_w contribute no more than half of B_w's potential. Therefore, the potential of the outer demands in B_w is $\mathrm{Pot}(M_w) \geq f/2$ and compensate for the cost charged to K. Formally,

$$\begin{aligned}\tfrac{x+1}{x} f + \Delta\Phi_K &\leq \tfrac{x+1}{x} f - \tfrac{\lambda+1}{\lambda}\Delta\Upsilon_K < \tfrac{x+1}{x} f + \tfrac{\lambda+1}{\lambda}\tfrac{\lambda}{\lambda-2} f \\ &\leq \tfrac{2(\lambda+1)}{\lambda-2} f \leq \tfrac{4(\lambda+1)}{\lambda-2}\,\mathrm{Pot}(M_w) \\ &\leq \tfrac{4(\lambda+1)}{\lambda-2}[(\lambda+1)\delta^*|B_w| + ((\lambda+1)\rho+1)\textstyle\sum_{u\in B_w} d_u^*]\end{aligned}$$

The first inequality holds because Ξ_K is non-increasing, the second inequality follows from Lemma 5, the third inequality holds for all $\lambda \geq 3x+2$, the fourth inequality follows from $\mathrm{Pot}(M_w) \geq f/2$, and the last inequality follows by bounding the potential of each outer demand in B_w using Ineq. (3).

We have also to consider the case that the inner demands in B_w contribute more than half of B_w's potential. Since K is an isolated active coalition and $d(\hat{w}, c_K) < d(F, c_K)$, by Proposition 3, every inner demand in B_w must be an inner demand of K. Therefore, $\mathrm{Pot}(\Lambda_w) = \mathrm{Pot}(B_w \cap \mathrm{in}(K)) > \mathrm{Pot}(B_w)/2$ and by Lemma 4, $d(\hat{w}, c_K) < 8\max\{D(K), r(\ell_K)\}$. Hence, either the current phase ends, in which case $\Delta\Xi_K^{(1)} \leq -f$, or K becomes broken, in which case $\Delta\Xi_K^{(3)} \leq -f$. In both cases, $\Delta\Xi_K \leq -f$.

On the other hand, using $-\Delta\Upsilon_K < \frac{\lambda}{\lambda-2} f$, we obtain that

$$\tfrac{x+1}{x} f - \tfrac{\lambda+1}{\lambda}\Delta\Upsilon_K < (\tfrac{x+1}{x} + \tfrac{\lambda+1}{\lambda-2}) f < \tfrac{(2x+1)(\lambda+1)}{x(\lambda-2)} f \tag{9}$$

Therefore, the decrease in Ξ_K compensates for both the total algorithm's cost and the decrease in the potential of c_K. Formally, $\frac{x+1}{x} f + \Delta\Phi_K \leq 0$.

For any coalition K' different from K which is active when w arrives, we prove that the potential function $\Phi_{K'}$ does not increase when $\hat{w}$ opens (i.e., $\Delta\Phi_{K'} \leq 0$).

If K' is a non-isolated active coalition when w arrives, the claim follows from Lemma 7 and the fact that the potential of $c_{K'}$ can only decrease when a new facility opens. Formally,

$$\Delta\Phi_{K'} = \tfrac{(2x+1)(\lambda+1)}{x(\lambda-2)}\Delta\Xi_{K'} - \tfrac{\lambda+1}{\lambda}\Delta\Upsilon_{K'} \leq \tfrac{(2x+1)(\lambda+1)}{x(\lambda-2)}\Delta\Xi_{K'} - \tfrac{(2x+1)(\lambda+1)}{x\lambda}\Delta\Upsilon_{K'} \leq 0 \tag{10}$$

The second inequality holds because $\frac{\lambda+1}{\lambda} \leq \frac{(2x+1)(\lambda+1)}{x\lambda}$ and $\Delta\Upsilon_{K'} \leq 0$ and the last inequality follows from Lemma 7.

If K' is an isolated active coalition when w arrives, the value of the potential function $\Phi_{K'}$ does not change (i.e., $\Delta\Phi_{K'} = 0$). More specifically, Proposition 2 implies that $d(\hat{w}, c_{K'}) > d(F, c_{K'})$ and $d(F', c_{K'}) = d(F, c_{K'})$. Furthermore, no inner demands of K' are included in B_w. Therefore, both the set of unsatisfied inner demands of K' and the distance between $c_{K'}$ and the algorithm's facility configuration are not affected by the new facility. Consequently, the values of $\Xi_{K'}$ and $\Upsilon_{K'}$ do not change and $\Delta\Phi_{K'} = 0$.

Case Typical. In this case, we assume that for every isolated active coalition K, $d(\hat{w}, c_K) \geq d(F, c_K)$. Hence, Proposition 1 implies that no inner demands of K are included in B_w, since it would be $d(\hat{w}, c_K) < \frac{1}{3}d(F, c_K)$ otherwise. As before, the values of Ξ_K and Υ_K do not change when $\hat{w}$ opens and $\Delta\Phi_K = 0$.

Then, we bound the increase in the potential function components associated with active coalitions being non-isolated at w's arrival time. We distinguish between the cases that w is an inner and w is an outer demand.

If w is an inner demand, let K be the unique non-isolated active coalition such that $w \in \mathrm{in}(K)$. We first show that the decrease in Ξ_K compensates for both the total algorithm's cost and the decrease in the potential of c_K.

Since B_w intersects the set of inner demands of K at w, it must be $d(\hat{w}, c_K) < \frac{1}{3}d(F, c_K)$ by Proposition 1. Using Equation (7), we obtain that $\Delta\Xi_K^{(2)} < -f$. Therefore,

$$\tfrac{x+1}{x} f + \Delta\Phi_K \leq \tfrac{(2x+1)(\lambda+1)}{x(\lambda-2)}\Delta\Xi_K^{(2)} + \tfrac{x+1}{x} f - \tfrac{\lambda+1}{\lambda}\Delta\Upsilon_K < -\tfrac{(2x+1)(\lambda+1)}{x(\lambda-2)} f + \tfrac{(2x+1)(\lambda+1)}{x(\lambda-2)} f = 0$$

where the second inequality follows from Ineq. (9).

For any other non-isolated active coalition K', Lemma 7 implies that $\Delta\Phi_{K'} \leq 0$ (see also Ineq. (10) above). This concludes the case that w is an inner demand.

If w is an outer demand, we use the following inequality:

$$\begin{aligned} \tfrac{x+1}{x} f &\leq \tfrac{x+1}{x}\mathrm{Pot}(B_w) = \tfrac{x+1}{x}[\mathrm{Pot}(M_w) + \mathrm{Pot}(\Lambda_w)] \\ &\leq \tfrac{x+1}{x}\left[(\lambda+1)\delta^*|B_w| + ((\lambda+1)\rho+1)\sum_{u \in B_w} d_u^*\right] - \tfrac{x+1}{x}\tfrac{\lambda+1}{\lambda}\sum_{\substack{K \text{ active} \\ \Lambda_K \cap \Lambda_w \neq \emptyset}} \Delta\Upsilon_K \end{aligned} \tag{11}$$

Before we prove Ineq. (11), we show that it indeed implies the lemma. We have already established that for every isolated active coalition K, $\Delta\Phi_K = 0$ and $\Lambda_K \cap \Lambda_w = \emptyset$. Hence, Ineq. (11) states that to account for the algorithm's cost, which is at most $\frac{x+1}{x} f$, it is sufficient to charge each non-isolated active coalition K with $-\frac{x+1}{x}\Delta\Upsilon_K$. This should be added to $\Delta\Phi_K$. Using Lemma 7, we conclude that for any non-isolated active coalition K, this sum cannot be positive.

$$\begin{aligned} \Delta\Phi_K - \tfrac{x+1}{x}\tfrac{\lambda+1}{\lambda}\Delta\Upsilon_K &= \tfrac{(2x+1)(\lambda+1)}{x(\lambda-2)}\Delta\Xi_K - (\tfrac{\lambda+1}{\lambda} + \tfrac{x+1}{x}\tfrac{\lambda+1}{\lambda})\Delta\Upsilon_K \\ &= \tfrac{(2x+1)(\lambda+1)}{x(\lambda-2)}\Delta\Xi_K - \tfrac{(2x+1)(\lambda+1)}{x\lambda}\Delta\Upsilon_K \leq 0 \end{aligned}$$

The Proof of Ineq. (11). We recall that Λ_w denotes the set of inner demands in B_w and M_w denotes the set of outer demands in B_w. Therefore, we can apply Ineq. (3) and obtain that

$$\mathrm{Pot}(M_w) \leq (\lambda+1)\delta^*|M_w| + ((\lambda+1)\rho+1)\sum_{u \in M_w} d_u^* \tag{12}$$

Since we have assumed that w is an outer demand, every inner demand in B_w belongs to the set of unsatisfied inner demands of some active coalition just before w arrives. Formally, $\Lambda_w \subseteq \bigcup_{K \text{ active}} \Lambda_K$. For every coalition K which is active at w's arrival time, let $\Lambda_{K,w} = \Lambda_K \cap \Lambda_w$. Thus, the collection $\{\Lambda_{K,w} : K \text{ is an active coalition}\}$ is a partitioning of Λ_w.

For every active coalition K such that $\Lambda_{K,w} = \Lambda_K \cap \Lambda_w \neq \emptyset$, the demands in $\Lambda_{K,w}$ become satisfied and are removed from the set Λ_K. Hence, the potential of c_K decreases by at least $|\Lambda_{K,w}| \cdot d(F, c_K)$. Formally,

$$-\Delta\Upsilon_K \geq |\Lambda_{K,w}| \cdot d(F, c_K)$$

Summing up the above inequalities over all active coalitions with $\Lambda_K \cap \Lambda_w \neq \emptyset$, we obtain that

$$-\sum_{\substack{K \text{ active} \\ \Lambda_K \cap \Lambda_w \neq \emptyset}} \Delta\Upsilon_K \geq \sum_{\substack{K \text{ active} \\ \Lambda_K \cap \Lambda_w \neq \emptyset}} |\Lambda_{K,w}| \cdot d(F, c_K) = \sum_{K \text{ active}} |\Lambda_{K,w}| \cdot d(F, c_K) \tag{13}$$

Using Ineq (2) and Ineq. (13), we obtain the following upper bound on the potential of the inner demands in B_w:

$$\begin{aligned} \mathrm{Pot}(\Lambda_w) &= \sum_{u \in \Lambda_w} d(F, u) = \sum_{K \text{ active}} \sum_{u \in \Lambda_{K,w}} d(F, u) \\ &\leq \tfrac{\lambda+1}{\lambda} \sum_{K \text{ active}} |\Lambda_{K,w}| \cdot d(F, c_K) + \sum_{K \text{ active}} \sum_{u \in \Lambda_{K,w}} d_u^* \\ &\leq \sum_{u \in \Lambda_w} d_u^* - \tfrac{\lambda+1}{\lambda} \sum_{\substack{K \text{ active} \\ \Lambda_K \cap \Lambda_w \neq \emptyset}} \Delta\Upsilon_K \end{aligned} \tag{14}$$

In the second equality and the last inequality, we have used the fact that $\{\Lambda_{K,w} : K \text{ is an active coalition}\}$ is a partitioning of Λ_w. In the last inequality, we have also used Ineq. (13). The third inequality follows from Ineq. (2).

Ineq. (11) follows by combining Ineq. (12) with Ineq. (14). $\square$

Concluding the Proof of Theorem 2. The potential function Φ allocates a credit of $\frac{(2x+1)(\lambda+1)}{x(\lambda-2)}[\psi + 3 + \ln(\frac{\lambda+1}{\lambda}\gamma^2)]\, f$ to each component $K \in \mathcal{K}$. Since there are at most $2k-1$ different components, this credit does not exceed $\frac{2(2x+1)(\lambda+1)}{x(\lambda-2)}[\psi + 3 + \ln(\frac{\lambda+1}{\lambda}\gamma^2)]\,\mathrm{Fac}^*$. Moreover, the optimal assignment cost of each demand is considered at most once by Lemma 6 (i.e., when the demand arrives and becomes unsatisfied) and at most once by Lemma 8 (i.e., when the demand becomes satisfied). Therefore, the total algorithm's cost is

$$\tfrac{2(2x+1)(\lambda+1)}{x(\lambda-2)}[\psi + 3 + \ln(\tfrac{\lambda+1}{\lambda}\gamma^2)]\,\mathrm{Fac}^* + \tfrac{5\lambda+2}{\lambda-2}[(\lambda+1)\rho + \lambda + 2]\,\mathrm{Asg}^*$$

Since $x \geq 10$, $\lambda = 3x + 2$, and $\gamma = 8\lambda$ are constants, setting $\psi = \rho = \mathrm{O}(\frac{\log n}{\log\log n})$ yields the claimed competitive ratio.

Remark. In the proof of Theorem 2, we have tried to simplify the presentation instead of optimizing the constants involved in the proven competitive ratio. By a more careful analysis, we can improve these constants and establish a competitive ratio of $\mathrm{O}(\frac{\log n}{\log\log n})$ for any $x > 3$.

4 A Deterministic Algorithm for Non-Uniform Facility Costs

The algorithm Non-Uniform Deterministic Facility Location - NDFL (Fig. 6) is a generalization of the algorithm DFL. In case of non-uniform facility costs, the metric space along with the cost of opening a facility at each point are given to the online algorithm in advance. Without loss of generality, we assume that the minimum facility cost is 1.

Round down the facility costs to the nearest integral power of 2.
Type of facility $\hat{v}: p_{\hat{v}} = \log f_{\hat{v}}$.
For each facility type $p \geq 0$, $Q_p = \{\text{facilities } \hat{v} : f_{\hat{v}} \leq 2^p\}$. $Q_{-1} = \emptyset$.
Let x be an appropriately chosen constant.
$F \leftarrow \emptyset;\ L \leftarrow \emptyset;$
For each new demand w:
    $L \leftarrow L \cup \{w\};\ r_w \leftarrow d(F,w)/x;\ B_w \leftarrow \mathrm{Ball}(w, r_w) \cap L;$
    $\mathrm{Pot}(B_w) \leftarrow \sum_{u \in B_w} d(F,u);\ p_w \leftarrow \max\{\lfloor \log \mathrm{Pot}(B_w) \rfloor, -1\};$
    if ($|B_w| > 1$) then `/* Compute` $\bar{w}$ `*/`
        let $\nu \geq 0$ be the smallest integer:
            either there exists exactly one point $u \in B_w$ such that
                $\mathrm{Pot}(B_w \cap \mathrm{Ball}(u, r_w/2^\nu)) > \mathrm{Pot}(B_w)/2$,
            or for every $u \in B_w$, $\mathrm{Pot}(B_w \cap \mathrm{Ball}(u, r_w/2^{\nu+1})) \leq \mathrm{Pot}(B_w)/2$.
            let $\bar{w}$ be any point in B_w: $\mathrm{Pot}(B_w \cap \mathrm{Ball}(\bar{w}, r_w/2^\nu)) > \mathrm{Pot}(B_w)/2$.
        if $((\mathrm{Ball}(\bar{w}, r_w) \cup \mathrm{Ball}(w, r_w)) \cap Q_{p_w} \neq \emptyset)$ then
            `/* Open the cheapest facility in` $\mathrm{Ball}(\bar{w}, r_w) \cup \mathrm{Ball}(w, r_w)$ `*/`
            let p be the smallest facility type in $\mathrm{Ball}(\bar{w}, r_w) \cup \mathrm{Ball}(w, r_w)$.
            let $\hat{w}$ be the facility in $(\mathrm{Ball}(\bar{w}, r_w) \cup \mathrm{Ball}(w, r_w)) \cap Q_p$ closest to $\bar{w}$.
            $F \leftarrow F \cup \{\hat{w}\};\ L \leftarrow L \setminus B_w;$
    else `/*` $|B_w| = 1$, $\bar{w}$ `coincides with` w `*/`
        $\mu \leftarrow 1;$
        while $(\mathrm{Ball}(w, d(F,w)/x) \cap Q_{p_w} \neq \emptyset)$ do
            let p be the smallest facility type in $\mathrm{Ball}(w, d(F,w)/x)$.
            let $\hat{w}_\mu$ be the facility in $\mathrm{Ball}(w, d(F,w)/x) \cap Q_p$ closest to w.
            $F \leftarrow F \cup \{\hat{w}_\mu\};\ \mu \leftarrow \mu + 1;$
    `/* Next iteration:` $r_w = d(F,w)/x$, $B_w = \{w\}$, `and` $\mathrm{Pot}(B_w) = d(F,w)$. `*/`
            $p_w \leftarrow \max\{\lfloor \log d(F,w) \rfloor, -1\};$
    Assign w to the nearest facility in F;

Fig. 6. The algorithm Non-Uniform Deterministic Facility Location - NDFL.

The algorithm starts by rounding down the facility costs to the nearest integral power of 2. This increases the actual facility cost incurred by the algorithm by a factor of 2. Then, instead of considering the (rounded down) facility costs, the algorithm considers each facility's type, which is the logarithm to base 2 of its cost.

Similarly to DFL, the algorithm's configuration consists of the facility configuration F and the set of unsatisfied demands L. Each new demand w becomes unsatisfied and is added to L upon arrival. Then, the algorithm computes w's unsatisfied neighborhood B_w and its potential $\mathrm{Pot}(B_w)$. NDFL opens the *cheapest facility* close to w if its cost does not exceed $\mathrm{Pot}(B_w)$. Since there may be many cheapest facilities close to w, the algorithm selects an appropriate one. Finally, the new demand w is assigned to the nearest facility in F.

If w is not the only demand in its unsatisfied neighborhood (i.e., $|B_w| > 1$), the algorithm computes the center of the smallest radius subset of B_w which contributes more than half of B_w's potential. This is the point denoted by $\bar{w}$. In DFL (Fig. 2), where the cost of opening a new facility is the same for all points, the location of the new facility coincides with $\bar{w}$. However, in case of non-uniform facility costs, opening the nearest facility to $\bar{w}$ may be much more expensive than opening the cheapest facility close to $\bar{w}$ and w. Thus, the algorithm opens the cheapest facility in $\mathrm{Ball}(w, r_w) \cup \mathrm{Ball}(\bar{w}, r_w)$ provided that its cost does not exceed $\mathrm{Pot}(B_w)$. To establish an (asymptotically) optimal competitive ratio, it is crucial that ties are always resolved in favour of $\bar{w}$ (cf. the proof of Lemma 11). Namely, if there are many facilities of the same (cheapest) cost in $\mathrm{Ball}(w, r_w) \cup \mathrm{Ball}(\bar{w}, r_w)$, preference is given to the facility closest to $\bar{w}$. After the new facility has opened, the demands in B_w become satisfied and are removed from L.

In case that w is the only demand in its unsatisfied neighborhood B_w (e.g., w is the first demand to arrive), w remains unsatisfied at its assignment time even if the algorithm opens some new facilities. In this case, the algorithm keeps opening the cheapest facility in $\mathrm{Ball}(w, r_w)$ as long as there is a facility of cost not exceeding $\mathrm{Pot}(B_w)$. Since the locations of w and $\bar{w}$ coincide, the cheapest facility closest to w is opened. After opening a new facility, the algorithm updates r_w and $\mathrm{Pot}(B_w)$ according to the new facility configuration and iterates. Hence, the algorithm ensures that (i) every new facility is the cheapest facility in w's neighborhood at its opening time, and (ii) at w's assignment time, the potential of w is less than the cost of any facility in w's neighborhood[5].

The remaining of this section is devoted to the proof of the following theorem.

Theorem 3. *For any $x \geq 12$, the competitive ratio of Non-Uniform Deterministic Facility Location is $\mathrm{O}(\frac{\log n}{\log \log n})$.*

4.1 Notation

In the following, we always assume that the facility costs are integral powers of 2. Without loss of generality, we assume that the minimum facility cost is 1. The cost of a facility $\hat{v}$ is denoted by $f_{\hat{v}}$. The type of a facility $\hat{v}$, denoted by $p_{\hat{v}}$, is the integral power of 2 equal to $\hat{v}$'s facility cost. Formally, $p_{\hat{v}} = \log f_{\hat{v}}$. For an integer p, f_p denotes the cost of a facility of type p, i.e., $f_p = 2^p$.

Let $f_{\max}$ be the maximum facility cost, and let $p_{\max} = \log f_{\max}$ be the maximum facility type. For $p = 0, \ldots, p_{\max}$, we let Q_p denote the set of facilities of type no greater than p. For any point v, $d(Q_p, v)$ is the distance of v from the nearest facility of type at most p. For uniformity of notation, we use the convention that $f_{-1} = 0$ and $Q_{-1} = \emptyset$. By definition, $\emptyset = Q_{-1} \subseteq Q_0 \subseteq Q_1 \subseteq \cdots \subseteq Q_{p_{\max}}$, and for any point v, $\infty = d(Q_{-1}, v) \geq d(Q_0, v) \geq \cdots \geq d(Q_{p_{\max}}, v)$.

We always let w denote the new demand currently considered by the algorithm, B_w denote the unsatisfied neighborhood of w, $\bar{w}$ denote the center of the smallest radius subset of B_w contributing a potential greater than $\mathrm{Pot}(B_w)/2$, and $\hat{w}$ denote (the location of) the new facility (if any). In case that $|B_w| = 1$ and the algorithm may open more than one facilities, $\hat{w}$ denotes the facility opened in the current iteration. We also let $\tilde{B}_w$ denote the set of demands in B_w different from w, namely $\tilde{B_w} = B_w \setminus \{w\}$.

[5] The case that $|B_w| = 1$ and some new facilities open corresponds to the case that $d(F, w) \geq f$ in DFL. In DFL, the locations of w, $\bar{w}$, and $\hat{w}$ coincide. Thus, it does not make any difference whether w remains unsatisfied at its assignment time or not, since w's potential is 0 in both cases. By opening a new facility at w, DFL fulfills both properties (i) and (ii) above.

In addition to the notation introduced in the next paragraphs, we use the notation introduced in Section 3.1 with exactly the same meaning. In the following, we briefly summarize the most important notation introduced in Section 3.1.

As in the proof of Theorem 2, we fix an arbitrary sequence of n demands and compare the algorithm's cost with the cost of the offline optimal solution. The optimal solution, denoted by F^*, consists of k centers $c_1, \dots, c_k$. For each optimal center c_i, ζ_i denotes the type of c_i and f_{ζ_i} denotes the cost of c_i. Therefore, the optimal facility cost is $\mathrm{Fac}^* = \sum_{i=1}^{k} f_{\zeta_i}$. As before, $\mathrm{Asg}^* = \sum_u d^*_u$ denotes the optimal assignment cost and $\delta^* = \mathrm{Asg}^*/n$ denotes the average optimal assignment cost per demand.

The analysis is based on the notions of isolated and non-isolated active coalitions defined in Section 3.1. The only difference is that we use $\lambda = 3x + 6$ throughout the proof of Theorem 3. As in the proof of Theorem 2, we let $\mathcal{K}$ denote a fixed hierarchical decomposition of F^* which satisfies the hypothesis of Lemma 1. For any $K \in \mathcal{K}$, we let $c_K \in K$ denote the representative of the coalition K. In case of non-uniform facility costs, c_K is always the cheapest facility in the coalition K. We also let ζ_K denote the type of K's representative. Equivalently, $\zeta_K = \min_{c_i \in K}\{\zeta_i\}$ is the type of the cheapest facility in K. For every facility type p, $-1 \leq p \leq p_{\max}$, $d(Q_p, c_K)$ is the distance of K's representative from the nearest facility of type not exceeding p. We remark that $d(Q_p, c_K) = 0$, for any $p \geq \zeta_K$.

As before, we let ρ, ψ be any fixed pair of integers such that $\rho^\psi > n$. For every integer j, $-1 \leq j \leq \psi$, $r(j)$ and $R(j)$ are defined as in Section 3.1 with the only difference that $\lambda = 3x + 6$. In addition, for any coalition K, any integer j, $0 \leq j \leq \psi$, and any integer p, $-1 \leq p \leq \zeta_K$, we define

$$P_K(j,p) = \min\{\lambda \max\{D(K), r(j)\} + x\, d(Q_p, c_K), R(j+1)\}$$

By definition, $P_K(j,-1) = R(j+1)$, $P_K(j,p) \leq R(j+1)$, and $P_K(j,\zeta_K) \leq \lambda \max\{D(K), r(j)\}$. In addition, for any integers j, p, $0 \leq j \leq \psi$, $-1 \leq p \leq \zeta_K$,

$$\max\{P_K(j,p+1), d(Q_p, c_K)\} \leq P_K(j,p) \leq P_K(j,p+1) + x\, d(Q_p, c_K) \tag{15}$$

The analysis of the cost charged to each active coalition is again broken down into phases. For an active coalition K, ℓ_K denotes the current phase of K and is equal to the integer j, $-1 \leq j \leq \psi$, such that $d(F, c_K) \in (R(j), R(j+1)]$. We recall that for the current phase ℓ_K of an active coalition K, the set of inner demands of K is defined as $\mathrm{in}(K) = \{u \in C_K : d^*_u < r(\ell_K)\}$ and the set of outer demands of K is defined as $\mathrm{out}(K) = C_K \setminus \mathrm{in}(K)$. Moreover, $\Lambda_K = \mathrm{in}(K) \cap L$ denotes the set of unsatisfied inner demands of the active coalition K.

To take into account non-uniform facility costs, the analysis of the cost charged to an active coalition in the current phase ℓ_K is broken down into $\zeta_K + 1$ disjoint *stages* according to the distance between the algorithm's facility configuration F and c_K. For an active coalition K in phase ℓ_K, the stage (ℓ_K, p) lasts as long as $d(F, c_K) \in (P_K(\ell_K, p+1), P_K(\ell_K, p)]$, $-1 \leq p \leq \zeta_K - 1$. The last phase (i.e., $\ell_K = -1$) is not divided into any stages. For an active coalition K in phase ℓ_K, $\ell_K \geq 0$, we let (ℓ_K, t_K), or simply t_K if ℓ_K is clear from the context, denote the current stage for K. Namely, t_K is equal to the integer p, $-1 \leq p \leq \zeta_K - 1$, such that $d(F, c_K) \in (P_K(\ell_K, p+1), P_K(\ell_K, p)]$.

The current phase ℓ_K starts with a stage of index -1 or greater because $P_K(\ell_K, -1) = R(\ell_K + 1)$. Since $P_K(\ell_K, \zeta_K) \leq \lambda \max\{D(K), r(\ell_K)\}$, either the coalition K becomes broken or the current phase ends no later than it becomes $d(F, c_K) \leq P_K(\ell_K, \zeta_K)$. Hence, the index of the current stage t_K never exceeds $\zeta_K - 1$.

The intuition behind the partitioning of the current phase into stages is that for any active coalition K in stage (ℓ_K, t_K) and any inner demand w of K, there is a facility of type $t_K + 1$ in $\mathrm{Ball}(w, r_w)$. Hence, the potential of c_K can never exceed f_{t_K+1} as long as K is in stage (ℓ_K, t_K) (cf. Lemma 9). In addition, no inner demand of K can open a facility of type greater than $t_K + 1$ while K is in stage (ℓ_K, t_K). On the other hand, the current stage ends if the unsatisfied inner demands of a single active coalition contribute more than half of B_w's potential and a new facility of type at least $t_K + 1$ opens (cf. Lemma 10 and Lemma 11).

The current phase ℓ_K and the current stage t_K depend on the algorithm's facility configuration F and may change each time the algorithm opens a new facility. Moreover, the sets $\mathrm{in}(K)$, $\mathrm{out}(K)$, and Λ_K depend on the facility configuration F and the set of unsatisfied demands L and may change each time a new demand is considered. Similarly to the proof of Theorem 2, we let plain symbols denote the algorithm's configuration and its functions at the arrival time of a new demand, and we let primed symbols denote the algorithm's configuration and its functions at the demand's assignment time. For example, when a demand w arrives and a new facility $\hat{w}$ opens, F denotes the algorithm's facility configuration before $\hat{w}$ opens and $F' = F \cup \{\hat{w}\}$ denotes the facility configuration after $\hat{w}$ opens.

4.2 Preliminaries

Let K be an active coalition in stage (ℓ_K, t_K), and let c_K be K's representative. We recall that as long as K is an active coalition in phase ℓ_K, $d(F, c_K) > \lambda \max\{D(K), r(\ell_K)\}$. Furthermore, as long as K is a non-isolated active coalition in phase ℓ_K, $d(F, c_K) > \lambda \max\{D_N(K), r(\ell_K)\}$ (recall that $D_N(K) = \max\{D(K), \frac{1}{2\lambda}\mathrm{sep}(K)\}$ and $\mathrm{sep}(K) = \mathrm{sep}_{F^*}(K) = d(K, F^* \setminus K)$).

By definition, as long as K is an active coalition in stage (ℓ_K, t_K), it is the case that

$$P_K(\ell_K, t_K + 1) < d(F, c_K) \le P_K(\ell_K, t_K) \le R(\ell_K + 1)$$

Therefore, if K is in stage (ℓ_K, t_K), $P_K(\ell_K, t_K + 1)$ is strictly smaller than $R(\ell_K + 1)$. Hence, for any active coalition K in stage (ℓ_K, t_K),

$$P_K(\ell_K, t_K + 1) = \lambda \max\{D(K), r(\ell_K)\} + x\, d(Q_{t_K+1}, c_K) \tag{16}$$

Inequalities (1), (2), and (3) established in Section 3.2 still hold because they have only to do with the current phase of an active coalition and its inner and outer demands whose definitions remain unchanged.

Next, we prove that Proposition 1 also holds for NDFL. We start with the following observation.

Proposition 4. *Let w be a demand opening a new facility $\hat{w}$. Then, $d(\hat{w}, w) \le \frac{2}{x} d(F, w)$.*

Proof. If $\hat{w} \in \mathrm{Ball}(w, r_w)$, the proposition follows from $r_w = d(F, w)/x$. Otherwise, it must be $\hat{w} \in \mathrm{Ball}(\bar{w}, r_w)$. Therefore,

$$d(\hat{w}, w) \le d(\hat{w}, \bar{w}) + d(\bar{w}, w) \le \tfrac{2}{x} d(F, w)$$

because $\bar{w} \in \mathrm{Ball}(w, r_w)$. □

The proof of Proposition 1 for NDFL. Let K be an active coalition, and let w be a demand opening a new facility $\hat{w}$. Proposition 1 states that if B_w includes some inner demands of K (i.e., $B_w \cap \mathrm{in}(K) \neq \emptyset$), then $d(\hat{w}, c_K) < \frac{1}{3} d(F, c_K)$.

Let u be an inner demand of K which also belongs to B_w. Then,

$$d(\hat{w}, c_K) \leq d(\hat{w}, w) + d(w, u) + d(u, c_K) < \tfrac{3}{x} d(F, w) + \tfrac{2}{\lambda} d(F, c_K) < \tfrac{1}{\lambda}(\tfrac{3(\lambda+2)}{x-1} + 2) d(F, c_K)$$

The second inequality follows from (i) $d(\hat{w}, w) \leq \frac{2}{x} d(F, w)$ by Proposition 4, (ii) $d(w, u) \leq \frac{1}{x} d(F, w)$ because $u \in B_w$, and (iii) $d(u, c_K) < \frac{2}{\lambda} d(F, c_K)$ by Ineq. (2). The third inequality follows from $d(F, w) < \frac{(\lambda+2)x}{\lambda(x-1)} d(F, c_K)$, which is derived from Ineq. (4) in the proof of Proposition 1.

Using $x \geq 12$ and $\lambda = 3x + 6$, we conclude that $d(\hat{w}, c_K) < \frac{1}{3} d(F, c_K)$. □

Propositions 2 and 3 also hold for NDFL because their proofs only depend on the definition of isolated active coalitions and Proposition 1. Therefore, we can still analyze each isolated active coalition independently of the remaining active coalitions and similarly to the special case that the optimal solution consists of a single center located at the coalition's representative.

If a demand w opens a new facility $\hat{w}$, the algorithm incurs a cost no greater than $f_{\hat{w}} + d(\hat{w}, w)$. The following proposition states that this cost cannot exceed $\frac{x+2}{x} \operatorname{Pot}(B_w)$. In addition, if w is not the only demand in B_w (i.e., $|B_w| > 1$), this cost can also be bounded in terms of $\operatorname{Pot}(\tilde{B}_w)$ (recall that $\tilde{B}_w = B_w \setminus \{w\}$).

Proposition 5. *Let w be a demand opening a new facility $\hat{w}$. Then, $f_{\hat{w}} + d(\hat{w}, w) \leq \frac{x+2}{x} \operatorname{Pot}(B_w)$. Moreover, if $|B_w| > 1$, then $f_{\hat{w}} + d(\hat{w}, w) \leq \frac{2x+6}{x} \operatorname{Pot}(\tilde{B}_w)$.*

Proof. The facility cost of $\hat{w}$ does not exceed $\operatorname{Pot}(B_w)$. Thus, the first part of the proposition follows from $d(\hat{w}, w) \leq \frac{2}{x} d(F, w)$ by Proposition 4 and $d(F, w) \leq \operatorname{Pot}(B_w)$ because $w \in B_w$.

To prove the second part, we first show that if $|B_w| > 1$, then $\operatorname{Pot}(B_w) \leq \frac{2x-1}{x-1} \operatorname{Pot}(\tilde{B}_w)$. If w is not the only demand in B_w, let $u \in B_w$ be any demand different from w. Then,

$$d(F, w) \leq \tfrac{x}{x-1} d(F, u) \leq \tfrac{x}{x-1} \operatorname{Pot}(\tilde{B}_w)$$

The first inequality follows from $d(w, u) \leq \frac{1}{x} d(F, w)$ and the second inequality from $u \in \tilde{B}_w$. Using $\tilde{B}_w = B_w \setminus \{w\}$, we obtain that

$$\operatorname{Pot}(B_w) = d(F, w) + \operatorname{Pot}(\tilde{B}_w) \leq \tfrac{2x-1}{x-1} \operatorname{Pot}(\tilde{B}_w)$$

Therefore,

$$f_{\hat{w}} + d(\hat{w}, w) \leq \tfrac{x+2}{x} \operatorname{Pot}(B_w) \leq \tfrac{x+2}{x} \tfrac{2x-1}{x-1} \operatorname{Pot}(\tilde{B}_w) \leq \tfrac{2x+6}{x} \operatorname{Pot}(\tilde{B}_w)$$

where the last inequality holds for any $x \geq 4$. □

The following proposition is used for bounding the total credit allocated to the components of the hierarchical decomposition $\mathcal{K}$ by the potential function. In particular, the proposition states that if the representative of each $K \in \mathcal{K}$ is the cheapest optimal center in K, then the sum of the representatives' costs over all $K \in \mathcal{K}$ does not exceed twice the optimal facility cost.

Proposition 6. *Let $\mathcal{K}$ be a hierarchical decomposition of F^*. For each $K \in \mathcal{K}$, let c_K be the cheapest center in K, and let ζ_K be the type of c_K (i.e., $f_{\zeta_K} = \min_{c_i \in K}\{f_{\zeta_i}\}$). Then,*

$$\sum_{K \in \mathcal{K}} f_{\zeta_K} \leq 2 \sum_{i=1}^{k} f_{\zeta_i} = 2 \operatorname{Fac}^*$$

Proof sketch. The proposition holds for uniform facility costs because the hierarchical decomposition $\mathcal{K}$ consists of no more than $2k-1$ components. The proposition cannot be violated if the facility costs are non-uniform and each representative is the cheapest center in its component. For completeness, we include a formal proof in the Appendix, Section A.3. $\square$

4.3 Basic Properties

The algorithm maintains the invariant that for every active coalition K in stage (ℓ_K, t_K), the potential of c_K is less than f_{t_K+1}.

Lemma 9. *Let K be an active coalition in stage (ℓ_K, t_K), and let Λ_K be the set of unsatisfied inner demands of K just before a new demand arrives. Then,* $\mathrm{Pot}(\Lambda_K) < f_{t_K+1}$.

Proof. The proof is a generalization of the proof of Lemma 3. In the last phase ($\ell_K = -1$), there are no inner demands and $\mathrm{Pot}(\Lambda_K) = 0$.

Otherwise (i.e., $\ell_K \geq 0$), let w be the demand in Λ_K arriving last. Similarly to the proof of Lemma 3, we can prove that B_w must have included every unsatisfied inner demand of K currently in Λ_K. Therefore, $\mathrm{Pot}(\Lambda_K) \leq \mathrm{Pot}(B_w)$.

If w is the not the only demand in B_w (i.e., $|B_w| > 1$), let F_w be the algorithm's facility configuration at w's arrival time, while if w is the only demand in B_w (i.e., $|B_w| = 1$), let F_w be the facility configuration at w's assignment time. In the former case, w must not have opened any new facilities because it remains unsatisfied. Hence, F_w is also the facility configuration at w's assignment time. Let also $r_w = d(F_w, w)/x$.

Next, we prove that $\mathrm{Ball}(w, r_w)$ must have included a facility of type no greater than $t_K + 1$. Since w is an inner demand of K,

$$\begin{aligned} d(F,w) &\geq d(F,c_K) - d(c_K,w) \\ &> (3x+6)\max\{D(K), r(\ell_K)\} + x\, d(Q_{t_K+1}, c_K) - 2\max\{D(K), r(\ell_K)\} \\ &\geq 3x\max\{D(K), r(\ell_K)\} + x\, d(Q_{t_K+1}, c_K) \end{aligned}$$

The second inequality follows from (i) the hypothesis that K is in stage (ℓ_K, t_K) and thus $d(F, c_K) > P_K(\ell_K, t_K+1)$, (ii) Eq. (16), (iii) the choice of $\lambda = 3x+6$, and (iv) Ineq. (2).

Since $F_w \subseteq F$, it must have been

$$r_w = \tfrac{d(F_w,w)}{x} \geq \tfrac{d(F,w)}{x} > 3\max\{D(K), r(\ell_K)\} + d(Q_{t_K+1}, c_K)$$

On the other hand,

$$d(Q_{t_K+1}, w) \leq d(c_K, w) + d(Q_{t_K+1}, c_K) < 2\max\{D(K), r(\ell_K)\} + d(Q_{t_K+1}, c_K) < r_w$$

where the second inequality follows from Ineq. (2).

Therefore, $\mathrm{Ball}(w, r_w)$ must have included at least one facility of cost not exceeding f_{t_K+1}. Since w did not open any new facilities after the algorithm's facility configuration had become F_w, it must have been $\mathrm{Pot}(B_w) < f_{t_K+1}$ and the lemma follows. $\square$

The following lemma is the equivalent of Lemma 4 in the proof of Theorem 2. It states that if the unsatisfied inner demands of a single active coalition K contribute more than half of B_w's potential, $\bar{w}$ (namely, the center of the smallest radius subset of B_w contributing a potential greater than $\mathrm{Pot}(B_w)/2$) is located very close to c_K.

Lemma 10. *Let w be a new demand. If there exists an active coalition K such that $\mathrm{Pot}(B_w \cap \mathrm{in}(K)) > \mathrm{Pot}(B_w)/2$, then $d(\bar{w}, c_K) < 8 \max\{D(K), r(\ell_K)\}$.*

Proof. We start with the case that w is the only demand in B_w (i.e., $|B_w| = 1$) and the location of $\bar{w}$ coincides with the location of w. If there exists an active coalition K such that $w \in \mathrm{in}(K)$, the lemma follows from Ineq. (2). For any active coalition K' such that $w \notin \mathrm{in}(K')$, $\mathrm{Pot}(B_w \cap \mathrm{in}(K')) = 0 \leq \mathrm{Pot}(B_w)/2$ because $B_w = \{w\}$.

In case that w is not the only demand in B_w (i.e., $|B_w| > 1$), the proof is identical to the proof of Lemma 4, case $d(F, w) < f$. In DFL, the location of the new facility $\hat{w}$ coincides with the location of $\bar{w}$. Hence, the proof follows by replacing $\hat{w}$ with $\bar{w}$ in the proof of Lemma 4, case $d(F, w) < f$. □

The following lemma complements Lemma 10 above. In NDFL, the distance between $\bar{w}$ and the location of the new facility $\hat{w}$ may be up to $2\,r_w$. Hence, it may happen that the unsatisfied inner demands of a single active coalition K contribute more than half of B_w's potential, in which case $\bar{w}$ is very close to c_K, and the new facility $\hat{w}$ is at a distance greater than $P_K(\ell_K, t_K + 1)$ from c_K (i.e., the new facility does not conclude the current stage of K). Nevertheless, this can only happen if $\hat{w}$ is a facility of type no greater than t_K (in other words, if $\hat{w}$ is a "cheap" facility).

Lemma 11. *Let w be a demand opening a new facility $\hat{w}$, and let K be an active coalition in stage (ℓ_K, t_K). If $d(\bar{w}, c_K) < 8 \max\{D(K), r(\ell_K)\}$, then either the current stage ends (formally, $d(\hat{w}, c_K) \leq P_K(\ell_K, t_K + 1)$) or $\hat{w}$'s type is no greater than t_K.*

Proof. We show that if $d(\bar{w}, c_K) < 8 \max\{D(K), r(\ell_K)\}$ and $\hat{w}$'s type is at least $t_K + 1$, then $d(\hat{w}, c_K) \leq P_K(\ell_K, t_K + 1)$ and the current stage ends. We distinguish between two cases depending on whether r_w is sufficiently larger than $d(Q_{t_K+1}, c_K)$ or not.

If $r_w \geq 8 \max\{D(K), r(\ell_K)\} + d(Q_{t_K+1}, c_K)$ (i.e., r_w is sufficiently larger than $d(Q_{t_K+1}, c_K)$), $\mathrm{Ball}(\bar{w}, r_w)$ must include a facility of type $t_K + 1$ (or less). This is true because

$$d(Q_{t_K+1}, \bar{w}) \leq d(\bar{w}, c_K) + d(Q_{t_K+1}, c_K) < 8 \max\{D(K), r(\ell_K)\} + d(Q_{t_K+1}, c_K) \leq r_w \qquad (17)$$

We also recall that the new facility is chosen as the facility of smallest type in $\mathrm{Ball}(w, r_w) \cup \mathrm{Ball}(\bar{w}, r_w)$ which is closest to $\bar{w}$. Therefore, if the new facility is of type $t_K + 1$ (or more), its distance from $\bar{w}$ must be no more than $d(Q_{t_K+1}, \bar{w})$. Since we have assumed that $\hat{w}$'s type is at least $t_K + 1$, we conclude that:

$$d(\hat{w}, c_K) \leq d(\hat{w}, \bar{w}) + d(\bar{w}, c_K) < 16 \max\{D(K), r(\ell_K)\} + d(Q_{t_K+1}, c_K) \leq P_K(\ell_K, t_K + 1)$$

The second inequality follows from Ineq. (17) and $d(\bar{w}, c_K) < 8 \max\{D(K), r(\ell_K)\}$. The third inequality follows from Eq. (16) and holds for any $x \geq 12$ and $\lambda \geq 3x + 6$.

We have also to consider the case that $r_w < 8 \max\{D(K), r(\ell_K)\} + d(Q_{t_K+1}, c_K)$ (i.e., r_w is not sufficiently larger than $d(Q_{t_K+1}, c_K)$). We recall that $d(\hat{w}, \bar{w}) \leq 2\,r_w$ (see also Proposition 4). Hence, we conclude that:

$$\begin{aligned} d(\hat{w}, c_K) &\le d(\hat{w}, \bar{w}) + d(\bar{w}, c_K) < 2\, r_w + 8 \max\{D(K), r(\ell_K)\} \\ &< 24 \max\{D(K), r(\ell_K)\} + 2\, d(Q_{t_K+1}, c_K) \le P_K(\ell_K, t_K + 1) \end{aligned}$$

The third inequality follows from $r_w < 8 \max\{D(K), r(\ell_K)\} + d(Q_{t_K+1}, c_K)$. The last inequality follows from Eq. (16), for any $x \ge 12$ and $\lambda \ge 3x + 6$. □

4.4 The Potential Function Argument

Next we generalize the potential function argument developed in the proof of Theorem 2 and prove that the competitive ratio of NDFL remains $\mathrm{O}(\frac{\log n}{\log \log n})$.

Outline. The demands not opening any new facilities are treated in the same way as in the proof of Theorem 2, Lemma 6. The assignment cost of outer demands is directly charged to their optimal assignment cost using Ineq. (3). On the other hand, the representative of each active coalition K holds a potential of $\Upsilon_K = |\Lambda_K| \cdot d(F, c_K)$. The assignment cost of each inner demand of K which does not open any new facilities is charged to the potential of c_K (cf. Lemma 13). As before, the total assignment cost of the unsatisfied inner demands of K is bounded by considering the decrease in the potential of c_K each time a new facility opens.

If a new demand w opens some new facilities, we distinguish between the case that w's unsatisfied neighborhood B_w contains more than one demands (i.e., $|B_w| > 1$) and a single new facility $\hat{w}$ opens, and the case that B_w contains only w (i.e., $|B_w| = 1$) and one or more new facilities open.

The former case is treated along the lines of the potential function argument in Section 3.4. If $|B_w| > 1$, the new demand w will have become satisfied by its assignment time. The algorithm incurs a facility cost of $f_{\hat{w}}$ and an assignment cost of $d(\hat{w}, w)$. The algorithm is also charged with the total decrease in the potential of the coalitions being active when w arrives.

We first observe that the potential of an active coalition K in stage (ℓ_K, t_K) cannot exceed $\frac{\lambda}{\lambda-2} f_{t_K+1}$ (cf. Lemma 12). Hence the decrease in the potential of K is always bounded by $\frac{\lambda}{\lambda-2} f_{t_K+1}$. The algorithm's cost cannot exceed the potential of the outer demands in B_w, which is charged to their optimal assignment cost using Ineq. (3), plus (approximately) twice the decrease in the potential of the active coalitions whose inner demands are included in $\tilde{B}_w = B_w \setminus \{w\}$ (cf. Ineq. (18)). Hence, to take into account the algorithm's cost, it suffices to charge the algorithm with the potential of the outer demands in B_w plus (approximately) thrice the decrease in the potential of each active coalition.

Non-isolated active coalitions are treated in the same way as in Section 3.4 (cf. the definition of the potential function component $\Xi_K^{(2)}$ and Lemma 14). As for isolated active coalitions, Proposition 2 implies that there can be at most one isolated active coalition K such that $d(\hat{w}, c_K) < d(F, c_K)$. The potential of any other isolated active coalition cannot change when $\hat{w}$ opens.

Let (ℓ_K, t_K) be the current stage of K. If $\hat{w}$'s type is greater than t_K, we proceed similarly to the proof of Lemma 8, Case Isolated. In particular, if the outer demands in B_w contribute at least half of B_w's potential, the decrease in the potential of c_K is charged to the optimal assignment cost of the outer demands in B_w using Ineq. (3). Otherwise, by Proposition 3, every inner demand in B_w is an unsatisfied inner demand of K. Hence, by Lemma 10 and Lemma 11, the new facility $\hat{w}$ is so close to c_K that the current stage ends. Then, the decrease in the credit allocated to K by the potential function compensates for the decrease in the potential of c_K (cf. the definition of the potential function component $\Xi_K^{(1)}$).

The most interesting case is when $\hat{w}$'s type is no greater than t_K (namely, $\hat{w}$ is a "cheap" facility) and the current stage of K does not end when $\hat{w}$ opens. Informally, the algorithm invests on a "cheap" facility which fails to make any significant progress towards converging to c_K. Nevertheless, the algorithm must pay for the cost of the new facility and the decrease in the potential of c_K. To deal with this case, we introduce a new potential function component, denoted by $\Xi_K^{(4)}$. Each time a "cheap" facility $\hat{w}$ opens and the current stage of K does not end, the new potential function component decreases by $\mathrm{O}(\ln(\frac{d(F,c_K)}{d(F\cup\{\hat{w}\},c_K)})f_{t_K+1})$ and compensates for the decrease in the potential of c_K (cf. Lemma 15 for a formal proof of this claim and Lemma 16 for the use of the new component).

The active coalition K enters the stage (ℓ_K, t_K) as soon as $d(F, c_K)$ drops below $P_K(\ell_K, t_K)$. By Ineq. (15), $P_K(\ell_K, t_K)$ cannot exceed $(x+1)\max\{P_K(\ell_K, t_K+1), d(Q_{t_K}, c_K)\}$. On the other hand, the potential function component $\Xi_K^{(4)}$ must decrease only if $\hat{w}$'s type is no greater than t_K and the current stage of K does not end. In other words, $\Xi_K^{(4)}$ must decrease only if $d(\hat{w}, c_K) \geq \max\{P_K(\ell_K, t_K+1), d(Q_{t_K}, c_K)\}$. Therefore, it suffices $\Xi_K^{(4)}$ to allocate a credit of $\mathrm{O}(\ln(x+1)f_{t_K+1})$ for the stage (ℓ_K, t_K), which sums up to a credit of $\mathrm{O}(2(\psi+1)\ln(x+1)f_{\zeta_K})$ in total.

The case that a new demand w is the only demand in its unsatisfied neighborhood B_w (i.e., $|B_w| = 1$) is exceptional. Then, the algorithm may open one or more facilities before w's assignment time. Moreover, w remains unsatisfied (i.e., it is not removed from L) even though some new facilities open. In this case, we analyze each iteration separately using the same ideas as above (cf. Lemma 18).

To account for w's assignment cost, we use the convention that w becomes unsatisfied (i.e., is added to the set L) at its assignment time. Hence, we can analyze w's assignment cost as if no new facilities opened. More specifically, if w is an outer demand at its assignment time, its assignment cost is charged to its optimal assignment cost using Ineq. (3). If w is an inner demand at its assignment time, the increase in the potential of the active coalition w is assigned to compensates for w's assignment cost[6].

For an iteration μ in which a new facility $\hat{w}_\mu$ opens, we again distinguish between the cases that w is an inner and w is an outer demand at the beginning of iteration μ. If w is an inner demand, we show that the potential function decreases and compensates for the cost of the new facility. If w is an outer demand, we show that the cost of the new facility plus the increase in the potential function is bounded by (a constant times) w's potential at the beginning of the iteration. The key observation is that w's potential drops by a factor of x in each iteration. Therefore, the total algorithm's cost plus the total increase in the potential function over all iterations in which w is an outer demand can be charged to w's potential at the beginning of the first iteration in which w is an outer demand. This in turn can be charged to w's optimal assignment cost using Ineq. (3).

The Potential Function. As in the proof of Theorem 2, let $\mathcal{K}$ be the hierarchical decomposition of F^* implied by Lemma 1. We generalize the potential function used in Section 3.4 to deal with non-uniform

[6] This actually explains the reason that the case $|B_w| = 1$ comprises an exceptional case. On the one hand, there is no other way to account for w's assignment cost if w is an inner demand at its assignment time. On the other hand, we must make sure that the invariant of Lemma 9 is not violated after w's assignment. Thus, the algorithm can perform w's assignment only after w's potential is not enough to pay for a new facility in $\mathrm{Ball}(w, d(F, w)/x)$ (see also the proof of Lemma 9). Since the analysis requires that the cheapest facility in $\mathrm{Ball}(w, d(F, w)/x)$ should always open, the algorithm must proceed in iterations. As long as there exists a facility in w's neighborhood of cost not exceeding w's potential, the algorithm opens the cheapest such facility and iterates.

facility costs. More specifically, we let

$$\Phi = \sum_{K \in \mathcal{K}} \Phi_K \ , \ \text{where} \ \Phi_K = \tfrac{(3x+6)(\lambda+1)}{x(\lambda-2)} \, \Xi_K - \tfrac{\lambda+1}{\lambda} \, \Upsilon_K \ .$$

The function Ξ_K is now the sum of four components $\Xi_K = \Xi_K^{(1)} + \Xi_K^{(2)} + \Xi_K^{(3)} + \Xi_K^{(4)}$. The definitions and the use of the first three components generalize the definitions and the use of the corresponding components in Section 3.4. The fourth component compensates for the case that a "cheap" new facility decreases the potential of an active coalition without either concluding its current stage or making the coalition broken (cf. Lemma 15).

$$\Xi_K^{(1)} = \sum_{j=0}^{\psi} \sum_{p=0}^{\zeta_K} \xi^{(1)}(K,j,p) \ , \ \ \xi^{(1)}(K,j,p) = \begin{cases} f_p & \text{if } d(F,c_K) > P_K(j,p). \\ 0 & \text{if } d(F,c_K) \leq P_K(j,p). \end{cases}$$

$$\Xi_K^{(2)} = \begin{cases} 0 & \text{if } K \text{ is the root of } \mathcal{K} \ (K = F^*). \\ f_{\zeta_K} \max\left\{ \ln\left(\dfrac{\min\{d(F,c_K), (\lambda+1)\gamma^2 D_N(K)\}}{\lambda D_N(K)} \right), 0 \right\} & \text{otherwise.} \end{cases}$$

$$\Xi_K^{(3)} = \begin{cases} 2f_{\zeta_K} & \text{if } K \text{ is a non-isolated coalition } (\lambda D_N(K) < d(F,c_K)). \\ f_{\zeta_K} & \text{if } K \text{ is an isolated coalition } (\lambda D(K) < d(F,c_K) \leq \frac{1}{2}\,\mathrm{sep}(K)). \\ 0 & \text{if } K \text{ has become broken } (d(F,c_K) \leq \lambda D(K)). \end{cases}$$

$$\Xi_K^{(4)} = \sum_{j=0}^{\psi} \sum_{p=0}^{\zeta_K - 1} \xi^{(4)}(K,j,p) \ , \ \ \xi^{(4)}(K,j,p) = f_{p+1} \, \max\left\{ \ln\left(\frac{\min\{d(F,c_K), P_K(j,p)\}}{\max\{P_K(j,p+1), d(Q_p,c_K)\}} \right), 0 \right\}$$

As before, Υ_K is equal to the potential of c_K: $\Upsilon_K = \begin{cases} |\Lambda_K| \cdot d(F,c_K) & \text{if } K \text{ is an active coalition.} \\ 0 & \text{otherwise.} \end{cases}$

The function Ξ_K is non-increasing and can decrease only if a new facility closer to c_K opens. On the other hand, the function Υ_K increases each time a new inner demand of K remains unsatisfied at its assignment time and decreases when a new facility closer to c_K opens. As before, we let $\Delta\Phi$ denote the change in the value of the potential function when a new demand w is considered. The same notation is also used with any of the potential function components above.

We first prove that the potential of each active coalition is always bounded (cf. Lemma 12). This implies that the potential function Φ_K remains non-negative. If a new demand w does not open any new facilities when it arrives, the algorithm incurs an assignment cost of $d(F,w)$. Lemma 13 states that

$$d(F,w) + \Delta\Phi \leq (\lambda+1)\delta^* + ((\lambda+1)\rho + 1)d_w^* \ .$$

Then, we focus on the most interesting case that a new demand w opens some new facilities when it arrives. We start with some technical lemmas which bound the increase in the potential function components when a new facility opens (Proposition 7, Lemma 14, Lemma 15, and Lemma 16). Finally, we establish the main results of this section which conclude the proof of Theorem 3. More specifically, if w is a new demand with $|B_w| > 1$ and opens a new facility $\hat{w}$, the algorithm incurs a cost of $f_{\hat{w}} + d(\hat{w},w)$[7].

[7] We recall that each new facility is significantly closer to w than any existing facilities (see also Proposition 4). Hence, w is always assigned to the most recent facility to open.

Lemma 17 establishes that

$$f_{\hat{w}} + d(\hat{w}, w) + \Delta\Phi \leq \tfrac{4x+6}{x}[(\lambda+1)\delta^*|B_w| + ((\lambda+1)\rho+1)\textstyle\sum_{u\in B_w} d_u^*] \,.$$

If w is a new demand with $|B_w| = 1$ and opens some new facilities $\hat{w}_1, \ldots, \hat{w}_m$, the algorithm incurs a cost of $\sum_{\mu=1}^{m} f_{\hat{w}_\mu} + d(\hat{w}_m, w)$. Lemma 18 establishes that

$$\sum_{\mu=1}^{m} f_{\hat{w}_\mu} + d(\hat{w}_m, w) + \Delta\Phi \leq \tfrac{2x+1}{x-1}[(\lambda+1)\delta^* + ((\lambda+1)\rho+2)d_w^*] \,.$$

Lemma 12. *Let K be an active coalition in stage (ℓ_K, t_K). Then, $\Upsilon_K < \frac{\lambda}{\lambda-2} f_{t_K+1}$ if $\ell_K \geq 0$, while $\Upsilon_K = 0$ if $\ell_K = -1$.*

Proof. The proof is similar to the proof of Lemma 5. In the last phase (i.e., $\ell_K = -1$), $\Upsilon_K = 0$ because there are no inner demands and $|\Lambda_K| = 0$.

Otherwise, the algorithm maintains the invariant that $\mathrm{Pot}(\Lambda_K) < f_{t_K+1}$ (Lemma 9). In addition, for any $u \in \Lambda_K$, $d(F, u) > \frac{\lambda-2}{\lambda} d(F, c_K)$ (see also Ineq. (6) in the proof of Lemma 5). Therefore, $\Upsilon_K < \frac{\lambda}{\lambda-2}\,\mathrm{Pot}(\Lambda_K) < \frac{\lambda}{\lambda-2} f_{t_K+1}$. □

To show that Φ_K remains non-negative, we recall that the index of the current stage cannot exceed $\zeta_K - 1$ because $P_K(\ell_K, \zeta_K) \leq \lambda \max\{D(K), r(\ell_K)\}$. Hence, if K is an active coalition not in the last phase (i.e., $\ell_K \geq 0$), then $\frac{\lambda+1}{\lambda}\,\Upsilon_K < \frac{\lambda+1}{\lambda-2} f_{\zeta_K} \leq \frac{(3x+6)(\lambda+1)}{x(\lambda-2)}\,\Xi_K^{(1)}$. On the other hand, if either K is not active or K is in the last phase (i.e., $\ell_K = -1$), then $\Upsilon_K = 0$. In any case, $\Phi_K \geq 0$ and $\Phi \geq 0$.

For every new demand w not opening any new facilities, the algorithm incurs an assignment cost of $d(F, w)$. The following lemma is identical to Lemma 6 and bounds the algorithm's cost plus the change in the potential function due to such a demand w.

Lemma 13. *Let w be a demand not opening any new facilities when it arrives. Then,*

$$d(F, w) + \Delta\Phi \leq (\lambda+1)\delta^* + ((\lambda+1)\rho+1)d_w^* \,.$$

Proof. We observe that the proof of Lemma 6, Section 3.4, has only to do with the definition of inner and outer demands and the definition of the potential function component Υ_K which remain unchanged. Hence, we can directly apply the proof of Lemma 6 to bound $\Delta\Phi + d(F, w)$. □

The remaining part of the proof is devoted to the case that a new demand w opens some new facilities when it arrives. The following lemmas (Proposition 7, Lemma 14, Lemma 15, and Lemma 16) bound the increase in the potential function and its components when a new facility $\hat{w}$ is opened by w. We should emphasize that all these lemmas also hold if $|B_w| = 1$ and w opens more than one new facilities. In this case, $\hat{w}$ should be thought of as the new facility opened in the current iteration (i.e., the facility not added to the algorithm's facility configuration yet).

As usual, we use plain symbols to denote the algorithm's configuration, its functions, and the values of the potential function components before $\hat{w}$'s opening, and primed symbols to denote the corresponding values after $\hat{w}$'s opening. In particular, we let F and F' denote the algorithm's facility configuration before and after $\hat{w}$'s opening (i.e., $F' = F \cup \{\hat{w}\}$).

If either K is not an active coalition or K is in the last phase ($\ell_K = -1$) just before $\hat{w}$ opens, the potential of c_K is 0 and Φ_K cannot increase. Hence, we restrict our attention to bounding the increase in Φ_K and its components for each active coalition K which is not in the last phase ($\ell_K \geq 0$).

Proposition 7. *Let $\hat{w}$ be a new facility, and let K be an isolated (resp. non-isolated) active coalition in stage (ℓ_K, t_K). If K remains an isolated (resp. non-isolated) active coalition in stage (ℓ_K, t_K) after $\hat{w}$ opens, then $-\Delta\Upsilon_K \leq \frac{\lambda}{\lambda-2}\min\{\ln(\frac{d(F,c_K)}{d(F\cup\{\hat{w}\},c_K)}), 1\} f_{t_K+1}$. Otherwise, $-\Delta\Upsilon_K < \frac{\lambda}{\lambda-2} f_{t_K+1}$.*

Proof. By Lemma 12, $\Upsilon_K < \frac{\lambda}{\lambda-2} f_{t_K+1}$. Therefore, the decrease in Υ_K due to the new facility $\hat{w}$ can never exceed $\frac{\lambda}{\lambda-2} f_{t_K+1}$. Formally, $-\Delta\Upsilon_K < \frac{\lambda}{\lambda-2} f_{t_K+1}$.

The remaining part of the proof resembles the second part of the proof of Lemma 7. We assume that K remains an isolated (resp. non-isolated) active coalition in stage (ℓ_K, t_K) after $\hat{w}$ opens. For brevity, let $\alpha = \frac{d(F,c_K)}{d(F\cup\{\hat{w}\},c_K)} \geq 1$ denote the factor by which the distance of c_K from the nearest algorithm's facility decreases when $\hat{w}$ opens. We conclude the proof by showing that $-\Delta\Upsilon_K \leq \frac{\lambda}{\lambda-2}\ln(\alpha) f_{t_K+1}$.

Let w be the new demand opening $\hat{w}$, and let B_w be the unsatisfied neighborhood of w contributing its potential towards $\hat{w}$'s opening. If no inner demands of K are included in B_w, the set of unsatisfied inner demands of K does not change when $\hat{w}$ opens (i.e., $\Lambda'_K = \Lambda_K$). Therefore,

$$\Upsilon'_K = |\Lambda'_K| \cdot d(F \cup \{\hat{w}\}, c_K) = |\Lambda_K| \cdot d(F \cup \{\hat{w}\}, c_K) = \tfrac{1}{\alpha}\, |\Lambda_K| \cdot d(F, c_K) = \tfrac{1}{\alpha}\, \Upsilon_K$$

Using the equality above and Lemma 12, we obtain that

$$-\Delta\Upsilon_K = \Upsilon_K - \Upsilon'_K = (1 - \tfrac{1}{\alpha})\Upsilon_K \leq \ln(\alpha)\Upsilon_K < \tfrac{\lambda}{\lambda-2}\ln(\alpha) f_{t_K+1}$$

The third inequality follows from $\ln(\alpha) \geq (1 - \frac{1}{\alpha})$, for all $\alpha \geq 1$.

If some inner demands of K are included in B_w, it must be $d(\hat{w}, c_K) < \frac{1}{3} d(F, c_K)$ by Proposition 1. Therefore, $\alpha > 3$ and the proposition follows. ◻

The analysis of non-isolated active coalition is based on the following lemma.

Lemma 14. *Let $\hat{w}$ be a new facility. For every non-isolated active coalition K,*

$$\Delta\Phi_K \leq \tfrac{(3x+6)(\lambda+1)}{x(\lambda-2)}\Delta\Xi_K - \tfrac{(3x+6)(\lambda+1)}{x\lambda}\Delta\Upsilon_K \leq 0$$

Proof. For the first inequality, we observe that

$$\Delta\Phi_K = \tfrac{(3x+6)(\lambda+1)}{x(\lambda-2)}\Delta\Xi_K - \tfrac{\lambda+1}{\lambda}\Delta\Upsilon_K \leq \tfrac{(3x+6)(\lambda+1)}{x(\lambda-2)}\Delta\Xi_K - \tfrac{(3x+6)(\lambda+1)}{x\lambda}\Delta\Upsilon_K$$

because Υ_K can only decrease when a new facility opens (i.e., $\Delta\Upsilon_K \leq 0$) and $\frac{\lambda+1}{\lambda} \leq \frac{(3x+6)(\lambda+1)}{x\lambda}$. The proof of the second inequality (i.e., $\frac{(3x+6)(\lambda+1)}{x(\lambda-2)}\Delta\Xi_K - \frac{(3x+6)(\lambda+1)}{x\lambda}\Delta\Upsilon_K \leq 0$) is essentially identical to the proof of Lemma 7.

Let (ℓ_K, t_K) be the current stage of K just before $\hat{w}$ opens. We recall that the index of the current stage t_K cannot exceed $\zeta_K - 1$. If either the current stage ends (i.e., $d(\hat{w}, c_K) \leq P_K(\ell_K, t_K + 1)$) or K stops being a non-isolated active coalition (i.e., $d(\hat{w}, c_K) \leq \lambda D_N(K)$) when $\hat{w}$ opens, then $\Delta\Xi_K \leq -f_{t_K+1}$. The lemma follows from $-\Delta\Upsilon_K < \frac{\lambda}{\lambda-2} f_{t_K+1}$ (Proposition 7).

Otherwise, let $\alpha = \frac{d(F,c_K)}{d(F',c_K)}$ be the factor by which the distance of c_K from the nearest algorithm's facility decreases when $\hat{w}$ opens. Using exactly the same argument as in the proof of Lemma 7, we prove that $\Delta\Xi^{(2)}_K = -\ln(\alpha) f_{\zeta_K}$. The lemma follows from $-\Delta\Upsilon_K \leq \frac{\lambda}{\lambda-2}\ln(\alpha) f_{t_K+1}$ (Proposition 7) and $f_{t_K+1} \leq f_{\zeta_K}$. ◻

The following lemma demonstrates the use of the new potential function component $\Xi_K^{(4)}$. Let $\hat{w}$ be a new facility, and let (ℓ_K, t_K) be the current stage of an active coalition K. If the current phase of K does not end when $\hat{w}$ opens (formally, $d(\hat{w}, c_K) > P_K(\ell_K, t_K + 1)$) and $\hat{w}$'s type is at most t_K (namely, $\hat{w}$ is a "cheap" facility), then $\Xi_K^{(4)}$ decreases by $-\ln(\frac{d(F,c_K)}{d(F',c_K)}) f_{t_K+1}$ and compensates for the decrease in the potential of c_K.

Lemma 15. *Let $\hat{w}$ be a new facility, let K be an active coalition in stage (ℓ_K, t_K), $\ell_K \geq 0$, just before $\hat{w}$ opens. If K remains in stage (ℓ_K, t_K) after $\hat{w}$ opens (formally, if $d(\hat{w}, c_K) > P_K(\ell_K, t_K + 1)$) and $\hat{w}$'s type is no greater than t_K, then $\Delta\Xi_K^{(4)} \leq -\ln(\frac{d(F,c_K)}{d(F',c_K)}) f_{t_K+1}$.*

Proof. If $d(\hat{w}, c_K) \geq d(F, c_K)$, the value of $\Xi_K^{(4)}$ does not change when $\hat{w}$ opens and the lemma holds because $d(F', c_K) = d(F, c_K)$. Hence, we restrict our attention to the most interesting case that $\hat{w}$ is closer to c_K than any existing facility. Formally, $d(\hat{w}, c_K) < d(F, c_K)$ and $d(\hat{w}, c_K) = d(F', c_K) < d(F, c_K)$ (recall that $F' = F \cup \{\hat{w}\}$).

Since K is in stage (ℓ_K, t_K) before and after $\hat{w}$ opens, it must be the case that[8]:

$$P_K(\ell_K, t_K + 1) < d(F', c_K) < d(F, c_K) \leq P_K(\ell_K, t_K)$$

In addition, $d(Q_{t_K}, c_K) \leq d(\hat{w}, c_K) = d(F', c_K) < d(F, c_K)$ because Q_{t_K} includes every facility of type no greater than t_K and thus $\hat{w} \in Q_{t_K}$. Combining the inequalities above, we obtain that

$$\max\{P_K(\ell_K, t_K + 1), d(Q_{t_K}, c_K)\} \leq d(F', c_K) < d(F, c_K) \leq P_K(\ell_K, t_K)$$

Using the definition of the potential function component $\Xi_K^{(4)}$ (recall that $\ell_K, t_K \geq 0$), we conclude that

$$\begin{aligned}
\Delta\Xi_K^{(4)} &\leq \Delta\xi^{(4)}(K, \ell_K, t_K) \\
&= \left[\ln\left(\tfrac{\min\{d(F',c_K), P_K(\ell_K,t_K)\}}{\max\{P_K(\ell_K,t_K+1), d(Q_{t_K},c_K)\}}\right) - \ln\left(\tfrac{\min\{d(F,c_K), P_K(\ell_K,t_K)\}}{\max\{P_K(\ell_K,t_K+1), d(Q_{t_K},c_K)\}}\right)\right] f_{t_K+1} \\
&= -\ln\left(\tfrac{d(F,c_K)}{d(F',c_K)}\right) f_{t_K+1}
\end{aligned}$$

□

The following lemma bounds the increase in the potential function Φ_K when a new facility opens.

Lemma 16. *Let $\hat{w}$ be a new facility, let w be the demand opening $\hat{w}$, and let B_w be the unsatisfied neighborhood of w contributing its potential towards $\hat{w}$'s opening. Let also K be an active coalition, and let (ℓ_K, t_K) be the stage of K just before $\hat{w}$ opens. If $\hat{w}$'s type is greater than t_K, then $\Delta\Phi_K < \frac{\lambda+1}{\lambda-2}\mathrm{Pot}(B_w)$, while if $\hat{w}$'s type is no greater than t_K, then $\Delta\Phi_K \leq 0$.*

Proof. We restrict our attention to the case that K is not in the last phase (i.e., $\ell_K \geq 0$) because Φ_K cannot increase if K is in the last phase.

If $\hat{w}$'s type is greater than t_K, then it must be $\mathrm{Pot}(B_w) \geq f_{\hat{w}} \geq f_{t_K+1}$ because the algorithm opens a new facility only if its cost does not exceed $\mathrm{Pot}(B_w)$. On the other hand, the function Ξ_K is non-increasing and $-\Delta\Upsilon_K \leq \frac{\lambda}{\lambda-2} f_{t_K+1}$ by Proposition 7. Therefore,

$$\Delta\Phi_K = \frac{(3x+6)(\lambda+1)}{x(\lambda-2)}\Delta\Xi_K - \frac{\lambda+1}{\lambda}\Delta\Upsilon_K \leq \frac{\lambda+1}{\lambda-2} f_{t_K+1} \leq \frac{\lambda+1}{\lambda-2}\mathrm{Pot}(B_w)\ .$$

[8] We should also notice that $t_K \geq 0$ because the minimum facility cost is 1 and $\hat{w}$'s type is no greater than t_K.

Let now the new facility $\hat{w}$ be of type no greater than t_K. We also recall that the index of the current stage t_K cannot exceed $\zeta_K - 1$. Hence, $f_{t_K+1} \leq f_{\zeta_K}$.

If the current stage of K ends (i.e., $d(\hat{w}, c_K) \leq P_K(\ell_K, t_K + 1)$) when $\hat{w}$ opens, then $\Delta\Xi_K^{(1)} \leq -f_{t_K+1}$. The lemma follows from $-\Delta\Upsilon_K < \frac{\lambda}{\lambda-2} f_{t_K+1}$ (Proposition 7).

Otherwise, it is the case that $d(\hat{w}, c_K) > P_K(\ell_K, t_K + 1)$ and $\hat{w}$'s type is no greater than t_K. Hence, we can apply Lemma 15 and obtain $\Delta\Xi_K^{(4)} \leq -\ln(\frac{d(F,c_K)}{d(F',c_K)}) f_{t_K+1}$. Then, the lemma follows from $-\Delta\Upsilon_K \leq \frac{\lambda}{\lambda-2} \min\{\ln(\frac{d(F,c_K)}{d(F',c_K)}), 1\} f_{t_K+1}$ by Proposition 7. □

We are ready to establish the main results of this section. Let w be a new demand with $|B_w| > 1$ which opens a new facility $\hat{w}$. In this case, the algorithm incurs a facility cost of $f_{\hat{w}}$ and an assignment cost of $d(\hat{w}, w)$. The following lemma bounds the algorithm's cost plus the change in the potential function when the new facility opens.

Lemma 17. *Let w be a demand with $|B_w| > 1$ which opens a new facility $\hat{w}$. Then,*

$$f_{\hat{w}} + d(\hat{w}, w) + \Delta\Phi \leq \tfrac{4x+6}{x}[(\lambda+1)\delta^*|B_w| + ((\lambda+1)\rho+1)\textstyle\sum_{u\in B_w} d_u^*] \,.$$

Proof. Throughout the proof of this lemma, let Λ_w denote the set of inner demands in w's unsatisfied neighborhood B_w, and let M_w denote the set of outer demands in B_w. In addition, let $\tilde{\Lambda}_w = \Lambda_w \setminus \{w\}$ denote the set of inner demands in $\tilde{B}_w = B_w \setminus \{w\}$, and let $\tilde{M}_w = M_w \setminus \{w\}$ denote the set of outer demands in $\tilde{B}_w$.

We use the following inequality to amortize the algorithm's cost among the active coalitions whose unsatisfied inner demands are included in $\tilde{B}_w$.

$$\begin{aligned} f_{\hat{w}} + d(\hat{w}, w) &\leq \tfrac{2x+6}{x} \operatorname{Pot}(\tilde{B}_w) \leq \tfrac{2x+6}{x}[\operatorname{Pot}(M_w) + \operatorname{Pot}(\tilde{\Lambda}_w)] \\ &\leq \tfrac{2x+6}{x}\left[(\lambda+1)\delta^*|B_w| + ((\lambda+1)\rho+1)\sum_{u\in B_w} d_u^*\right] - \tfrac{2x+6}{x}\tfrac{\lambda+1}{\lambda} \sum_{\substack{K \text{ active} \\ \Lambda_K \cap \tilde{\Lambda}_w \neq \emptyset}} \Delta\Upsilon_K \end{aligned} \tag{18}$$

The Proof of Ineq. (18). The proof is similar to the proof of Ineq. (11), Lemma 8.

The first inequality follows from Proposition 5. The second inequality follows from $\tilde{B}_w \subseteq M_w \cup \tilde{\Lambda}_w$. For the third inequality, we first bound $\operatorname{Pot}(M_w)$ using Ineq. (3).

$$\operatorname{Pot}(M_w) \leq (\lambda+1)\delta^*|M_w| + ((\lambda+1)\rho+1)\sum_{u\in M_w} d_u^* \tag{19}$$

Since $w \notin \tilde{\Lambda}_w$, we can prove the following inequality using exactly the same arguments as in the proof of Ineq. (14), Lemma 8.

$$\operatorname{Pot}(\tilde{\Lambda}_w) \leq \sum_{u\in\tilde{\Lambda}_w} d_u^* - \tfrac{\lambda+1}{\lambda} \sum_{\substack{K \text{ active} \\ \Lambda_K \cap \tilde{\Lambda}_w \neq \emptyset}} \Delta\Upsilon_K \tag{20}$$

Then, Ineq. (18) follows by combining Ineq. (19) with Ineq. (20). □

As in the proof of Lemma 8, we distinguish between the case that there exists an isolated active coalition K such that $d(\hat{w}, c_K) < d(F, c_K)$ and the case that for all isolated active coalitions K, $d(\hat{w}, c_K) \geq d(F, c_K)$.

Case Isolated. Let K be an isolated active coalition such that $d(\hat{w}, c_K) < d(F, c_K)$. We first establish that the optimal assignment cost of the outer demands in B_w and the decrease in the function Ξ_K compensate for both the total algorithm's cost and the decrease in the potential of c_K.

We start by assuming that $\mathrm{Pot}(\Lambda_w) \leq \mathrm{Pot}(B_w)/2$. Thus, the outer demands in B_w contribute at least half of the potential of B_w (formally, $\mathrm{Pot}(B_w) \leq 2\,\mathrm{Pot}(M_w)$). By Proposition 5, the total algorithm's cost cannot exceed $\frac{x+2}{x}\,\mathrm{Pot}(B_w)$. By Lemma 16, if $\hat{w}$'s type is greater than t_K, then the increase in Φ_K is less than $\frac{\lambda+1}{\lambda-2}\,\mathrm{Pot}(B_w)$, while if $\hat{w}$'s type is no greater than t_K, then Φ_K does not increase at all. In both cases, $\Delta\Phi_K \leq \frac{\lambda+1}{\lambda-2}\,\mathrm{Pot}(B_w)$. Therefore,

$$\begin{aligned} f_{\hat{w}} + d(\hat{w}, w) + \Delta\Phi_K &\leq \tfrac{x+2}{x}\,\mathrm{Pot}(B_w) + \tfrac{\lambda+1}{\lambda-2}\,\mathrm{Pot}(B_w) \\ &\leq \tfrac{2x+3}{x}\,\mathrm{Pot}(B_w) \leq \tfrac{4x+6}{x}\,\mathrm{Pot}(M_w) \\ &\leq \tfrac{4x+6}{x}[(\lambda+1)\delta^*|B_w| + ((\lambda+1)\rho+1)\textstyle\sum_{u\in B_w} d^*_u] \end{aligned}$$

The second inequality holds for any $\lambda \geq 3x+6$ and the last inequality follows from Ineq. (3) since M_w is the set of the outer demands in B_w.

We have also to consider the case that $\mathrm{Pot}(\Lambda_w) > \mathrm{Pot}(B_w)/2$. In the following, let (ℓ_K, t_K) be the current stage of the isolated active coalition K just before $\hat{w}$ opens. We also recall that the index of the current stage t_K can never exceed $\zeta_K - 1$. Hence, $f_{t_K+1} \leq f_{\zeta_K}$.

Since K is an isolated active coalition and $d(\hat{w}, c_K) < d(F, c_K)$, by Proposition 3, every inner demand in B_w must be an inner demand of K. Formally, $\Lambda_w = B_w \cap \mathrm{in}(K)$. We use this observation and show that $\Delta\Xi_K \leq -f_{t_K+1}$.

Since every inner demand in B_w is an inner demand of K and $\mathrm{Pot}(\Lambda_w) = \mathrm{Pot}(B_w \cap \mathrm{in}(K)) > \mathrm{Pot}(B_w)/2$, we can apply Lemma 10 and obtain that $d(\bar{w}, c_K) < 8\max\{D(K), r(\ell_K)\}$. Then, we can apply Lemma 11 and conclude that either the current stage of K ends or $\hat{w}$'s type is no greater than t_K.

If the current stage of K ends (i.e., $d(\hat{w}, c_K) \leq P_K(\ell_K, t_K+1)$) when $\hat{w}$ opens, then $\Xi^{(1)}_K$ decreases by at least f_{t_K+1}. Otherwise, we can apply Lemma 15 since $d(\hat{w}, c_K) > P_K(\ell_K, t_K+1)$ and $\hat{w}$'s type must be no greater than t_K. Hence, we obtain that $\Delta\Xi^{(4)}_K \leq -\ln(\frac{d(F,c_K)}{d(F',c_K)})f_{t_K+1}$. Furthermore, $B_w \cap \mathrm{in}(K) \neq \emptyset$ because $\mathrm{Pot}(\Lambda_w) = \mathrm{Pot}(B_w \cap \mathrm{in}(K)) > 0$, and Proposition 1 implies that $d(\hat{w}, c_K) < \frac{1}{3}d(F, c_K)$. Therefore, $\frac{d(F,c_K)}{d(F',c_K)} > 3$ and $\Delta\Xi^{(4)}_K < -f_{t_K+1}$. This concludes the proof of the claim that $\Delta\Xi_K \leq -f_{t_K+1}$.

Since every inner demand of B_w is an inner demand of K, for any active coalition K' different from K, it must be $\Lambda_{K'} \cap \tilde{\Lambda}_w = \emptyset$. Using this observation and applying Ineq. (18), we obtain that

$$\begin{aligned} f_{\hat{w}} + d(\hat{w}, w) + \Delta\Phi_K &\leq \tfrac{2x+6}{x}[(\lambda+1)\delta^*|B_w| + ((\lambda+1)\rho+1)\textstyle\sum_{u\in B_w} d^*_u] \\ &\qquad - \tfrac{(2x+6)(\lambda+1)}{x\lambda}\Delta\Upsilon_K + \tfrac{(3x+6)(\lambda+1)}{x(\lambda-2)}\Delta\Xi_K - \tfrac{\lambda+1}{\lambda}\Delta\Upsilon_K \\ &= \tfrac{2x+6}{x}[(\lambda+1)\delta^*|B_w| + ((\lambda+1)\rho+1)\textstyle\sum_{u\in B_w} d^*_u] \\ &\qquad + \tfrac{(3x+6)(\lambda+1)}{x(\lambda-2)}\Delta\Xi_K - \tfrac{(3x+6)(\lambda+1)}{x\lambda}\Delta\Upsilon_K \\ &\leq \tfrac{2x+6}{x}[(\lambda+1)\delta^*|B_w| + ((\lambda+1)\rho+1)\textstyle\sum_{u\in B_w} d^*_u] \\ &\qquad - \tfrac{(3x+6)(\lambda+1)}{x(\lambda-2)}f_{t_K+1} + \tfrac{(3x+6)(\lambda+1)}{x\lambda}\tfrac{\lambda}{\lambda-2}f_{t_K+1} \\ &= \tfrac{2x+6}{x}[(\lambda+1)\delta^*|B_w| + ((\lambda+1)\rho+1)\textstyle\sum_{u\in B_w} d^*_u] \end{aligned}$$

The third inequality follows from $\Delta\Xi_K < -f_{t_K+1}$ and $-\Delta\Upsilon_K \leq \frac{\lambda}{\lambda-2}f_{t_K+1}$, by Proposition 7.

Similarly to the proof of Lemma 8, we show that for any active coalition K' different from K, the potential function $\Phi_{K'}$ cannot increase when $\hat{w}$ opens (i.e., $\Delta\Phi_{K'} \le 0$). More specifically, if K' is an isolated active coalition, then it must be $d(\hat{w}, c_{K'}) \ge d(F, c_{K'})$ (Proposition 2) and $\Lambda_{K'} \cap B_w = \emptyset$ (Proposition 3). Hence, none of the potential function components associated with K' is affected by the new facility $\hat{w}$, and $\Delta\Xi_{K'} = \Delta\Upsilon_{K'} = \Delta\Phi_{K'} = 0$. If K' is a non-isolated active coalition, the claim follows from Lemma 14.

Case Typical. We have also to consider the case that for any isolated active coalition K, $d(\hat{w}, c_K) \ge d(F, c_K)$. Therefore, every unsatisfied inner demand in B_w is an unsatisfied inner demand of some non-isolated active coalition (otherwise, Proposition 1 would lead to a contradiction). Thus, for any isolated active coalition K, none of the potential function components associated with K is affected by $\hat{w}$, and $\Delta\Xi_K = \Delta\Upsilon_K = \Delta\Phi_K = 0$.

Next, we use Ineq. (18) and bound the algorithm's cost plus the increase in the potential function components associated with non-isolated active coalitions.

$$\begin{aligned}
f_{\hat{w}} + d(\hat{w}, w) + \sum_{K \text{ non-isol.}} \Delta\Phi_K &\le \tfrac{2x+6}{x}[(\lambda+1)\delta^*|B_w| + ((\lambda+1)\rho+1)\textstyle\sum_{u\in B_w} d_u^*] \\
&\quad - \tfrac{2x+6}{x} \sum_{\substack{K \text{ active} \\ \Lambda_K \cap \tilde{\Lambda}_w \neq \emptyset}} \Delta\Upsilon_K + \sum_{K \text{ non-isol.}} \left[\tfrac{(3x+6)(\lambda+1)}{x(\lambda-2)}\Delta\Xi_K - \tfrac{\lambda+1}{\lambda}\Delta\Upsilon_K\right] \\
&\le \tfrac{2x+6}{x}[(\lambda+1)\delta^*|B_w| + ((\lambda+1)\rho+1)\textstyle\sum_{u\in B_w} d_u^*] \\
&\quad + \sum_{K \text{ non-isol.}} \left[\tfrac{(3x+6)(\lambda+1)}{x(\lambda-2)}\Delta\Xi_K - \tfrac{(3x+6)(\lambda+1)}{x\lambda}\Delta\Upsilon_K\right] \\
&\le \tfrac{2x+6}{x}[(\lambda+1)\delta^*|B_w| + ((\lambda+1)\rho+1)\textstyle\sum_{u\in B_w} d_u^*]
\end{aligned}$$

The first inequality follows from Ineq. (18). The second inequality follows from the fact that every unsatisfied inner demand in B_w is an unsatisfied inner demand of some non-isolated active coalition. The last inequality follows from Lemma 14. □

We have also to consider the case that a new demand w is the only the demand in its unsatisfied neighborhood B_w (in other words, $|B_w| = 1$) and the algorithm opens one or more facilities. Then, w is assigned to the most recent facility to open and remains unsatisfied (formally, w is not removed from the set L). Let $m \ge 1$ be the number of while-iterations, and let $\hat{w}_1, \dots, \hat{w}_m$ be the new facilities opened in these iteration. The total algorithm's cost is $\sum_{\mu=1}^{m} f_{\hat{w}_\mu} + d(\hat{w}_m, \hat{w})$. The following lemma bounds the total algorithm's cost and the total increase in the potential function in case that $|B_w| = 1$.

Lemma 18. *Let w be a demand with $|B_w| = 1$ which opens some new facilities, let $m \ge 1$ be the number of iterations, and let $\hat{w}_1, \dots, \hat{w}_m$ be the new facilities. Then,*

$$\sum_{\mu=1}^{m} f_{\hat{w}_\mu} + d(\hat{w}_m, w) + \Delta\Phi \le \tfrac{2x+1}{x-1}[(\lambda+1)\delta^* + ((\lambda+1)\rho+2)d_w^*] \,.$$

Proof. For uniformity, we also treat w's assignment, which takes place after the last while-iteration, as a fictitious iteration with index $m+1$. The algorithm's cost for this iteration is $d(\hat{w}_m, w)$.

We establish the lemma by analyzing each iteration separately. More specifically, we prove that if w is an inner demand with respect to the algorithm's facility configuration at the beginning of iteration

μ, the potential function decreases and its decrease compensates for the algorithm's cost in iteration μ. On the other hand, if w is an outer demand with respect to the facility configuration at the beginning of iteration μ, the algorithm's cost and the increase in the potential function is bounded by (a constant times) w's potential at the beginning of iteration μ. The key observation is that w's potential drops by a factor of x in each iteration. Therefore, the total algorithm's cost plus the total increase in the potential function over all iterations in which w is an outer demand can be charged to w's potential at the beginning of the first iteration in which w is an outer demand. The potential of w at this point is bounded using Ineq. (3).

We start by introducing some notation. Let $F_0 = F$ denote the algorithm's facility configuration at w's arrival time. For $\mu = 1, \ldots, m$, let $F_\mu = F_{\mu-1} \cup \{\hat{w}_\mu\}$ denote the algorithm's facility configuration at the beginning of iteration $\mu + 1$.

To analyze iteration μ, we regard $F_{\mu-1}$ as the current facility configuration. Hence, at the beginning of iteration μ, the values of the potential function components, the set of active coalitions, the current stage of each active coalition, the sets of inner and outer demands, etc. are defined with respect to $F_{\mu-1}$. We use the convention that w is added to the set of unsatisfied demands L at its assignment time.

We use $\Delta\Phi(\mu) = \Phi(\mu + 1) - \Phi(\mu)$ to denote the change in the value of the potential function between the beginning of iteration μ and the beginning of iteration $\mu + 1$. We eventually calculate the total change in the potential function by summing up $\Delta\Phi(\mu)$ over all iterations.

We also observe that the distance of w from the nearest algorithm's facility drops by a factor no less than x in each iteration. In particular, for every $\mu = 1, \ldots, m$, $d(F_\mu, w) \leq d(F_{\mu-1}, w)/x$ because $\hat{w}_\mu$ is within a distance of $d(F_{\mu-1}, w)/x$ from w. As a result, w's unsatisfied neighborhood B_w contains only w in all iterations. Moreover, the potential of B_w in iteration μ is equal to $d(F_{\mu-1}, w)$ and drops by a factor of x in each iteration.

We distinguish between the cases that w is an inner and w is an outer demand at the beginning of iteration μ (i.e., with respect to the facility configuration $F_{\mu-1}$).

Case Inner. Let w be an inner demand at the beginning of iteration μ. If a new facility $\hat{w}_\mu$ opens in iteration μ, then the decrease in the potential function compensates for the cost of $\hat{w}_\mu$ (formally, $f_{\hat{w}_\mu} + \Delta\Phi(\mu) \leq 0$). If w is an inner demand at its assignment time (i.e., with respect to the facility configuration F_m), then $d(\hat{w}_\mu, w) + \Delta\Phi(m + 1) \leq d^*_w$.

Let K be the unique active coalition at the beginning of iteration μ to which w is assigned in the optimal solution. Let also (ℓ_K, t_K) be the current stage of K at the beginning of iteration μ Let us first consider an iteration μ, $1 \leq \mu \leq m$, in which a new facility $\hat{w}_\mu$ opens. Using the hypothesis that w is an inner demand of K at the beginning of iteration μ, we conclude that:

(a) The set Λ_K of the unsatisfied inner demands of K is currently empty and the potential of c_K is 0 (formally, $\Upsilon_K = 0$). More specifically, in the proof of Lemma 9, we have shown that the unsatisfied neighborhood B_w of an inner demand w assigned to an active coalition K includes every unsatisfied inner demand of K. Since B_w only contains w, Λ_K is currently empty and $\Upsilon_K = 0$. Hence, the value of Υ_K cannot change in iteration μ and $\Delta\Upsilon_K(\mu) = 0$.
(b) The type of the new facility $\hat{w}_\mu$ is no greater than $t_K + 1$. More specifically, in the proof of Lemma 9, we have shown that for any inner demand w of an active coalition K in stage (ℓ_K, t_K), $\mathrm{Ball}(w, d(F, w)/x)$ includes a facility of type $t_K + 1$. Since the algorithm opens the cheapest facility in $\mathrm{Ball}(w, d(F_{\mu-1}, w)/x)$, the type of $\hat{w}_\mu$ cannot exceed $t_K + 1$. Therefore, $f_{\hat{w}_\mu} \leq f_{t_K+1}$.

We then show that $\Delta\Xi_K(\mu) \leq -f_{t_K+1}$ and hence, the decrease in the potential function Φ_K compensates for the algorithm's cost in iteration μ. The proof is similar to the second case in the proof of Lemma 17, Case Isolated, even though we do not assume that K is an isolated coalition.

If the current stage of K ends (i.e., $d(\hat{w}_\mu, c_K) \leq P_K(\ell_K, t_K+1)$) when $\hat{w}$ opens, $\Xi_K^{(1)}$ decreases by at least f_{t_K+1}. Otherwise, Lemma 11 implies that $\hat{w}_\mu$'s type is at most t_K. In particular, the locations of w and $\bar{w}$ coincide with each other. Thus, $d(w, c_K) = d(\bar{w}, c_K) < 2\max\{D(K), r(\ell_K)\}$ because w is an inner demand of K (Ineq. (2)). Since the current stage of K does not end, it must be $d(\hat{w}_\mu, c_K) > P_K(\ell_K, t_K+1)$. Hence, we can apply Lemma 11 and obtain that $\hat{w}_\mu$'s type is no greater than t_K.

Furthermore, since $d(\hat{w}_\mu, c_K) > P_K(\ell_K, t_K+1)$ and $\hat{w}_\mu$'s type is no greater than t_K, we can apply Lemma 15 and conclude that $\Xi_K^{(4)}$ decreases by $\ln(\frac{d(F_{\mu-1}, c_K)}{d(F_{\mu-1}\cup\{\hat{w}_\mu\}, c_K)})f_{t_K+1}$. In addition, Proposition 1 implies that $\frac{d(F_{\mu-1}, c_K)}{d(F_{\mu-1}\cup\{\hat{w}_\mu\}, c_K)} > 3$ because B_w intersects the set of inner demands of K at w. Therefore, $\Delta\Xi_K^{(4)}(\mu) < -f_{t_K+1}$. This concludes the proof that $\Delta\Xi_K(\mu) \leq -f_{t_K+1}$.

We also observe that for any active coalition K' different from K, $\Delta\Phi_{K'}(\mu) \leq 0$. If K' is a non-isolated active coalition, the claim follows from Lemma 14. If K' is an isolated active coalition, it must be $d(\hat{w}_\mu, c_{K'}) \geq d(F_{\mu-1}, c_{K'})$ and $\Delta\Phi_{K'}(\mu) = 0$. Otherwise (i.e., if $d(\hat{w}_\mu, c_{K'}) < d(F_{\mu-1}, c_{K'})$), Proposition 3 implies that w must have been an inner demand of K' instead of K.

Hence, we have shown that for any iteration μ in which w is an inner demand and a new facility $\hat{w}_\mu$ opens, $f_{\hat{w}_\mu} + \Delta\Phi(\mu) \leq 0$.

It remains to consider the fictitious iteration $m+1$ in which w is assigned to $\hat{w}_m$. The algorithm incurs an assignment cost of $d(F_m, w) = d(\hat{w}_m, w)$. The analysis is essentially identical to the proof of Lemma 13. We use the convention that w is added to the set of unsatisfied demands L at its assignment time. No new facilities open and none of the unsatisfied demands become satisfied at w's assignment time. Hence, only the value of Υ_K can change. In particular, since w is an inner demand of K, the potential of c_K increases by $d(F_m, c_K)$ because w is added to Λ_K. Using $d(\hat{w}_m, w) = d(F_m, w)$ and Ineq. (1), we conclude that:

$$d(\hat{w}_m, w) + \Delta\Phi(m+1) \leq \tfrac{\lambda+1}{\lambda} d(F_m, c_K) + d_w^* - \tfrac{\lambda+1}{\lambda} d(F_m, c_K) = d_w^*$$

Putting everything together, we conclude that the total algorithm's cost plus the total change in the potential function over all iterations in which w is an inner demand cannot exceed d_w^*.

Case Outer. We then assume that w is an outer demand at the beginning of iteration μ. We show that the algorithm's cost plus the increase in the potential function can be charged to w's potential at the beginning of iteration μ.

We first recall that the potential of w's unsatisfied neighborhood B_w at the beginning of iteration μ is equal to the potential of w (i.e., $d(F_{\mu-1}, w)$) since w is the only demand in B_w.

Let μ, $1 \leq \mu \leq m$, be an iteration in which w is an outer demand and a new facility $\hat{w}_\mu$ opens. The algorithm incurs a cost of $f_{\hat{w}_\mu}$ in iteration μ. This cost cannot exceed w's potential. By Proposition 2, there can be at most one isolated active coalition K such that $d(\hat{w}_\mu, c_K) < d(F_{\mu-1}, c_K)$. If such an isolated active coalition K exists, we can apply Lemma 16 and obtain that $\Delta\Phi_K(\mu) < \frac{\lambda+1}{\lambda-2} d(F_{\mu-1}, w)$. For any other active coalition K', it must be $\Delta\Phi_{K'}(\mu) \leq 0$. In particular, if K' is a non-isolated active coalition, the claim follows from Lemma 14. If K' is an isolated active coalition such that $d(\hat{w}_\mu, c_{K'}) \geq d(F_{\mu-1}, c_{K'})$, then the potential function$\Phi_{K'}$ is not affected by $\hat{w}_\mu$ and $\Delta\Phi_{K'}(\mu) = 0$.

Therefore, for any iteration μ in which w is an outer demand and a new facility $\hat{w}_\mu$ opens,

$$f_{\hat{w}_\mu} + \Delta\Phi(\mu) \leq (1 + \tfrac{\lambda+1}{\lambda-2}) d(F_{\mu-1}, w) \leq \tfrac{2x+1}{x} d(F_{\mu-1}, w)$$

where the last inequality holds for any $\lambda \geq 3x + 6$.

If w is an outer demand at its assignment time, the value of the potential function does not change because no new facilities open and the demand w, which becomes unsatisfied, is an outer demand. Hence, $d(\hat{w}_m, w) + \Delta\Phi(\mu + 1) = d(F_m, w)$.

Therefore, for any iteration μ in which w is an outer demand, the algorithm's cost plus the increase in the potential function in iteration μ cannot exceed $\frac{2x+1}{x} d(F_{\mu-1}, w)$. Let μ^* be the first iteration (if any) in which w is an outer demand. Using the fact that w's potential drops by a factor of x in each iteration, we conclude that

$$\tfrac{2x+1}{x} \sum_{\mu=\mu^*}^{m+1} d(F_{\mu-1}, w) \leq \tfrac{2x+1}{x-1} d(F_{\mu^*-1}, w) \leq \tfrac{2x+1}{x-1} [(\lambda + 1)\delta^* + ((\lambda + 1)\rho + 1) d^*_w]$$

The last inequality follows from Ineq. (3) because w is an outer demand with respect to the facility configuration F_{μ^*-1}.

Putting everything together, we conclude that the total algorithm's cost plus the total increase in the potential function over all iterations in which w is an outer demand does not exceed $\frac{2x+1}{x-1}[(\lambda + 1)\delta^* + ((\lambda + 1)\rho + 1) d^*_w]$.

The lemma follows by combining the results of the analysis in case that w is an inner demand and w is an outer demand. $\square$

Concluding the Proof of Theorem 3. We have to bound the total credit allocated to each $K \in \mathcal{K}$ by the potential function Φ_K. We first observe that $\sum_{p=0}^{\zeta_K} f_p \leq 2 f_{\zeta_K}$ because the facility costs scale by a factor of 2. Therefore, the potential function component $\Xi^{(1)}_K$ allocates a credit of $2(\psi + 1) f_{\zeta_K}$ to each $K \in \mathcal{K}$. In addition, the components $\Xi^{(2)}_K$ and $\Xi^{(3)}_K$ allocate a credit of $\ln((1+\frac{1}{\lambda})\gamma^2) f_{\zeta_K}$ and $2 f_{\zeta_K}$ respectively. As for the potential function component $\Xi^{(4)}_K$, it allocates a credit of $2(\psi + 1)\ln(x + 1) f_{\zeta_K}$ because for each stage (j, p), $P_K(j, p) \leq (x + 1)\max\{P_K(j, p + 1), d(Q_p, c_K)\}$, by Ineq. (15). Hence, the total credit allocated to each $K \in \mathcal{K}$ by the potential function Φ_K is $[2(\psi+1)(\ln(x+1)+1)+\ln((1+\frac{1}{\lambda})\gamma^2)+2]\, f_{\zeta_K}$.

We recall that for each $K \in \mathcal{K}$, K's representative c_K is chosen as the cheapest optimal center in K. In other words, $\zeta_K = \min_{c_i \in K}\{\zeta_i\}$, where ζ_i denotes the type of the optimal center c_i. Hence, we can apply Proposition 6 and conclude that the total credit allocated by the potential function Φ is

$$\tfrac{2(3x+6)(\lambda+1)}{x(\lambda-2)} [2(\psi + 1)(\ln(x + 1) + 1) + \ln((1 + \tfrac{1}{\lambda})\gamma^2) + 2]\, \mathrm{Fac}^*$$

Furthermore, the optimal assignment cost of each demand is considered at most once either by Lemma 13 or by Lemma 18 (i.e., when the demand becomes unsatisfied) and at most once by Lemma 17 (i.e., when the demand becomes satisfied). Therefore, the total optimal assignment cost is charged at most $7\,[(\lambda + 1)\rho + \lambda + 3]$ times, where we have used $\frac{2x+1}{x-1} \leq \frac{2x+4}{x}$, for any $x \geq 4$, and $\frac{6x+10}{x} \leq 7$, for any $x \geq 10$.

Finally, we have to multiply the algorithm's cost by a factor of 2 because we have rounded down the actual facility costs to the nearest integral power of 2. Therefore, the total algorithm's cost cannot exceed

$$\tfrac{4(3x+6)(\lambda+1)}{x(\lambda-2)} [2(\psi + 1)(\ln(x + 1) + 1) + \ln((1 + \tfrac{1}{\lambda})\gamma^2) + 2]\, \mathrm{Fac}^* + 14\,[(\lambda + 1)\rho + \lambda + 3]\, \mathrm{Asg}^*$$

Since $x \geq 12$, $\lambda = 3x + 6$, and $\gamma = 8\lambda$ are constants, setting $\psi = \rho = \mathrm{O}(\frac{\log n}{\log \log n})$ yields the claimed competitive ratio.

5 Concluding Remarks

In this paper, we considered the problem of Online Facility Location and proved that its competitive ratio is $\Theta(\frac{\log n}{\log \log n})$. On the negative side, we presented a lower bound of $\Omega(\frac{\log n}{\log \log n})$ on the competitive ratio of any randomized algorithm against an oblivious adversary which holds even for uniform facility costs, instances whose optimal solution consists of a single facility, and very simple metric spaces. On the positive side, we formulated a deterministic algorithm and proved that its competitive ratio is $\mathrm{O}(\frac{\log n}{\log \log n})$ for every metric space and non-uniform facility costs. The analysis is based on a fundamental property of metric spaces, namely, that every metric space has a hierarchical decomposition such that each component either is well-separated or has a large diameter.

An interesting direction for further research is to determine whether there exists a deterministic algorithm which uses a small (e.g., poly-logarithmic) amount of memory and achieves an (asymptotically) optimal competitive ratio. Such an algorithm would be more efficient than the algorithm presented in this paper and could be used to obtain an efficient deterministic streaming algorithm for k-Median along the lines of [10].

To establish an (asymptotically) optimal competitive ratio, we let the online algorithm store the location of each unsatisfied demand and exploit this information to determine the location of a new facility. Hence, the algorithm requires a linear amount of memory in the worst case. On the other hand, we provide evidence that any (similar) algorithm which determines the location of the new facility without carefully considering the locations of previous demands cannot achieve a sub-logarithmic competitive ratio.

Recently, Anagnostopoulos et al. [2] presented a deterministic algorithm for d-dimensional Euclidean spaces which uses only $\mathrm{O}(2^d k \log n)$ memory per optimal center. However, the competitive ratio of this algorithm is $\Theta(2^d \log n)$ and its efficiency crucially depends on the dimension of the metric space. At the other end, the randomized algorithm of Meyerson [24] is memoryless, works for every metric space, and achieves an asymptotically optimal competitive ratio.

References

1. S. Albers and H. Koga. New Online Algorithms for the Page Replication Problem. *J. of Algorithms*, 27(1):75–96, 1998.
2. A. Anagnostopoulos, R. Bent, E. Upfal, and P. Van Hentenryck. A Simple and Deterministic Competitive Algorithm for Online Facility Location. Technical Report CS 03-16, Brown University, 2003.
3. V. Arya, N. Garg, R. Khandekar, A. Meyerson, K. Munagala, and V. Pandit. Local Search Heuristics for k-Median and Facility Location Problems. In *Proc. of the 33rd ACM Symp. on Theory of Computing (STOC '01)*, pp. 21–29, 2001.
4. B. Awerbuch, Y. Bartal, and A. Fiat. Competitive Distributed File Allocation. In *Proc. of the 25th ACM Symp. on Theory of Computing (STOC '93)*, pp. 164–173, 1993.
5. Y. Bartal. Probabilistic Approximations of Metric Spaces and its Algorithmic Applications. In *Proc. of the 37th IEEE Symp. on Foundations of Computer Science (FOCS '96)*, pp. 184–193, 1996.
6. Y. Bartal, A. Fiat, and Y. Rabani. Competitive Algorithms for Distributed Data Management. *J. of Computer and System Sciences*, 51(3):341–358, 1995.
7. A. Borodin and R. El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, 1998.
8. M. Charikar, C. Chekuri, T. Feder, and R. Motwani. Incremental Clustering and Dynamic Information Retrieval. In *Proc. of the 29th ACM Symp. on Theory of Computing (STOC '97)*, pp. 626–635, 1997.

9. M. Charikar and S. Guha. Improved Combinatorial Algorithms for the Facility Location and k-Median Problems. In *Proc. of the 40th IEEE Symp. on Foundations of Computer Science (FOCS '99)*, pp. 378–388, 1999.
10. M. Charikar, L. O'Callaghan, and R. Panigrahy. Better Streaming Algorithms for Clustering Problems. In *Proc. of the 35th ACM Symp. on Theory of Computing (STOC '03)*, pp. 30–39, 2003.
11. M. Charikar and R. Panigrahy. Clustering to Minimize the Sum of Cluster Diameters. In *Proc. of the 33rd ACM Symp. on Theory of Computing (STOC '01)*, pp. 1–10, 2001.
12. R. Fleischer and S. Seiden. New results for Online Page Replication. In *Proc. of APPROX '00*, LNCS 1913, pp. 144–154, 2000.
13. D. Fotakis. On the Competitive Ratio for Online Facility Location. In *Proc. of ICALP '03*, LNCS 2719, pp. 637–652, 2003.
14. D. Fotakis. Incremental Algorithms for Facility Location and k-Median. Manuscipt available from `http://users.auth.gr/~fotakis/data/incremen.pdf`, 2004.
15. S. Guha. *Approximation Algorithms for Facility Location Problems*. PhD thesis, Stanford University, 2000.
16. S. Guha and S. Khuller. Greedy Strikes Back: Improved Facility Location Algorithms. In *Proc. of the 9th ACM-SIAM Symposium on Discrete Algorithms (SODA '98)*, pp. 649–657, 1998.
17. M. Imase and B.M. Waxman. Dynamic Steiner Tree Problem. *SIAM J. on Discrete Mathematics*, 4(3):369–384, 1991.
18. K. Jain, M. Mahdian, and A. Saberi. A New Greedy Approach for Facility Location Problems. In *Proc. of the 34th ACM Symp. on Theory of Computing (STOC '02)*, pp. 731–740, 2002.
19. 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 Association for Computing Machinery*, 48(2):274–296, 2001.
20. M.R. Korupolu, C.G. Plaxton, and R. Rajaraman. Analysis of a Local Search Heuristic for Facility Location Problems. In *Proc. of the 9th ACM-SIAM Symposium on Discrete Algorithms (SODA '98)*, pp. 1–10, 1998.
21. C. Lund, N. Reingold, J. Westbrook, and D. Yan. Competitive Online Algorithms for Distributed Data Management. *SIAM J. on Computing*, 28(3):1086–1111, 1999.
22. M. Mahdian, Y. Ye, and J. Zhang. Improved Approximation Algorithms for Metric Facility Location Problems. In *Proc. of APPROX '02*, LNCS 2462, pp. 229–242, 2002.
23. R.R. Mettu and C.G. Plaxton. The Online Median Problem. In *Proc. of the 41st IEEE Symp. on Foundations of Computer Science (FOCS '00)*, pp. 339–348, 2000.
24. A. Meyerson. Online Facility Location. In *Proc. of the 42nd IEEE Symp. on Foundations of Computer Science (FOCS '01)*, pp. 426–431, 2001.
25. A. Meyerson, K. Munagala, and S. Plotkin. Designing Networks Incrementally. In *Proc. of the 42nd IEEE Symp. on Foundations of Computer Science (FOCS '01)*, pp. 406–415, 2001.
26. D. Shmoys. Approximation Algorithms for Facility Location Problems. In *Proc. of APPROX '00*, LNCS 1913, pp. 27–33, 2000.
27. D. Shmoys, E. Tardos, and K. Aardal. Approximation Algorithms for Facility Location Problems. In *Proc. of the 29th ACM Symp. on Theory of Computing (STOC '97)*, pp. 265–274, 1997.
28. M. Sviridenko. An Improved Approximation Algorithm for the Metric Uncapacitated Facility Location Problem. In *Proc. of IPCO '02*, LNCS 2337, pp. 240–257, 2002.

A Appendix

A.1 Opening the New Facility at an Arbitrary Location

In this section, we consider the variant of DFL which opens a new facility at the location of each new demand w with $\mathrm{Pot}(B_w) \geq f$ (see also Fig. 2). We call this algorithm Simple Facility Location – SFL. We prove that the competitive ratio of SFL is $\Omega(\log n)$. At an intuitive level, this provides evidence that a deterministic algorithm cannot achieve a sub-logarithmic competitive ratio unless it carefully takes into account the locations of the previous demands to select the location of a new facility.

Lemma 19. *For any $x > 1$, the competitive ratio of* SFL *is* $\Omega(\log n)$.

Proof. We consider a simple instance whose optimal solution consists of a single facility c. Let $m = \lceil \log_x((x-1)n) \rceil$. The metric space is a star consisting of the center c and m points $v_1, \ldots, v_m$. For each point v_i, the distance between v_i and c is $d(v_i, c) = \frac{f}{x^i}$.

The demand sequence is divided into m phases. Phase i, $i = 1, \ldots, m$, consists of x^{i-1} unit demands. Phase i starts with $x^{i-1} - 1$ unit demands located at c and ends with a single demand located at v_i. The total number of demands is at most $\frac{x^m - 1}{x-1} \leq n$. The optimal solution opens a facility at c and incurs a total cost no greater than $f + \sum_{i=1}^{m} \frac{1}{x^i} f \leq \frac{x}{x-1} f$.

Then, we consider the cost of SFL on this demand sequence. By induction on i, we show that (i) at the end of phase i, there are no unsatisfied demands (i.e., $L = \emptyset$), and (ii) while in phase i, the algorithm opens a single new facility at v_i when it considers the last demand of the phase.

The claim is true for the first phase, since SFL opens a new facility at v_1 when it considers the one and only demand of phase 1. We inductively assume that the claim remains true up to the end of phase i. By the inductive hypothesis, at the beginning of phase $i+1$, it is the case that $L = \emptyset$ and $d(F, c) = d(v_i, c) = \frac{f}{x^i}$. Hence, the first $x^i - 1$ demands of phase $i+1$ increase the potential of unsatisfied demands to $\frac{f}{x^i}(x^i - 1) < f$. Therefore, SFL must wait for the last demand of phase $i+1$ before it opens a new facility. The last demand of phase $i+1$ is located at v_{i+1}. Since $d(F, v_{i+1}) = d(v_i, v_{i+1}) > \frac{f}{x^i}$ and $d(v_{i+1}, c) = \frac{f}{x^{i+1}}$, $B_{v_{i+1}}$ includes c and thus every demand which arrives in phase $i+1$. Therefore, $\mathrm{Pot}(B_{v_{i+1}}) > f$ and the algorithm opens a new facility at v_{i+1}. Furthermore, the demands arriving in phase $i+1$ become satisfied and L becomes empty.

By the discussion above, the facility cost of SFL is exactly mf and the assignment cost of SFL is $\sum_{i=1}^{m-1} \frac{x^i - 1}{x^i} f \geq (m - \frac{x}{x-1}) f$. Therefore, the competitive ratio of SFL is at least $\frac{2(x-1)}{x} m - 1 = \frac{2(x-1)}{x} \lceil \log_x((x-1)n) \rceil - 1$, which is $\Omega(\log n)$ for every constant $x > 1$. $\square$

Remark. We focus on $x > 1$ because it comprises the most interesting case. A similar construction yields a lower bound of $\Omega(\log n)$ on the competitive ratio of SFL for any $x \in (0, 1]$.

A.2 The Proof of Lemma 2

Let M be a metric space, and let $D = D(M)$ denote the diameter of M. We start by hierarchically partitioning M into groups such that (i) the groups are well-separated from each other and (ii) no large diameter group contains any small diameter sub-groups which are well-separated from the remaining points in the group and can form a group on their own. Formally, for every integer $i \geq 0$, we partition M into *level i groups* $G_1^i, \ldots, G_m^i$ such that (i) for every $j_1 \neq j_2$, $d(G_{j_1}^i, G_{j_2}^i) > \frac{D}{4\gamma^i}$, and (ii) for every level

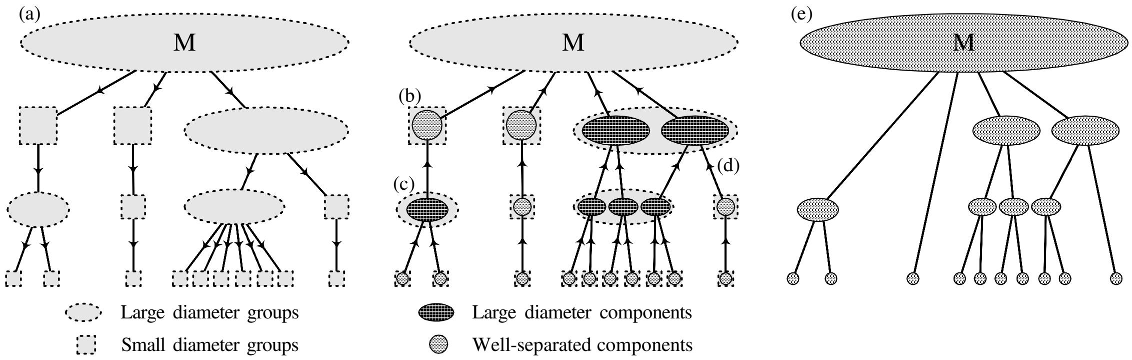

Fig. 7. The hierarchical decomposition of Lemma 2.

i group G^i_j, if $D(G^i_j) > \frac{D}{\gamma^i}$, then G^i_j does not contain any subsets $G \subseteq G^i_j$ such that both $D(G) \le \frac{D}{\gamma^{i+1}}$ and $d(G, G^i_j \setminus G) > \frac{D}{4\gamma^i}$. Since level i groups form a partitioning of M, for every level i group G^i_j, $\mathrm{sep}_M(G^i_j) > \frac{D}{4\gamma^i}$ (recall that $\mathrm{sep}_M(G^i_j) = d(G^i_j, M \setminus G^i_j)$).

We inductively prove that a simple greedy procedure which introduces new groups at the next level as long as condition (ii) is violated results in a partitioning with the desired properties. To begin with, M is the only level 0 group. Given the partitioning at the i-th level, the partitioning at the next level is constructed as follows: For each group G^i_j of $D(G^i_j) > \frac{D}{\gamma^{i+1}}$ (Fig. 7.(a), large diameter groups), we set $\tilde{G}^i_j \leftarrow G^i_j$. While $\tilde{G}^i_j$ violates (ii) for the $(i+1)$-th level (i.e., while $D(\tilde{G}^i_j) > \frac{D}{\gamma^{i+1}}$ and there exists a sub-group $G \subseteq \tilde{G}^i_j$ such that $D(G) \le \frac{D}{\gamma^{i+2}}$ and $d(G, \tilde{G}^i_j \setminus G) > \frac{D}{4\gamma^{i+1}}$), we create a new level $i+1$ group for G, remove G from $\tilde{G}^i_j$ ($\tilde{G}^i_j \leftarrow \tilde{G}^i_j \setminus G$), and iterate. We also create a new level $i+1$ group consisting of the points remaining in $\tilde{G}^i_j$ after the loop has terminated. The level $i+1$ groups created from G^i_j form a partitioning of it and comprise G^i_j's children. If $D(G^i_j) \le \frac{D}{\gamma^{i+1}}$, G^i_j has a single child at level $i+1$ which is identical to it (Fig. 7.(a), small diameter groups).

It is straight-forward that the above procedure results in a partitioning into level $i+1$ groups which fulfill condition (ii). Moreover, since condition (i) holds for level i, the children of different level i groups are also well-separated from each other. Formally, since for every $j_1 \ne j_2$, $d(G^i_{j_1}, G^i_{j_2}) > \frac{D}{4\gamma^i}$, each child of $G^i_{j_1}$ must be at distance greater than $\frac{D}{4\gamma^i} > \frac{D}{4\gamma^{i+1}}$ from any child of $G^i_{j_2}$. As for the distance between different children of G^i_j, when a subset G is removed, the points in G are at distance greater than $\frac{D}{4\gamma^{i+1}}$ from the points remaining in $\tilde{G}^i_j$. For every G' removed from G^i_j before G, it must be $d(G', G) > \frac{D}{4\gamma^{i+1}}$, because the distance between G' and G was considered when G' was removed.

Since the quantity $\frac{D}{\gamma^i}$ is strictly decreasing with i, we eventually reach a level ν such that every level ν group consists of a single point.

Level i groups are further partitioned into *level i components* $K^i_1, \dots, K^i_m$ such that for every j, (i) $D(K^i_j) \le \frac{D}{\gamma^i}$, and (ii) either $D(K^i_j) > \frac{D}{\gamma^{i+1}}$ or $\mathrm{sep}_M(K^i_j) > \frac{D}{4\gamma^i}$. By the definition of level i groups, only large diameter groups (namely, groups of diameter greater than $\frac{D}{\gamma^i}$) should be partitioned. To end up with a hierarchical decomposition, we proceed inductively in a bottom-up fashion where the level $i+1$ components are merged with each other to form the level i components.

Each level ν group consists of a single point and is a small diameter group. Thus, we create a single level ν component for each level ν group. We inductively assume the collection $\mathcal{K}^{i+1}_j = \{K_1, \dots, K_m\}$

consisting of the level $i+1$ components the children of the level i group G_j^i are partitioned in. $\mathcal{K}_j^{i+1}$ is a partitioning of G_j^i and for each $K \in \mathcal{K}_j^{i+1}$, $D(K) \leq \frac{D}{\gamma^{i+1}}$ because K is a level $i+1$ component.

If G_j^i is a small diameter level i group (i.e., $D(G_j^i) \leq \frac{D}{\gamma^i}$), we create a single level i component for it (Fig. 7, (b) and (c)). We recall that $\mathrm{sep}_M(G_j^i) > \frac{D}{4\gamma^i}$. If G_j^i is a large diameter level i group (i.e., $D(G_j^i) > \frac{D}{\gamma^i}$), we show that G_j^i can be partitioned into level i components of diameter in $(\frac{D}{\gamma^{i+1}}, \frac{D}{\gamma^i}]$ which are obtained by merging the level $i+1$ components in $\mathcal{K}_j^{i+1}$ (Fig. 7, (d)). Intuitively, this is true because G_j^i does not contain any small diameter sub-groups which are well-separated from each other. Formally, for each $G \subseteq G_j^i$ either $D(G) > \frac{D}{\gamma^{i+1}}$ or $d(G, G_j^i \setminus G) > \frac{D}{4\gamma^i}$. As a result, for each $K \in \mathcal{K}_j^{i+1}$, there exists a $K' \in \mathcal{K}_j^{i+1}$ such that $d(K, K') \leq \frac{D}{4\gamma^i}$.

To merge the level $i+1$ components in $\mathcal{K}_j^{i+1}$, we maintain two disjoint collections Z_S and Z_L. Z_S contains the components of diameter no greater than $\frac{D}{\gamma^{i+1}}$ and Z_L contains the components of diameter greater than $\frac{D}{\gamma^{i+1}}$ obtained by merging some of the components in Z_S. Initially, $Z_S = \mathcal{K}_j^{i+1}$.

While there exist $K_1, K_2 \in Z_S$ such that $d(K_1, K_2) \leq \frac{D}{4\gamma^i}$, K_1 and K_2 are removed from Z_S and merged into a new component $K = K_1 \cup K_2$. If $D(K) \leq \frac{D}{\gamma^{i+1}}$, K is put in Z_S. Otherwise, K is put in Z_L. Due to the choice of K_1 and K_2, for every $K \in Z_L$, $D(K) \leq \frac{D}{\gamma^i}(\frac{1}{4} + \frac{2}{\gamma})$.

The merge procedure above cannot terminate with an empty collection Z_L. If Z_L was empty, Z_S would contain more than one components, because $D(G_j^i) > \frac{D}{\gamma^i}$ and for every $K \in Z_S$, $D(K) \leq \frac{D}{\gamma^{i+1}}$. In addition, for every $K_1, K_2 \in Z_S$, it must be $d(K_1, K_2) > \frac{D}{4\gamma^i}$, otherwise the merge procedure would not have terminated. Hence, G_j^i would contain a subset $K \subseteq G_j^i$ such that both $D(K) \leq \frac{D}{\gamma^{i+1}}$ and $d(K, G_j^i \setminus K) > \frac{D}{4\gamma^i}$, which contradicts to the hypothesis that G_j^i is a level i group.

Therefore, the merge procedure always terminates with a non-empty collection Z_L. If Z_S is non-empty, every $K \in Z_S$ is associated with a component $K' \in Z_L$ such that $d(K, K') \leq \frac{D}{4\gamma^i}$. For every $K \in Z_S$, such a $K' \in Z_L$ must exist because otherwise, it would be both $D(K) \leq \frac{D}{\gamma^{i+1}}$ and $d(K, G_j^i \setminus K) > \frac{D}{4\gamma^i}$, which is a contradiction.

Each component in Z_L is merged with the components of Z_S associated with it. The resulting components have a diameter greater than $\frac{D}{\gamma^{i+1}}$ because each of them includes a single component of Z_L. In addition, their diameter cannot exceed $\frac{D}{\gamma^{i+1}} + \frac{D}{4\gamma^i} + \frac{D}{\gamma^i}(\frac{1}{4} + \frac{2}{\gamma}) + \frac{D}{4\gamma^i} + \frac{D}{\gamma^{i+1}} = \frac{D}{\gamma^i}(\frac{3}{4} + \frac{4}{\gamma})$, because the diameter of the components in Z_S is at most $\frac{D}{\gamma^{i+1}}$, the diameter of each component in Z_L is at most $\frac{D}{\gamma^i}(\frac{1}{4} + \frac{2}{\gamma})$, and every component of Z_S has been associated with a component of Z_L at distance no greater than $\frac{D}{4\gamma^i}$. For $\gamma \geq 16$, $\frac{D}{\gamma^i}(\frac{3}{4} + \frac{4}{\gamma})$ is no greater than $\frac{D}{\gamma^i}$.

By eliminating multiple occurrences of the same component at different levels (Fig. 7, (e)), we obtain a hierarchical decomposition/complete laminar set system on M. It remains to establish that this decomposition has the desired properties. Let K' be any component different from the root and let $i+1$, $i \geq 0$, be the first level (namely, the level with the smallest index) at which K' appears before the elimination of any multiple occurrences of it. Since K' is a level $i+1$ component, either $D(K') > \frac{D}{\gamma^{i+2}}$ or $\mathrm{sep}_M(K') > \frac{D}{4\gamma^{i+1}}$. To conclude the proof, we show that the diameter of the parent of K' in the hierarchical decomposition is at most $\frac{D}{\gamma^i}$. Since K' appears as a level $i+1$ component, but it does not appear as a level i component, there exists a level i component K such that $K' \subset K$ and K' is a child of K in the hierarchical decomposition. It must be $D(K) \leq \frac{D}{\gamma^i}$ because K is a level i component. $\square$

A.3 The Proof of Proposition 6

Proof. We work with the decomposition tree $T_{\mathcal{K}}$ representing the hierarchical decomposition $\mathcal{K}$ (see also the proof of Lemma 1). For a component/node $K \in T_{\mathcal{K}}$, we let $T_{\mathcal{K}}(K)$ denote the subtree of $T_{\mathcal{K}}$ rooted at K. In addition, we let $f_{\max}(K) = \max_{c_i \in K}\{f_{\zeta_i}\}$ denote the cost of the most expensive center in K, and we let $\mathrm{Fac}^*(K) = \sum_{c_i \in K} f_{\zeta_i}$ denote the total cost of the centers in K. We also recall that $f_{\zeta_K} = \min_{c_i \in K}\{f_{\zeta_i}\}$ denotes the cost of the cheapest center in K.

We prove the following stronger claim: For any component/node $K \in T_{\mathcal{K}}$,

$$\sum_{K' \in T_{\mathcal{K}}(K)} f_{\zeta_{K'}} \leq 2\,\mathrm{Fac}^*(K) - f_{\max}(K) \tag{21}$$

The proposition follows by applying the claim for $K = F^*$ because $T_{\mathcal{K}}(F^*) = T_{\mathcal{K}}$ and $\mathrm{Fac}^*(F^*) = \mathrm{Fac}^*$.

We prove Ineq. (21) by an inductive argument proceeding from the leaves of $T_{\mathcal{K}}$ to the root. The leaves of $T_{\mathcal{K}}$ consist of a single center each. Hence, Ineq. (21) holds if K is a leaf of $T_{\mathcal{K}}$ because K consists of a single center and $f_{\zeta_K} = \mathrm{Fac}^*(K) = f_{\max}(K)$.

Let K be any non-leaf component/node of the decomposition tree $T_{\mathcal{K}}$. Since $\mathcal{K}$ is a hierarchical decomposition of F^*, each non-leaf component/node of $T_{\mathcal{K}}$ must have at least two children. Hence, let $K_1, \dots, K_m$, $m \geq 2$, be the children of K in $T_{\mathcal{K}}$ indexed in non-decreasing order of their $f_{\max}$ values, namely, $f_{\max}(K_1) \leq \dots \leq f_{\max}(K_m)$. Thus,

$$\begin{aligned}
\sum_{K' \in T_{\mathcal{K}}(K)} f_{\zeta_{K'}} &= f_{\zeta_K} + \sum_{j=1}^{m} \sum_{K' \in T_{\mathcal{K}}(K_j)} f_{\zeta_{K'}} \\
&\leq f_{\zeta_K} + \sum_{j=1}^{m} [2\,\mathrm{Fac}^*(K_j) - f_{\max}(K_j)] \\
&= 2\,\mathrm{Fac}^*(K) + f_{\zeta_K} - \sum_{j=1}^{m} f_{\max}(K_j)
\end{aligned} \tag{22}$$

The second inequality follows from the inductive hypothesis. The last equality holds because the components $K_1, \dots, K_m$ form a partitioning of K.

Since the components $K_1, \dots, K_m$ form a partitioning of K, it must be $f_{\zeta_K} \leq f_{\max}(K_1)$. In addition, $f_{\max}(K) = f_{\max}(K_m)$, since the children of K are indexed in non-decreasing order of their $f_{\max}$ values. Combining these observations with Ineq. (22), we conclude that

$$\sum_{K' \in T_{\mathcal{K}}(K)} f_{\zeta_{K'}} \leq 2\,\mathrm{Fac}^*(K) + f_{\zeta_K} - f_{\max}(K_1) - f_{\max}(K_m) \leq 2\,\mathrm{Fac}^*(K) - f_{\max}(K)$$

In the first inequality, we have also used the fact that K has at least two children (i.e., $m \geq 2$) and K_1 is different from K_m. $\square$