Incremental Algorithms for Facility Location and k-Median*

Dimitris Fotakis

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

Abstract. In the incremental versions of Facility Location and k-Median, the demand points arrive one at a time and the algorithm must maintain a good solution by either adding each new demand to an existing cluster or placing it in a new singleton cluster. The algorithm can also merge some of the existing clusters at any point in time. We present the first incremental algorithm for Facility Location which achieves a constant performance ratio and the first incremental algorithm for k-Median which achieves a constant performance ratio using $\mathrm{O}(k)$ medians, thus resolving an open question of [7]. The algorithm is based on a novel merge rule which ensures that the algorithm's configuration monotonically converges to the optimal facility locations according to a certain notion of distance. Using this property, we reduce the general case to the special case that the optimal solution consists of a single facility.

1 Introduction

The model of incremental algorithms for data clustering is motivated by practical applications where the demand sequence is not known in advance and the algorithm must maintain a consistently good clustering using a restricted set of operations which result in a solution of hierarchical structure. The framework of incremental clustering was introduced by Charikar et al. [4]. In this paper, we consider the incremental versions of metric Facility Location and k-Median. The problems of Facility Location and k-Median find many applications in the areas of network design and data clustering and have been the subject of intensive research over the last decade (e.g., [21] for a survey and [10] for approximation algorithms and applications). In addition to the offline setting, there are many applications where the demand points arrive online and the solution must be constructed incrementally using no information about future demands (e.g., [19]).

In *Incremental k-Median* [7], the demand points arrive one at a time. Each new demand must be either added to an existing cluster or placed in a new singleton cluster upon arrival. At any point in time, the algorithm can also merge some of the existing clusters. Each cluster is represented by its median whose location is determined at the cluster's creation time. When some clusters are merged with each other, the median of the new cluster must be selected among the medians of its components. The goal is to maintain a solution consisting of at most k clusters/medians which minimize the total assignment cost of the demands considered so far. The assignment cost of a demand is its distance from the median of the cluster the demand is currently included in.

The definition of *Incremental Facility Location* is similar. Demand points arrive one at a time and must be assigned to either an existing or a new facility upon arrival. At any point in time, the algorithm can also merge a facility with another one by closing the first facility and re-assigning all the demands currently assigned to it to the second facility. The objective is to maintain a solution which minimizes the sum of facility and assignment costs. As before, the assignment cost of a demand is its distance from the facility the demand is currently assigned to.

We evaluate the performance of incremental algorithms using the *performance ratio* [4]. An incremental algorithm achieves a performance ratio of c if for all demand sequences, the cost incurred by the

* This work was partially supported by the Future and Emerging Technologies programme of the EU under contract numbers IST-1999-14186 (ALCOM-FT) and IST-2001-33116 (FLAGS). Part of this work was done while the author was at the Max-Planck-Institut für Informatik, Saarbrücken, Germany.

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.

Comparison to Online and Streaming Algorithms. Similarly to online algorithms, incremental algorithms commit themselves to irrevocable decisions made without any knowledge about future demands. More specifically, when a new demand arrives, the algorithm may decide to add the demand to an existing cluster or merge some clusters with each other. These decisions are irrevocable because once formed, clusters cannot be broken up. In addition, the definition of the performance ratio is essentially identical to the definition of the competitive ratio (e.g., [3]). However, we have avoided casting Incremental k-Median as "Online k-Median". The most important reason is that we are not aware of any simple and natural notion of irrevocable cost which could be associated with the irrevocable decision that a demand is clustered together with some other demands.

Incremental algorithms also bear a resemblance to one-pass streaming algorithms for clustering problems (e.g., see [13] for a formulation of the streaming model and [12, 6] for applications to k-Median). However, in case of streaming algorithms, the emphasis is on space and time efficient algorithms which achieve a small approximation ratio by ideally performing a single scan over the input data. A streaming algorithm for k-Median is not restricted in terms of the solution's structure or the set of operations available. On the other hand, incremental algorithms must maintain a good hierarchical clustering by making irrevocable decisions. As for time and space efficiency, we only require explicitly that incremental algorithms should run in polynomial time. Nevertheless, all known incremental algorithms for clustering problems can be either directly regarded as or easily transformed to time and space efficient one-pass streaming algorithms (e.g., [4, 12, 7, 6]).

Previous Work. Charikar et al. [4] introduced the framework of incremental clustering and presented incremental algorithms for k-Center (i.e., minimize the maximum cluster radius) which achieve a constant performance ratio using k clusters. Charikar and Panigrahy [7] presented an incremental algorithm for Sum k-Radius (i.e., minimize the sum of cluster radii) which achieves a constant performance ratio using $\mathrm{O}(k)$ clusters.

The incremental version of k-Median was first considered by Charikar and Panigrahy [7], where it is shown that no deterministic algorithm which maintains at most k clusters can achieve a performance ratio better than $\Omega(k)$. Hence, we relax the requirement on the number of clusters allowing the algorithm to maintain $\mathrm{O}(k)$ clusters. Determining whether there exists an incremental algorithm for k-Median which achieves a constant performance ratio using $\mathrm{O}(k)$ medians is suggested as an open problem in [7].

The only known incremental algorithms for k-Median are the one-pass streaming algorithms of [12] and [6]. More specifically, the streaming algorithms of Guha et al. [12] can be regarded as incremental algorithms under the assumption that the number of demands n is known in advance. For k much smaller than n^ϵ, their algorithms achieve a performance ratio of $2^{\mathrm{O}(1/\epsilon)}$ using n^ϵ medians and run in $\mathrm{O}(nk\,\mathrm{poly}(\log n))$ time and n^ϵ space. The best known streaming algorithm for k-Median is the one-pass algorithm of Charikar et al. [6]. Under the assumption that n is known in advance, this algorithm can be easily transformed to an incremental algorithm which achieves a constant performance ratio with high probability (whp.[1]) using $\mathrm{O}(k \log^2 n)$ medians and runs in $\mathrm{O}(nk \log^2 n)$ time and $\mathrm{O}(k \log^2 n)$ space.

The only known incremental algorithms for Facility Location are the online algorithms of [19, 8, 1]. Meyerson [19] was the first to consider the online version of Facility Location, where the demand points arrive one at a time and must be irrevocably assigned to either an existing or a new facility upon arrival. In [19], a randomized $\mathrm{O}(\frac{\log n}{\log\log n})$-competitive algorithm and a lower bound of $\omega(1)$ are presented. In [8], the lower bound is improved to $\Omega(\frac{\log n}{\log\log n})$ and a deterministic $\mathrm{O}(\frac{\log n}{\log\log n})$-competitive algorithm is given. In [1], it is presented a simpler deterministic $\mathrm{O}(2^d \log n)$-competitive algorithm for d-dimensional Euclidean spaces.

[1] Throughout this paper, "whp." means "with probability at least $1 - \mathrm{O}(1/n)$".

The lower bounds of [19, 8] hold only if the decision of opening a facility at a particular location is irrevocable. Hence, they do not apply to the incremental version of Facility Location. However, the lower bound of [8] implies that every algorithm which maintains $\mathrm{o}(k \log n)$ facilities must incur a total *initial* assignment cost of $\omega(1)$ times the optimal cost, where the initial assignment cost of a demand is its distance from the first facility the demand is assigned to. Therefore, every algorithm treating merge as a black-box operation cannot approximate the optimal assignment cost within a constant factor unless it uses $\Omega(k \log n)$ facilities (e.g., the algorithm of [6]). In other words, to establish a constant performance ratio, one must prove that merge operations can also *decrease* the algorithm's assignment cost.

Related Work on Facility Location and k-Median. In the offline case, where the demand set is fully known in advance, there are constant factor approximation algorithms for Facility Location based on Linear Programming rounding (e.g., [22, 23]), local search (e.g., [16, 5, 2]), and the primal-dual method (e.g., [15, 14]). The best known polynomial-time algorithm achieves an approximation ratio of 1.52 [17], while no polynomial-time algorithm can achieve an approximation ratio less than 1.463 unless $\mathrm{NP} = \mathrm{DTIME}(n^{\mathrm{O}(\log \log n)})$ [11]. For k-Median, the best known polynomial-time algorithm achieves an approximation ratio of $3 + \mathrm{o}(1)$ [2], while no polynomial-time algorithm can achieve an approximation ratio less than $1 + 2/e$ unless $\mathrm{NP} = \mathrm{DTIME}(n^{\mathrm{O}(\log \log n)})$ [14]. As it is also observed in [19], our setting should not be confused with the setting of [18, 20], where the demand set is fully known in advance and the number of facilities/medians increases online.

Contribution. We present the first incremental algorithm for metric Facility Location which achieves a constant performance ratio. The algorithm combines a simple rule for opening new facilities with a novel merge rule based on distance instead of cost considerations. We use a new technique to prove that a case similar to the special case where the optimal solution consists of a single facility is the dominating case in the analysis. This technique is also implicit in [8] and may find applications to other online problems. To overcome the limitation imposed by the lower bound of [8], we also establish that in the dominating case, merge operations decrease the total assignment cost.

Using the algorithm for Facility Location as a building block, we obtain the first incremental algorithm for k-Median which achieves a constant performance ratio using $\mathrm{O}(k)$ medians, thus resolving the open question of [7]. Our algorithm is deterministic, runs in $\mathrm{O}(n^2 k)$ time and $\mathrm{O}(n)$ space, and is the first incremental algorithm for k-Median which does not assume any advance knowledge of n. Combining our techniques with the techniques of [6], we obtain a randomized incremental algorithm which achieves a constant performance ratio whp. using $\mathrm{O}(k)$ medians and runs in $\mathrm{O}(nk^2 \log^2 n)$ time and $\mathrm{O}(k^2 \log^2 n)$ space. This algorithm can also be regarded as an one-pass streaming algorithm for k-Median. Similarly to the algorithms of [12, 6], the randomized version of our algorithm assumes that a constant factor approximation to $\log n$ is known in advance.

Notation. We only consider unit demands by allowing multiple demands to be located at the same point. We always use n to denote the total number of demands. For Incremental Facility Location, we restrict our attention to the special case of uniform facility costs, where the cost of opening a facility, denoted by f, is the same for all points. We also use the terms facility, median and cluster interchangeably.

A metric space $\mathcal{M} = (M, d)$ is usually identified by its point set M. The distance function d is non-negative, symmetric, and satisfies the triangle inequality. 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 the nearest point in 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 a point in M' and a point in M''. For a subspace $M' \subseteq M$, $\mathrm{sep}(M') = d(M', M \setminus M')$ denotes the distance separating the points in M' from the points not in M'. It is $\mathrm{sep}(\emptyset) = \mathrm{sep}(M) = \infty$. For a point $u \in M$ and a non-negative number r, $\mathrm{Ball}(u, r)$ denotes the ball of center u and radius r, $\mathrm{Ball}(u, r) = \{v \in M : d(u, v) \leq r\}$.

```
Let x, β, and ψ be appropriately chosen constants
F ← ∅; L ← ∅; /* Initialization */
For each new demand u:
    L ← L ∪ {u}; r_u ← d(F,u)/x;
    B_u ← Ball(u, r_u) ∩ L;
    Pot(B_u) = Σ_{v∈B_u} d(F, v);
    if Pot(B_u) ≥ βf then
        Let w' be the location of u;
        open(w'); L ← L \ B_u;
        for each w ∈ F \ {w'} do
            if d(w, w') ≤ m(w) then
                merge(w → w');
    Let w be the facility in F closest to u;
    update_merge_radius(m(w));
    initial_assignment(u, w);

open(w')
    F ← F ∪ {w'}; Init(w') ← ∅;
    C(w') ← ∅; m^(1)(w') ← 3 r_u;
merge(w → w')
    F ← F \ {w}; C(w') ← C(w') ∪ C(w);
update_merge_radius(m(w))
    m^(2)(w) =
max{r : |Ball(w, r/ψ) ∩ (Init(w) ∪ {u})| · r ≤ βf};
    m(w) = min{m^(1)(w), m^(2)(w)};
initial_assignment(u, w)
    Init(w) ← Init(w) ∪ {u}; C(w) ← C(w) ∪ {u};
```

Fig. 1. The algorithm Incremental Facility Location – IFL .

2 An Incremental Algorithm for Facility Location

The algorithm Incremental Facility Location - IFL (Fig. 1) maintains its *facility configuration* F, its *merge configuration* consisting of a *merge ball* $\mathrm{Ball}(w, m(w))$ for each facility $w \in F$, and the set L of *unsatisfied demands*.

We use a simpler version of the deterministic algorithm of [8] for opening new facilities. The notion of unsatisfied demands (the set L) ensures that each demand contributes to the facility cost at most once. 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 neighborhood B_u of a new demand u has accumulated a potential of βf, a new facility located at the same point with u opens. Then, the unsatisfied demands in B_u lose their potential, become satisfied, and are removed from L.

Each facility $w \in F$ maintains the set $C(w)$ of the demands currently assigned to w and the set $\mathrm{Init}(w) \subseteq C(w)$ of the demands *initially assigned* to w. The demands in $\mathrm{Init}(w)$ are assigned to w when they arrive, while the demands in $C(w) \setminus \mathrm{Init}(w)$ have been initially assigned to a facility different from w. Each facility $w \in F$ also maintains its *merge radius* $m(w)$ and the corresponding *merge ball* $\mathrm{Ball}(w, m(w))$. The algorithm ensures that w is the only facility in its merge ball. When w opens, the merge radius of w is initialized to a fraction of the distance between w and the nearest existing facility. Then, if a new facility w' is included in w's merge ball, w is merged with w'. Namely, w is closed and removed from F, and every demand currently assigned to w is re-assigned to w'. The algorithm keeps decreasing $m(w)$ to ensure that no merge operation can dramatically increase the total assignment cost of the demands in $\mathrm{Init}(w)$. More specifically, the algorithm maintains the invariant that

$$|\mathrm{Init}(w) \cap \mathrm{Ball}(w, \tfrac{m(w)}{\psi})| \cdot m(w) \leq \beta\, f \tag{1}$$

After the algorithm has updated its configuration, it initially assigns the new demand to the nearest facility. We always distinguish between the arrival and the assignment time of a demand because the algorithm's configuration may have changed in between.

If the demands considered by IFL occupy m different locations, a crude analysis shows that IFL can be implemented in $\mathrm{O}(nm|F_{\max}|)$ time and $\mathrm{O}(\min\{n, m|F_{\max}|\})$ space, where $|F_{\max}|$ is the maximum number of facilities in F at any point in time. The remaining of this section is devoted to the proof of the following theorem.

Theorem 1. *For every* $x \geq 18$, $\beta \geq \frac{4(x+1)}{x-8}$, *and* $\psi \in [\max\{\frac{6\beta}{2\beta-3}, 4\}, 5]$, IFL *achieves a constant performance ratio.*

Preliminaries. For an arbitrary fixed sequence of demands, we compare the algorithm's cost with the cost of a fixed add-optimal facility configuration[2]. We denote this solution by F^* and refer to it as *the optimal solution*. To avoid confusing the algorithm's facilities with the 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 $c_1, c_2, \dots, c_k$. Each demand is assigned to the nearest center in F^*. For a demand u, c_u denotes the optimal center u is assigned to. We use the clustering induced by F^* to map the demands and the algorithm's facilities to optimal centers. In particular, a demand u is always mapped to c_u, i.e., the optimal center u is assigned to. Similarly, a facility w is mapped to the nearest optimal center denoted by c_w. Also, let $d^*_u = d(c_u, u) = d(F^*, u)$ denote the optimal assignment cost of u, let $\mathrm{Fac}^* = kf$ be the optimal facility cost, and let $\mathrm{Asg}^* = \sum_{i=1}^{n} d^*_u$ be the optimal assignment cost.

In addition to x, β, and ψ, let $\lambda = 3x + 2$, $\rho = (\psi + 2)(\lambda + 2)$, and $\gamma = 12\rho$ be also constants. Let also $u_1, \dots, u_n$ be the demand sequence considered by IFL. We show that after the demand u_j has been considered, $1 \le j \le n$, the facility cost of IFL does not exceed $a_1\mathrm{Fac}^* + b_1\mathrm{Asg}^*_j$ and the assignment cost of IFL does not exceed $a_2\mathrm{Fac}^* + b_2\mathrm{Asg}^*_j$, where $\mathrm{Asg}^*_j = \sum_{i=1}^{j} d^*_{u_i}$, and $a_1 = 1$, $a_2 = 2\beta\ln(3\gamma^2)(5(\psi+4)\gamma^2+3)$, $b_1 = \frac{3x}{\beta}$, and $b_2 = 4((\rho+1)\gamma^2+2)+14x$. With a more careful analysis, we can improve a_2 and b_2 to $a'_2 = 4\beta\log(\gamma)(12(\psi+2)+3)$ and $b'_2 = (\lambda+2)(8\psi+25)$. Moreover, we can remove the assumption that F^* is add-optimal by replacing the bound on the algorithm's assignment cost with $\max\{a_2, b_2\}(\mathrm{Fac}^* + \mathrm{Asg}^*_j)$ (see also Section A.9, in the Appendix).

Every time we want to explicitly refer to the algorithm's configuration (or some function of it) at the moment a demand is considered/facility opens, we use the demand's/facility's identifier as a subscript. Moreover, we use the convention that the algorithm first updates its configuration and then performs the demand's initial assignment. Hence, we distinguish between the algorithm's configuration at the demand's arrival and assignment times using plain symbols to refer to the former and primed symbols to refer to the latter time. For example, for a demand u, F_u/F'_u is the facility configuration at u's arrival/assignment time. Similarly, for a facility w, F_w/F'_w is the facility configuration just before/after w opens. Saying that an existing facility w is merged with a new facility w', we mean that the existing facility w is closed and the demands currently assigned to w are re-assigned to the new facility w' (and not the other way around). We proceed to establish the basic properties of IFL.

Lemma 1. *Let $\beta \ge \frac{4(x+1)}{x-8}$. Then, for every facility w mapped to c_w, $d(c_w, w) \le \frac{d(F_w, c_w)}{3}$.*

Proof Sketch. To reach a contradiction, let us assume that w is a facility such that $d(c_w, w) > \frac{d(F_w, c_w)}{3}$. Since $B_w \subseteq \mathrm{Ball}(w, \frac{d(F_w, w)}{x})$, we can show that for each $u \in B_w$, $d(u, w) < \frac{4}{x-4}\, d^*_u$ and $d(F_w, u) < \frac{4(x+1)}{x-4}\, d^*_u$. Using $\mathrm{Pot}(B_w) \ge \beta f$, we conclude that for every $\beta \ge \frac{4(x+1)}{x-8}$, $f + \sum_{u \in B_w} d(u, w) < \sum_{u \in B_w} d^*_u$, which contradicts to the add-optimality of F^*. The full proof can be found in the Appendix, Section A.1. □

Proposition 1. *For every facility w, there will always exist a facility in $\mathrm{Ball}(w, \frac{x}{x-3}\, m(w))$ and each demand currently assigned to w will remain assigned to a facility in $\mathrm{Ball}(w, \frac{x}{x-3}\, m(w))$.*

Proof. The proposition is true as long as w remains open. If w is merged with a new facility w', we inductively assume that the proposition is true for w'. Then, the proposition follows from the observation that $\mathrm{Ball}(w', \frac{x}{x-3}\, m(w'))$ is included in $\mathrm{Ball}(w, \frac{x}{x-3}\, m(w))$ (see also Proposition 2, Section A.1). □

Facility Cost. It is not difficult to prove that in contrast to the online algorithms for Facility Location [19, 8, 1], IFL does not suffer from facility proliferation. We distinguish between *supported* facilities,

[2] A facility configuration F is *add-optimal* if its total cost cannot decrease by adding a new facility to F. Formally, for every w, $f + \sum_u d(F \cup \{w\}, u) \ge \sum_u d(F, u)$.

whose opening cost can be charged to the optimal assignment cost, and *unsupported* facilities. A facility w is *supported* if $\mathrm{Asg}^*(B_w) = \sum_{u \in B_w} d_u^* \geq \frac{\beta}{3x} f$, and *unsupported* otherwise. Since each demand contributes to the facility cost at most once, the total cost of supported facilities is at most $\frac{3x}{\beta} \mathrm{Asg}^*$. Next, we prove that there always exists at most one unsupported facility mapped to each optimal center. Therefore, the algorithm's facility cost does not exceed $\mathrm{Fac}^* + \frac{3x}{\beta} \mathrm{Asg}^*$.

Lemma 2. *Let w be an unsupported facility mapped to an optimal center c_w, and let w' be a new facility also mapped to c_w. If w' opens while w is still open, then w is merged with w'.*

Proof. By Lemma 1, it must be $d(c_w, w') \leq \frac{1}{3} d(F_{w'}, c_w) \leq \frac{1}{3} d(c_w, w)$, because w' is mapped to c_w and $w \in F_{w'}$ by hypothesis. We also prove that $m(w) \geq \frac{3}{2} d(c_w, w)$ which implies the lemma because $d(w, w') \leq d(c_w, w) + d(c_w, w') \leq \frac{4}{3} d(c_w, w) \leq m(w)$ and w must be merged with w'.

To prove that $m^{(1)}(w) \geq \frac{3}{2} d(c_w, w)$, we show that for every unsupported facility w, $d(c_w, w) < \frac{4\, d(F_w, c_w)}{3(x-4)}$ (Proposition 3). We first observe that there must be a demand $u \in B_w$ such that $d_u^* < \frac{1}{3x} d(F_w, u)$, because w is an unsupported facility. Then, the claim follows from the fact that the radius of B_w is $\frac{d(F_w, w)}{x}$. Since $m^{(1)}(w) = 3 \frac{d(F_w, w)}{x}$, we obtain that $m^{(1)}(w) \geq \frac{3}{2} d(c_w, w)$ (Proposition 4, $x \geq 16$). The prove that $m^{(2)}(w) \geq \frac{3}{2} d(c_w, w)$ (Proposition 5), we first observe that $m^{(2)}(w)$ cannot become smaller than $\frac{3}{2} d(c_w, w)$ unless the number of demands in $\mathrm{Ball}(w, \frac{3\, d(c_w, w)}{2\psi})$ becomes greater than $\frac{2\beta f}{3 d(c_w, w)}$. This contradicts to the add-optimality of F^* because for every $\psi \geq \frac{6\beta}{2\beta - 3}$, these demands are closer to each other than to any optimal center in F^*. The details can be found in Section A.2. □

Assignment Cost. Bounding the algorithm's assignment cost is technically involved because we have to consider many different cases. We first distinguish between *inner* and *outer* demands. If the initial assignment cost of a new demand u is within a constant factor from its optimal assignment cost d_u^* (outer demand), then despite the merge operations, the assignment cost of u will remain within a constant factor from d_u^* (Lemma 10 and Lemma 13) . Our main concern is to bound the total assignment cost of the remaining demands (inner demands) throughout the execution of the algorithm.

To provide some intuition, we consider the special case that the optimal solution consists of a single center c. In this case, we further distinguish between *good* and *bad* inner demands. Intuitively, inner demands start as good ones and remain good as long as their assignment cost converges to their optimal assignment cost. Then, they become bad and never become good again. While an inner demand remains good, it is charged with its actual assignment cost. When it becomes bad, it is charged with an irrevocable cost which is an upper bound on its assignment cost at any future point in time (*final assignment cost*).

Let w be the facility which is currently the nearest one to c (in this case, w coincides with the most recent facility to open). Each new inner demand must be assigned to w because new facilities are much closer to c than any of the existing facilities (Lemma 1). Because of the rule for opening new facilities, the total initial assignment cost of the inner demands considered while w is the nearest facility to c cannot exceed βf. Let w' be the first new facility to open after w (w' becomes the nearest facility to c). If w is merged with w', the assignment cost of the inner demands assigned to w decreases by a factor of 2 (Lemma 1). Therefore, the total assignment cost of the inner demands which have always (i.e., from their arrival time until the present time) been assigned to the nearest facility to c (good inner demands) keeps converging to their optimal assignment cost. Therefore, the total assignment cost of good inner demands cannot exceed $2\beta f$ plus thrice their optimal assignment cost (Lemma 3). If w is not merged with w', w must be a supported facility (Lemma 2) and $\mathrm{Asg}^*(B_w) \geq \frac{\beta}{3x} f$ compensates for the (final) assignment cost of the good inner demands assigned to w at the moment w' opens (these demands become bad). From now on, no additional inner demands are assigned to w. Therefore, the total assignment cost of inner demands always remains within a constant factor from the total optimal cost.

We proceed to define formally the basic notions used in the analysis of the assignment cost.

Configuration Distance. For an optimal center $c \in F^*$ and a facility $w \in F$, *the configuration distance between c and w*, denoted by $g(c, w)$, is $g(c, w) = d(c, w) + \frac{x}{x-3} m(w)$. For an optimal center $c \in F^*$, the *configuration distance of* c, denoted by $g(c)$, is $g(c) = \min_{w \in F}\{g(c,w)\} = \min_{w \in F}\{d(c,w) + \frac{x}{x-3}\, m(w)\}$. The configuration distance $g(c)$ is non-increasing with time (Section A.3) and there always exists a facility within a distance of $g(c)$ from c (Proposition 1).

Coalitions. A set of optimal centers $K \subseteq F^*$ with representative $c_K \in K$ forms a *coalition* as long as $g(c_K) \geq \rho D(K)$. A coalition K becomes *broken* as soon as $g(c_K) < \rho D(K)$. A coalition K is *isolated* if $g(c_K) \leq \frac{1}{3}\operatorname{sep}(K)$ and *non-isolated* otherwise. Intuitively, as long as K's diameter is much smaller than $g(c_K)$ (K is a coalition), the algorithm behaves as if K was a single optimal center. If the algorithm is bound to have a facility which is closer to K than any optimal center not in K (K is isolated), then as far as K is concerned, the algorithm behaves as if there were no optimal centers outside K.

A *hierarchical decomposition* $\mathcal{K}$ of F^* is a complete laminar set system[3] on F^*. Every hierarchical decomposition of F^* contains at most $2\,|F^*| - 1$ distinct sets. Given a hierarchical decomposition $\mathcal{K}$ of F^*, we can fix an arbitrary representative c_K for each $K \in \mathcal{K}$ and regard $\mathcal{K}$ as a system of coalitions which hierarchically covers F^*. Formally, given a hierarchical decomposition $\mathcal{K}$ of F^* and the current algorithm's configuration, a set $K \in \mathcal{K}$ is *an active coalition* if K is still a coalition (i.e., $g(c_K) \geq \rho D(K)$), while every superset of K in $\mathcal{K}$ has become broken (i.e., for every $K' \in \mathcal{K}, K \subset K'$, it is $g(c_{K'}) < \rho D(K')$). The current algorithm's configuration induces a collection of active coalitions which form a partitioning of F^*. Since $g(c_K)$ is non-increasing, no coalition which has become broken (isolated) can become active (resp. non-isolated) again.

Let $D_N(K) = \max\{D(K), \frac{1}{3\rho}\operatorname{sep}(K)\}$. By definition, K becomes either isolated or broken as soon as $g(c_K) < \rho D_N(K)$. Using [8, Lemma 1], we show that there is a hierarchical decomposition of F^* such that no coalition K becomes active before $g(c_K) < (\rho+1)\gamma^2 D_N(K)$ (Lemma 7, Section A.4). In the following, we assume that the set of active coalitions is given by a fixed hierarchical decomposition $\mathcal{K}$ of F^* such that for every non-isolated active coalition K, $\rho D_N(K) \leq g(c_K) < (\rho+1)\gamma^2 D_N(K)$.

We use the notions of isolated and non-isolated active coalitions to establish a constant performance ratio for the general case that the optimal solution consists of k centers. More specifically, we prove that (i) isolated active coalitions can be analyzed similarly to the special case that the optimal solution consists of a single facility, and (ii) non-isolated active coalitions, where merge operations do not decrease the assignment cost, can only increase the performance ratio by a constant additive term.

A new demand u makes a coalition K broken/isolated if K has been active/non-isolated before u and becomes broken/isolated after u. Each new demand u is mapped to the unique active coalition K_u containing c_u when u arrives. If K_u is isolated (non-isolated) when u arrives, we say that u is a demand of the isolated (resp. non-isolated) active coalition K_u. Each new facility w is mapped to the unique active coalition containing c_w just before w opens. For an isolated active coalition K, we use w_K to denote the *nearest facility* to K's representative c_K at any given point in time. In other words, w_K is a function always mapping the isolated active coalition K to the facility in F which is currently the nearest facility to c_K. Lemma 1 and Proposition 1 imply that as long as K is an isolated active coalition, w_K is much closer to c_K than any other facility and converges to c_K.

Inner and Outer Demands. A demand u mapped to a non-isolated active coalition K is *inner* if $d_u^* < D_N(K)$, and *outer* otherwise. Let $\operatorname{in}_N(K)$ denote the set of inner demands and $\operatorname{out}_N(K)$ denote the set of outer demands mapped to K as long as K is a non-isolated active coalition. A demand u mapped to an isolated active coalition K is *inner* if $d_u^* < \frac{1}{\lambda}\max\{d(c_K, w'_K), \lambda D(K)\}$, and *outer* otherwise. In this definition, w'_K denotes the nearest facility to c_K at u's assignment time. Thereby, the characterization of a demand u as inner or outer is determined according to the updated algorithm's configuration at u's assignment time, in contrast to the active coalition u is mapped to, which is determined according to the algorithm's configuration at u's arrival time. Let $\operatorname{in}_I(K)$ denote the set of inner demands and $\operatorname{out}_I(K)$ denote the set of outer demands mapped to K as long as K is an isolated active coalition.

[3] 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 $\{c\}$, $c \in F^*$.

Good and Bad Inner Demands. The set of *good demands* of an *isolated* active coalition K, denoted by G_K, consists of the inner demands of K which have *always* (i.e., from their assignment time until the present time) been assigned to w_K (i.e., the nearest facility to c_K). G_K is empty as long as K is either not active or non-isolated. We call *bad* every inner demand of K which is not good. Each new inner demand mapped to an isolated active coalition K is initially assigned to the nearest facility to c_K, because this facility is much closer to c_K than any other facility. Hence, each new inner demand mapped to K becomes good and is added to G_K. An inner demand remains good until either K becomes broken or the location of w_K changes and the facility at the former location w_K is not merged with the facility at the new location w'_K. Then, the demand becomes bad and can never become good again. Since w_K converges to c_K, the actual assignment cost of good inner demands should converge to their optimal assignment assignment cost.

Final Assignment Cost. Let u be a demand currently assigned to a facility w with merge radius $m(w)$. The *final assignment cost* of u, denoted by $\overline{d}_u$, is equal to $\min\{d(u,w) + \frac{x}{x-3}\,m(w), (1+\frac{1}{\lambda})\max\{d(c_K,w), \lambda D(K)\} + d^*_u\}$ if u is mapped to an isolated active coalition K and w is currently the nearest facility to c_K, and equal to $d(u,w) + \frac{x}{x-3}\,m(w)$ otherwise. If a demand u is currently assigned to a facility w, then u will remain assigned to a facility in $\mathrm{Ball}(w, \frac{x}{x-3}\,m(w))$ (Proposition 1). We can also prove that if u is mapped to an isolated active coalition K and is currently assigned to w_K, then u's assignment cost can never exceed $(1+\frac{1}{\lambda})\max\{d(c_K,w_K), \lambda D(K)\} + d^*_u$ (Proposition 19, Section A.6). Therefore, the final assignment cost of u according to the current algorithm's configuration is an upper bound on its actual assignment cost at any future point in time.

With the exception of good demands, each demand is *irrevocably* charged with its final assignment cost at its assignment time. Then, we do not have to worry about the demand's actual assignment cost anymore. On the other hand, we keep track of the actual assignment cost of good demands until they become bad. This is possible because the good demands of an isolated active coalition K are always assigned to the nearest facility to c_K. Good demands are irrevocably charged with their final assignment cost at the moment they become bad. Fig. 3 in the Appendix summarizes the potential function argument used in the analysis of the assignment cost.

Isolated Coalitions. Let K be an isolated active coalition with representative c_K. In the Appendix, we show that (i) every facility spending some time as the nearest facility to c_K is mapped to K and can be merged only with a new facility mapped to K (Proposition 18), (ii) each new facility mapped to K either makes K broken or is at least 2.5 times closer to c_K than the current location of w_K (Proposition 20, see also Lemma 1), (iii) an unsupported facility mapped to K is merged with the next facility mapped to K (Proposition 8, see also Lemma 2), (iv) a new demand/facility not mapped to K cannot change either the location of w_K or the value of $g(c_K)$ (Lemma 8), (v) each inner demand of K which does not make K broken is initially assigned to the nearest facility to c_K (Lemma 9), and (vi) for every outer demand u which is mapped to K and does not make K broken, $\overline{d}_u \le 4(\lambda+2)d^*_u$ (Lemma 10). The properties (i)-(v) imply that the assignment cost of the inner demands of K can be analyzed independently from other active coalitions and similarly to the special case that there is a single optimal center.

Lemma 3. *Let K be an isolated active coalition. The total actual and the total final assignment cost of the good demands of K can be bounded as:* $\sum_{u\in \mathrm{G}_K} d(u,w_K) < 2\beta f + 3\sum_{u\in \mathrm{G}_K} d^*_u$ *and* $\sum_{u\in \mathrm{G}_K} \overline{d}_u < 4.5\beta f + 7\sum_{u\in \mathrm{G}_K} d^*_u$.

Proof Sketch. We sketch the proof of the first inequality. The second inequality can be derived from the first one using the definition of the final assignment cost. The full proof can be found in Section A.6.

The proof is by induction over a sequence of merge operations where the former nearest facility to c_K is merged with the new nearest facility to c_K. Let w be the nearest facility to c_K, i.e., $w_K = w$. By (i) and (iv), the location of w_K cannot change until a new facility mapped to K opens. Let w' be the next facility mapped to K and let G_K be the set of good demands just before w' opens. Wlog. we can assume that w' does not make K broken, since G_K becomes empty otherwise. Then, by (ii), $d(c_K, w') \le \frac{2}{5}\,d(c_K, w)$, and the location of the nearest facility to c_K must change from $w_K = w$ to $w'_K = w'$ as soon as w' opens. We inductively assume that the inequality holds just before w' opens, and we show that it remains valid until either the location of the nearest facility to c_K changes again or K becomes broken.

If w is not merged with w', the set of good demands becomes empty. Then, $\sum_{u \in \mathrm{G}_K} d(u, w') = 0$ just after w' opens. If w is merged with w', (ii) implies that for every $u \in \mathrm{G}_K$, $d(u, w') \leq \frac{1}{2} d(u, w) + \frac{3}{2} d^*_u$. Just after w has been merged with w', the set of good demands remains G_K, but each $u \in \mathrm{G}_K$ is now assigned to w'. Hence, $\sum_{u \in \mathrm{G}_K} d(u, w') \leq \beta f + 3 \sum_{u \in \mathrm{G}_K} d^*_u$. As long as w' remains the nearest facility to c_K and K remains an isolated active coalition, each new inner demand of K is initially assigned to w' (cf. (v) above) and becomes a good demand. Let $\mathrm{G}_K(w')$ be the set of good inner demands of K whose initial assignment takes place while w' is the nearest facility to c_K. We prove that the total initial assignment cost of the demands in $\mathrm{G}_K(w')$ is at most βf.

More specifically, if $d(c_K, w') \geq \lambda D(K)$, we observe that the demands in $\mathrm{G}_K(w')$ remain unsatisfied and hold a potential equal to their initial assignment cost (i.e., their distance from w') until either a new facility mapped to K opens or K becomes broken. This is true because every facility which makes some of them satisfied or decrease their potential must be mapped to K (see also (ii) above). In addition, since $d(c_K, w') \geq \lambda D(K)$, for each new inner demand v which is added to $\mathrm{G}_K(w')$, $d^*_v \leq \frac{1}{\lambda} d(c_K, w')$ and $d(c_K, c_v) \leq D(K) \leq \frac{1}{\lambda} d(c_K, w')$. Therefore, v's neighborhood $B_v = \mathrm{Ball}(v, \frac{d(F_v, v)}{x}) \cap L$ includes every demand in $\mathrm{G}_K(w')$, because $\frac{d(F_v,v)}{x} = \frac{d(w',v)}{x} > \frac{3}{\lambda} d(c_K, w')$ (recall that $\lambda = 3x + 2$) and for every $u \in \mathrm{G}_K(w')$, $d(u, v) < \frac{3}{\lambda} d(c_K, w')$. Consequently, the potential accumulated by $\mathrm{G}_K(w')$ is at most βf (Lemma 11). On the other hand, if $d(c_K, w') < \lambda D(K)$, we observe that as long as K remains active, it must be $m(w') \geq \psi(\lambda + 2) D(K)$ ($\rho = (\psi + 2)(\lambda + 2)$, Proposition 12). Hence, $\mathrm{G}_K(w') \subseteq \mathrm{Init}(w') \cap \mathrm{Ball}(w', \frac{m(w')}{\psi})$, and we can use Ineq. (1) (Lemma 12). □

Using a potential function argument based on Lemma 3 and the claims (i) - (vi) above, we can bound the assignment cost of the demands mapped to K (see also Section A.8). When K becomes an isolated active coalition, it receives a credit of $7\beta f$, which is not used until K becomes broken. Let u be a new demand mapped to the isolated active coalition K. If u makes K broken, K's credit is charged with u's final assignment cost, which cannot exceed $2.5\beta f$ (Proposition 14). If u is an outer demand and does not make K broken, its final assignment cost cannot exceed $4(\lambda + 2) d^*_u$ by (vi). If u is an inner demand and does not make K broken, it is initially assigned to the nearest facility to c_K (cf. (v) above) and becomes a good demand. As long as u remains a good demand, its actual assignment cost is equal to $d(u, w_K)$. By Lemma 3, the actual assignment cost of the demands in G_K never exceeds $2\beta f$ plus 3 times their optimal assignment cost. As long as K remains active, its credit can absorb the additional cost of $2\beta f$.

The good inner demands of K are charged with their final assignment cost as soon as they become bad and G_K becomes empty. By Lemma 3, the total final assignment cost of the demands in G_K does not exceed $4.5\beta f$ plus 7 times their optimal assignment cost. If G_K becomes empty because K becomes broken, the additional cost of $4.5\beta f$ is charged to K's credit. Otherwise, G_K becomes empty because the location of the nearest facility to c_K has changed and the facility w at the previous location w_K is not merged with the new facility w' at the new location w'_K. By (i), both w and w' are mapped to K. Then, by (iii), the facility w must be a supported facility. Hence, the additional cost of $4.5\beta f$ can be charged to the optimal assignment cost of the demands contributing to the opening cost of w, since $3x\, \mathrm{Asg}^*(B_w) \geq \beta f$. We also prove that each supported facility is charged with the final assignment cost of some good demands which become bad at most once (Proposition 21, see also Section A.8).

Since we consider at most $2\,|F^*| - 1$ different coalitions and each of them can become isolated at most once, the total assignment cost of the demands in $C_I = \bigcup_{K \in \mathcal{K}} \mathrm{in}_I(K) \cup \mathrm{out}_I(K)$ (i.e., the demands mapped to isolated active coalitions) is at most $14\beta\, \mathrm{Fac}^* + 4(\lambda + 2)\mathrm{Asg}^*(C_I) + 14x \mathrm{Asg}^*$.

Non-isolated Coalitions. The demands mapped to non-isolated active coalitions are irrevocably charged with their final assignment cost at their assignment time. The analysis of the assignment cost is based on the notion of unsatisfied inner demands. The set of *unsatisfied inner demands* of a non-isolated active coalition K, denoted by N_K, consists of the inner demands of K which are currently unsatisfied. N_K is equal to $\mathrm{in}_N(K) \cap L$ as long as K is a non-isolated active coalition, and empty otherwise.

Lemma 4. *For every non-isolated active coalition K, $|\mathrm{N}_K| \cdot g(c_K)$ never exceeds $(\psi+4)\gamma^2\beta f$.*

Proof Sketch. Since $g(c_K)$ is non-increasing, the product $|\mathrm{N}_K| \cdot g(c_K)$ can only increase if a new demand is added to the set of unsatisfied inner demands $\mathrm{N}_K = \mathrm{in}_N(K) \cap L$. We recall that for each $u \in \mathrm{in}_N(K)$, $d_u^* < D_N(K)$ and $d(c_K, c_u) \leq D_N(K)$. Let v be the last demand added to N_K.

If $d(F_v, c_K) \geq \lambda D_N(K)$, then $B_v = \mathrm{Ball}(v, \frac{d(F_v,v)}{x}) \cap L$ includes every demand in N_K because $\frac{d(F_v,v)}{x} > 3D_N(K)$ ($\lambda = 3x+2$) and $d(u,v) < 3D_N(K)$ for every $u \in \mathrm{in}_N(K)$. In addition, it must be $\mathrm{Pot}(B_v) < \beta f$ because v remains unsatisfied. Since for every $u \in \mathrm{in}_N(K)$, $d(F_v, u) > (\lambda-2)D_N(K)$ and $\mathrm{N}_K \subseteq B_v$, we conclude that $|\mathrm{N}_K|(\lambda-2)D_N(K) < \beta f$. Then, the lemma follows from $g(c_K) < (\rho+1)\gamma^2 D_N(K)$, because K is active when v arrives.

On the other hand, if $d(F_v, c_K) < \lambda D_N(K)$, let w be the nearest facility to c_K when v arrives. By hypothesis, $d(c_K, w) < \lambda D_N(K)$. We observe that as long as K remains a non-isolated active coalition, it must be $m(w) \geq \psi(\lambda+2)D_N(K)$ ($\rho = (\psi+2)(\lambda+2)$, Proposition 12). Since for every $u \in \mathrm{in}_N(K)$, $d(w,u) \leq (\lambda+2)D_N(K)$, we obtain that $\mathrm{N}_K \subseteq \mathrm{Ball}(w, \frac{m(w)}{\psi})$. At the moment w opened, it must have been $d(F_w, w) \geq x(\lambda+2)D_N(K)$, because otherwise, w would have made K either isolated or broken. Hence, the set of the unsatisfied inner demands of K became empty when w opened. In addition, for each new facility w' which opens while w is still open, either $d(c_K, w') < 2(\lambda+2)D_N(K)$ or w' makes K either isolated or broken. Therefore, as long as K remains a non-isolated active coalition, w is much closer to c_K than any other facility. Hence, every demand which is in N_K at the moment v arrives, it must have been initially assigned to w. Consequently, $\mathrm{N}_K \subseteq \mathrm{Init}(w) \cap \mathrm{Ball}(w, \frac{m(w)}{\psi})$. Then, the lemma follows from Ineq. (1). The full proof can be found in Section A.7. □

In addition to Lemma 4, we use the following properties: (i) for every non-isolated active coalition K, $g(c_K) \in [\,\rho D_N(K), (\rho+1)\gamma^2 D_N(K)\,)$ (Lemma 7), and (ii) for each new facility w', if $B_{w'} \cap \mathrm{in}_N(K) \neq \emptyset$, then $g'(c_K) \leq \frac{1}{3}\, g(c_K)$ (Proposition 22).

As long as K is a non-isolated active coalition, it holds a credit of $(5(\psi+4)\gamma^2 + 2.5)(\ln(\frac{g(c_K)}{\rho D_N(K)}) + 1)\beta f$. In addition, the function $-\Upsilon_K^{(N)} = -5\,|\mathrm{N}_K| \cdot g(c_K)$ accounts for the final assignment cost of the demands in N_K which has not charged to K's credit yet (see also Section A.8). By Lemma 4, $\Upsilon_K^{(N)}$ is always bounded by $5(\psi+4)\gamma^2\beta f$.

If the new demand u either makes K isolated or broken or opens a new facility w' such that $B_{w'} \cap \mathrm{in}_N(K) \neq \emptyset$, the final assignment cost of u is bounded by $2.5\beta f$ and is charged to K's credit. In this case, the function $-\Upsilon_K^{(N)}$ may increase because some demands may be removed from N_K. However, the increase in $-\Upsilon_K^{(N)}$ cannot exceed $5(\psi+4)\gamma^2\beta f$. If K becomes isolated or broken, K's credit become 0. Hence, it decreases by at least $(5(\psi+4)\gamma^2 + 2.5)\beta f$. If u opens a new facility w' such that $B_{w'} \cap \mathrm{in}_N(K) \neq \emptyset$, then $g(c_K)$ decreases by a factor of 3 (see also (ii) above) and K's credit decreases by $(5(\psi+4)\gamma^2 + 2.5)\beta f$. In both cases, the decrease in K's credit compensates for the final assignment cost of u and the increase in $-\Upsilon_K^{(N)}$. Otherwise, if u is an outer demand, then $\overline{d}_u \leq 4[(\rho+1)\gamma^2 + 2]\, d_u^*$ (Lemma 13) and its final assignment cost is charged to its optimal assignment cost. If u is an inner demand, then $\overline{d}_u \leq 5\, g_u'(c_K)$ (Lemma 14). In this case, u is added to the set of unsatisfied inner demands N_K and the function $-\Upsilon_K^{(N)} = -5\,|\mathrm{N}_K| \cdot g(c_K)$ decreases and compensates for the final assignment cost of u.

Consequently, the total assignment cost of the demands in $C_N = \bigcup_{K \in \mathcal{K}} \mathrm{in}_N(K) \cup \mathrm{out}_N(K)$ (i.e., the demands mapped to non-isolated active coalitions) is at most $2\beta\,[5(\psi+4)\gamma^2 + 2.5][\ln((1+\frac{1}{\rho})\gamma^2) + 1]\,\mathrm{Fac}^* + 4\,[(\rho+1)\gamma^2 + 2]\,\mathrm{Asg}^*(C_N)$. By partitioning the interval $[\,\rho D_N(K), (\rho+1)\gamma^2 D_N(K)\,)$ into disjoint sub-intervals $[\,2^i \rho D_N(K), 2^{i+1}\rho D_N(K)\,)$ and considering different phases according to the sub-interval $g(c_K)$ belongs to, we can improve the previous bound to $4\beta\log(\gamma)(12(\psi+2) + 2.5)\mathrm{Fac}^* + 8(\rho+1)\mathrm{Asg}^*(C_N)$.

3 An Incremental Algorithm for k-Median

To obtain an incremental algorithm for k-Median, we are based on the following standard lemma which is proven in the Appendix, Section A.9. We recall that a_1, a_2, b_1, and b_2 denote the constants in the performance ratio of IFL.

Lemma 5. *Let* Asg^* *be the cost of a feasible solution for an instance of k-Median, let* Λ *be an estimation of* Asg^**, and let* $\delta = \frac{a_2}{b_2}$*. Then,* IFL *with facility cost* $f = \frac{\Lambda}{\delta k}$ *maintains a solution of cost no greater than* $(a_2 + b_2)\,\mathrm{Asg}^* + b_2\,\Lambda$ *which consists of no more than* $(a_1 + a_2 \frac{b_1}{b_2} \frac{\mathrm{Asg}^*}{\Lambda})\,k$ *medians.*

The algorithm $\mathrm{IM}(k)$ (Fig. 4 in the Appendix) operates in phases using IFL as a building block. Phase i is characterized by an upper bound Λ_i on the optimal assignment cost of the demands considered in the current phase. $\mathrm{IM}(k)$ invokes IFL with facility cost $f_i = \frac{\Lambda_i}{\delta k}$, where $\delta = \frac{a_2}{b_2}$. Lemma 5 implies that as long as Λ_i is a valid upper bound, IFL maintains a solution consisting of no more than νk medians and costing at most $\mu \Lambda_i$, where ν, μ are appropriately chosen constants. Therefore, as soon as either the number of medians exceeds νk or the cost exceeds $\mu \Lambda_i$, we can be sure that the optimal cost has also exceeded Λ_i. Then, $\mathrm{IM}(k)$ merges the medians produced by the current phase with the medians produced by the previous phases, increases the upper bound by a constant factor, and proceeds with the next phase. The algorithm $\mathrm{IM}(k)$ is deterministic and runs in $\mathrm{O}(n^2 k)$ time and $\mathrm{O}(n)$ space. The proof of the following theorem follows from Lemma 5. The details can be found in the Appendix, Section A.10.

Theorem 2. *The algorithm* $\mathrm{IM}(k)$ *achieves a constant performance ratio using* $\mathrm{O}(k)$ *medians.*

The randomized algorithm $\mathrm{RIM}(k)$ (Fig. 5 in the Appendix) uses IFL and Gather as building blocks. The algorithm Gather (Fig. 6 in the Appendix), which is actually a modification of the algorithm PARA_CLUSTER [6], is made up of $\mathrm{O}(\log n)$ independent invocations of Meyerson's randomized algorithm [19] with facility cost $\hat{f}_i = \frac{\Lambda_i}{k(\log n + 1)}$. The algorithm $\mathrm{RIM}(k)$ uses Gather to generate a modified instance which can be represented in a space efficient manner. The modified instance contains the same number of different unit demands, which now occupy only $\mathrm{O}(k \log^2 n)$ different locations. Then, $\mathrm{RIM}(k)$ uses IFL with facility cost $f_i = \frac{\Lambda_i}{\delta k}$ to cluster the modified instance.

For an incremental implementation, each new demand is first moved to a gathering point by Gather. Then, a new demand located at the corresponding gathering point is given to IFL, which assigns it to a median. Both actions are completed before the next demand is considered. The current phase ends if either the number of gathering points, the gathering cost, the number of medians, or the assignment cost on the modified instance become too large.

We should emphasize that IFL treats the demands moved to the same gathering point by Gather as *different demands* and may put them in different clusters[4]. In other words, the output of Gather is regarded as a sample taken from the points of the metric space and not as a first-level clustering. This sample is only used to improve the time and space efficiency of IFL. On the other hand, the solution produced by IFL on the modified instance can be directly translated into a hierarchical clustering of the original instance.

Since the demands considered by IFL occupy only $\mathrm{O}(k \log^2 n)$ different locations, $\mathrm{RIM}(k)$ can be implemented in $\mathrm{O}(nk^2 \log^2 n)$ time and $\mathrm{O}(k^2 \log^2 n)$ space. Similarly to the analysis of [6], we can prove that the gathering step increases the performance ratio by no more than a constant factor whp. In contrast to $\mathrm{IM}(k)$ which does not require any advance knowledge of n, $\mathrm{RIM}(k)$ needs to know a constant factor approximation to $\log n$ in advance. The details can be found in Section A.11.

Theorem 3. *The algorithm* $\mathrm{RIM}(k)$ *runs in* $\mathrm{O}(nk^2 \log^2 n)$ *time and* $\mathrm{O}(k^2 \log^2 n)$ *space and achieves a constant performance ratio whp. using* $\mathrm{O}(k)$ *medians.*

[4] This actually increases the algorithm's time and space complexity by an additional factor of k.

4 Open Problems

An interesting open problem is to determine whether there exists a time and space efficient incremental algorithm for k-Median which does not assume any advance knowledge of n and achieves a constant performance ratio using $\mathrm{O}(k)$ medians. Another interesting research direction is to improve the constants involved in the performance ratio of IFL. For isolated coalitions, the performance ratio can be significantly improved by a careful analysis. On the other hand, the analysis of non-isolated coalitions increases the performance ratio by a large constant additive term. In addition to a really careful analysis, some new ideas concerning the analysis of non-isolated coalitions are required for establishing a performance ratio of practical interest.

References

1. A. Anagnostopoulos, R. Bent, E. Upfal, and P. Van Hentenryck. A Simple and Deterministic Competitive Algorithm for Online Facility Location. TR CS-03-16, Brown University, 2003.
2. V. Arya, N. Garg, R. Khandekar, A. Meyerson, K. Munagala, and V. Pandit. Local Search Heuristics for k-Median and Facility Location Problems. In *Proc. of the 33rd ACM Symp. on Theory of Computing (STOC '01)*, pp. 21–29, 2001.
3. A. Borodin and R. El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, 1998.
4. 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.
5. 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.
6. 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.
7. 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.
8. D. Fotakis. On the Competitive Ratio for Online Facility Location. In *Proc. of ICALP '03*, LNCS 2719, pp. 637–652, 2003.
9. D. Fotakis. Incremental Algorithms for Facility Location and k-Median. Manuscript available from `http://students.ceid.upatras.gr/~fotakis/data/incremen.pdf`, 2004.
10. S. Guha. *Approximation Algorithms for Facility Location Problems*. PhD thesis, Stanford University, 2000.
11. 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.
12. S. Guha, N. Mishra, R. Motwani, and L. O'Callaghan. Clustering Data Streams. In *Proc. of the 41st IEEE Symp. on Foundations of Computer Science (FOCS '00)*, pp. 359–366, 2000.
13. M.R. Henzinger, P. Raghavan, and S. Rajagopalan. Computing on Data Streams. TR 1998-011, DEC Systems Research Center, 1998.
14. 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.
15. 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.
16. 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.
17. 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.
18. 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.
19. A. Meyerson. Online Facility Location. In *Proc. of the 42nd IEEE Symp. on Foundations of Computer Science (FOCS '01)*, pp. 426–431, 2001.
20. C.G. Plaxton. Approximation Algorithms for Hierarchical Location Problems. In *Proc. of the 35th ACM Symp. on Theory of Computing (STOC '03)*, pp. 40–49, 2003.
21. D. Shmoys. Approximation Algorithms for Facility Location Problems. In *Proc. of APPROX '00*, LNCS 1913, pp. 27–33, 2000.
22. 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.
23. 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 Basic Properties

The Proof of Lemma 1. We assume that there exists a facility w such that $d(c_w, w) > \frac{d(F_w,c_w)}{3}$, and we show that this assumption contradicts to the add-optimality of F^*, i.e., if there exists such a facility w, then

$$f + \sum_{u \in B_w} d(u,w) \le \sum_{u \in B_w} \Big(\tfrac{1}{\beta} d(F_w,u) + d(u,w)\Big) < \sum_{u \in B_w} d_u^* \tag{2}$$

where the first inequality follows from $\mathrm{Pot}(B_w) = \sum_{u \in B_w} d(F_w, u) \ge \beta f$. In other words, we could decrease the total cost of F^* by opening a new facility at w. For every $u \in B_w$, we bound d_u^* from below in terms of $d(u,w)$ and $d(F_w,u)$. We recall that $B_w = \mathrm{Ball}(w, \frac{d(F_w,w)}{x}) \cap L$ is the set of unsatisfied demands contributing to the opening cost of w. Thus, for every $u \in B_w$,

$$d(u,w) \le \tfrac{d(F_w,w)}{x} \le \tfrac{d(c_w,w)}{x} + \tfrac{d(F_w,c_w)}{x} < \tfrac{4}{x} d(c_w,w) \tag{3}$$

where the last inequality follows from the assumption that $d(w,c_w) > \frac{d(F_w,c_w)}{3}$. In addition, for every $u \in B_w$,

$$d_u^* = d(c_u,u) \ge d(c_u,w) - d(u,w) > d(c_w,w) - \tfrac{4}{x} d(c_w,w) = \tfrac{x-4}{x} d(c_w,w) \tag{4}$$

where the second inequality follows from Ineq. (3) and the fact that w is mapped to c_w instead of c_u. Using Ineq. (3) and Ineq. (4) , we obtain the following lower bound on d_u^* in terms of $d(u,w)$.

$$d(u,w) < \tfrac{4}{x} d(c_w,w) < \tfrac{4}{x}\tfrac{x}{x-4} d_u^* = \tfrac{4}{x-4} d_u^* \tag{5}$$

We also obtain the following lower bound on d_u^* in terms of $d(F_w,u)$.

$$d(F_w,u) \le d(F_w,c_w) + d(c_w,w) + d(u,w) < \tfrac{4(x+1)}{x} d(c_w,w) < \tfrac{4(x+1)}{x-4} d_u^* \tag{6}$$

where the second inequality follows from the Ineq. (3) and the hypothesis that $d(c_w,w) > \frac{d(F_w,c_w)}{3}$, and the third inequality from Ineq. (4). Using inequalities (5) and (6) and assuming that $\frac{1}{\beta}\frac{4(x+1)}{x-4} + \frac{4}{x-4} \le 1$, we obtain Ineq. (2), which contradicts to the add-optimality of F^*. $\square$

The Proof of Proposition 1. To show that $\mathrm{Ball}(w', \frac{x}{x-3} m(w'))$ is included in $\mathrm{Ball}(w, \frac{x}{x-3} m(w))$, we apply the following proposition for $p = w$.

Proposition 2. *Let w be a facility merged with a new facility w'. Then, for every point p, $d(w',p) + \frac{x}{x-3}\, m(w') \le d(w,p) + \frac{x}{x-3}\, m(w)$.*

Proof. Since w is merged with w', $d(w,w') \le m(w)$. Therefore,

$$\begin{aligned} d(w',p) + \tfrac{x}{x-3}\, m(w') &\le d(w,p) + d(w,w') + \tfrac{x}{x-3}\,\tfrac{3}{x}\, d(w,w') \\ &= d(w,p) + (1 + \tfrac{3}{x-3})\, d(w,w') \le d(w,p) + \tfrac{x}{x-3} m(w)\,, \end{aligned}$$

where the first inequality follows from $m(w') \le m^{(1)}(w') \le \frac{3}{x}\, d(w,w')$, since $w \in F_{w'}$. $\square$

A.2 Facility Cost

Proposition 3. *For every unsupported facility w mapped to c_w, $d(c_w,w) < \frac{4}{3(x-4)}\, d(F_w,c_w)$.*

Proof. Since w is an unsupported facility and $\sum_{u \in B_w} d_u^* < \frac{1}{3x}\,\beta f \le \frac{1}{3x} \sum_{u \in B_w} d(F_w,u)$, there must be at least one demand $u \in B_w$ such that $d_u^* < \frac{1}{3x}\, d(F_w,u)$. Let c_u be the optimal center the demand u is mapped to. Then, $d_u^* = d(c_u,u) < \frac{1}{3x-1}\, d(F_w,c_u)$. We first establish that both $d(c_w,w)$ and $d(c_u,w)$ are bounded by $\frac{4x-1}{(3x-1)(x-1)}\, d(F_w,c_u)$.

$$d(c_u,w) \le d(w,u) + d(c_u,u) \le \tfrac{d(F_w,w)}{x} + d(c_u,u) < \tfrac{d(F_w,c_u)}{x} + \tfrac{d(c_u,w)}{x} + \tfrac{d(F_w,c_u)}{3x-1}\,.$$

Therefore, $d(c_w, w) \le d(c_u, w) < \frac{4x-1}{(3x-1)(x-1)} d(F_w, c_u)$, where the first inequality holds because w is mapped to c_w instead of c_u. Since both $d(c_w, w)$ and $d(c_u, w)$ are small fractions of $d(F_w, c_u)$, the distances $d(F_w, c_w)$ and $d(F_w, c_u)$ cannot differ by two much. In particular,

$$d(F_w, c_u) \le d(F_w, c_w) + d(c_w, w) + d(c_u, w) < d(F_w, c_w) + 2\,\frac{4x-1}{(3x-1)(x-1)}\,d(F_w, c_u)\,.$$

Therefore, $d(c_w, w) < \frac{4x-1}{3x^2-12x+3}\,d(F_w, c_w) < \frac{4}{3(x-4)}\,d(F_w, c_w)$. □

Proposition 4. *Let $x \ge 16$. For every unsupported facility w mapped to c_w, $m^{(1)}(w) \ge \frac{3}{2}\,d(c_w, w)$.*

Proof. We first bound r_w, i.e., the radius of B_w, from below:

$$r_w = \frac{d(F_w,w)}{x} \ge \frac{1}{x}(d(F_w, c_w) - d(c_w, w)) > \frac{1}{x}(\frac{3(x-4)}{4} - 1)d(c_w, w)\,,$$

where the last inequality follows from Proposition 3. Since $m^{(1)}(w)$ is equal to $3\,r_w$, we conclude that $m^{(1)}(w) > \frac{9x-48}{4x}\,d(c_w, w) \ge \frac{3}{2}d(c_w, w)$, where the last inequality holds for every $x \ge 16$. □

Proposition 5. *Let $\psi \ge \frac{6\beta}{2\beta-3}$. For every facility w, $m^{(2)}(w) \ge \frac{3}{2}\,d(c_w, w)$.*

Proof. The proof is similar to the proof of Lemma 1. To reach a contradiction, we assume that $m^{(2)}(w)$ is less than $\frac{3}{2}d(c_w, w)$. Let $B^{(2)} = \mathrm{Ball}(w, \frac{3d(c_w,w)}{2\psi}) \cap \mathrm{Init}(w)$. Since we have assumed that $m^{(2)}(w) < \frac{3}{2}\,d(c_w, w)$, it must be $|B^{(2)}| \cdot \frac{3d(c_w,w)}{2} > \beta f$ by the definition of $m^{(2)}(w)$ (see also Fig. 1). Using this inequality, we will establish that

$$f + \textstyle\sum_{u\in B^{(2)}} d(u, w) < \sum_{u\in B^{(2)}} (\frac{3}{2\beta}d(c_w, w) + d(u, w)) \le \sum_{u\in B^{(2)}} d^*_u \tag{7}$$

which contradicts to the add-optimality of F^*. We first observe that

$$d^*_u = d(c_u, u) \ge d(c_u, w) - d(u, w) \ge d(c_w, w) - d(u, w)\,.$$

Since for every $u \in B^{(2)}$, $d(u, w) \le \frac{3d(c_w,w)}{2\psi}$, we obtain that $d(c_w, w) \le \frac{2\psi}{2\psi-3}d^*_u$ and $d(u, w) \le \frac{3}{2\psi-3}d^*_u$. Using the inequalities above and assuming that $\frac{3}{2\beta}\frac{2\psi}{2\psi-3} + \frac{3}{2\psi-3} \le 1$, we obtain Ineq. (7). □

A.3 The Configuration Distance is Non-Increasing

Proposition 6. *For every point p, the quantity $g(p) = \min_{w\in F}\{d(w, p) + \frac{x}{x-3}\,m(w)\}$ is non-increasing with time.*

Proof. Let w be a facility in F. As long as w remains open, the quantity $d(w, p) + \frac{x}{x-3}\,m(w)$ cannot increase because the algorithm keeps decreasing $m(w)$ to maintain Ineq. (1). If w is merged with a new facility w', it must be $d(w', p) + \frac{x}{x-3}\,m(w') \le d(w, p) + \frac{x}{x-3}\,m(w)$ by Proposition 2. □

A.4 No Coalition Becomes Active before $g(c_K) < (\rho + 1)\gamma^2 D_N(K)$

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 component $K \in \mathcal{K}$ different from the root of the decomposition tree $T_{\mathcal{K}}$, we use p_K to denote the immediate ancestor/parent of K in $T_{\mathcal{K}}$. Throughout this section, we use the hierarchical decomposition $\mathcal{K}$ and its tree representation $T_{\mathcal{K}}$ interchangeably.

Next, we prove that there is a hierarchical decomposition $\mathcal{K}$ of F^* such that for any non-isolated active coalition $K \in \mathcal{K}$, $\rho D_N(K) \le g(c_K) < (\rho + 1)\gamma^2 D_N(K)$ (Lemma 7). The proof of Lemma 7 follows from the fact that any metric space has a hierarchical decomposition such that each component either is relatively well-separated or has a relatively large diameter (see also [8]). For completeness, we give a proof of this claim before we establish Lemma 7.

Lemma 6. *For any metric space M and any $\gamma \geq 16$, there is a hierarchical decomposition $\mathcal{K}$ of M such that for any set $K \in \mathcal{K}$ different from M, either $D(K) > \frac{D(\mathsf{p}_K)}{\gamma^2}$ or $\mathrm{sep}(K) > \frac{D(\mathsf{p}_K)}{4\gamma}$.*

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

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. For $i = 0$, M is the only level 0 group. Given the i-th level, the next level is constructed as follows: For each G_j^i of $D(G_j^i) > \frac{D}{\gamma^{i+1}}$ (Fig. 2, large diameter groups), we set $\tilde{G}_j^i \leftarrow G_j^i$. While $\tilde{G}_j^i$ violates (ii) for the $(i+1)$-th level (i.e., while $D(\tilde{G}_j^i) > \frac{D}{\gamma^{i+1}}$ and there exists a $G \subseteq \tilde{G}_j^i$ such that $D(G) \leq \frac{D}{\gamma^{i+2}}$ and $d(G, \tilde{G}_j^i \setminus G) > \frac{D}{4\gamma^{i+1}}$), we create a new level $i+1$ group for G, remove G from $\tilde{G}_j^i$ ($\tilde{G}_j^i \leftarrow \tilde{G}_j^i \setminus G$), and iterate. We also create a new level $i+1$ group consisting of the points remaining in $\tilde{G}_j^i$ after the loop. The level $i+1$ groups created from G_j^i form a partitioning of it and are G_j^i's children. If $D(G_j^i) \leq \frac{D}{\gamma^{i+1}}$, G_j^i has a single child at level $i+1$ which is identical to it (Fig. 2, small diameter groups).

It is straight-forward that the above procedure results in level $i+1$ groups which fulfill condition (ii). Moreover, since for any $j_1 \neq j_2$, $d(G_{j_1}^i, G_{j_2}^i) > \frac{D}{4\gamma^i}$, any child of $G_{j_1}^i$ must be at distance greater than $\frac{D}{4\gamma^i} > \frac{D}{4\gamma^{i+1}}$ from any child of $G_{j_2}^i$. As for the distance between different children of G_j^i, when a subset G is removed, any point in G is at distance greater than $\frac{D}{4\gamma^{i+1}}$ from the points remaining in $\tilde{G}_j^i$. For any G' removed 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 decreasing with i, we eventually reach a level ν such that all the level ν groups consist of a single point.

Level i groups are further partitioned into *level i components* $K_1^i, \ldots, K_m^i$ such that for any j, (i) $D(K_j^i) \leq \frac{D}{\gamma^i}$, and (ii) either $D(K_j^i) > \frac{D}{\gamma^{i+1}}$ or $\mathrm{sep}(K_j^i) > \frac{D}{4\gamma^i}$. To ensure the hierarchical structure, we proceed inductively in a bottom-up fashion. Each level ν group consists of a single point. Hence, we create a single level ν component for each level ν group. We inductively assume the collection $\mathcal{K}_j^{i+1} = \{K_1, \ldots, K_m\}$ consisting of all the level $i+1$ components the children of 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 $D(G_j^i) \leq \frac{D}{\gamma^i}$, we create a single level i component for the level i group G_j^i (Fig. 2, (a) and (b)). We recall that $\mathrm{sep}(G_j^i) > \frac{D}{4\gamma^i}$. If $D(G_j^i) > \frac{D}{\gamma^i}$, 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 of $\mathcal{K}_j^{i+1}$. Intuitively, this is true because G_j^i does not contain any subset $G \subseteq G_j^i$ such that both $D(G) \leq \frac{D}{\gamma^{i+1}}$ and $d(G, G_j^i \setminus G) > \frac{D}{4\gamma^i}$ (i.e., G_j^i does not contain any small-diameter subsets which are well-separated, Fig. 2, (c)). In other words, 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 of $\mathcal{K}_j^{i+1}$, we maintain two disjoint collections Z_1 and Z_2. Z_1 contains components of diameter no greater than $\frac{D}{\gamma^{i+1}}$, while Z_2 contains components of diameter greater than $\frac{D}{\gamma^{i+1}}$ obtained by merging some of the components in Z_1. Initially, $Z_1 = \mathcal{K}_j^{i+1}$.

While there exist $K_1, K_2 \in Z_1$ such that $d(K_1, K_2) \leq \frac{D}{4\gamma^i}$, K_1 and K_2 are removed from Z_1 and merged into a new component K, $K = K_1 \cup K_2$. If $D(K) \leq \frac{D}{\gamma^{i+1}}$, K is put in Z_1, otherwise, K is put in Z_2. Due to the choice of K_1 and K_2, for any $K \in Z_2$, $D(K) \leq \frac{D}{\gamma^i}\left(\frac{1}{4} + \frac{2}{\gamma}\right)$.

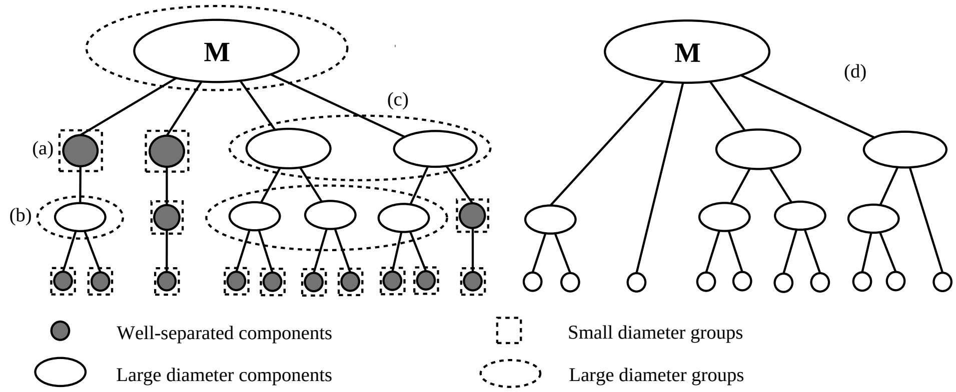

Fig. 2. The hierarchical decomposition of Lemma 6.

The merge procedure above cannot terminate with an empty collection Z_2. If Z_2 were empty, Z_1 would contain more than one component, because $D(G_j^i) > \frac{D}{\gamma^i}$ and, for any $K \in Z_1$, $D(K) \leq \frac{D}{\gamma^{i+1}}$. In addition, for any $K_1, K_2 \in Z_1$, $d(K_1, K_2) > \frac{D}{4\gamma^i}$, 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 is a contradiction.

Therefore, the merge procedure always terminates with a non-empty collection Z_2. If Z_1 is non-empty, every $K \in Z_1$ is associated with a component $K' \in Z_2$ such that $d(K, K') \leq \frac{D}{4\gamma^i}$. For any $K \in Z_1$, such a $K' \in Z_2$ 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.

Then, each component in Z_2 is merged with the components of Z_1 associated with it. The resulting components have diameter greater than $\frac{D}{\gamma^{i+1}}$, because all of them include a single component of Z_2. In addition, their diameter cannot exceed $\frac{D}{\gamma^{i+1}} + \frac{D}{4\gamma^i} + \frac{D}{\gamma^i}\left(\frac{1}{4} + \frac{2}{\gamma}\right) + \frac{D}{4\gamma^i} + \frac{D}{\gamma^{i+1}} = \frac{D}{\gamma^i}\left(\frac{3}{4} + \frac{4}{\gamma}\right)$, because the diameter of any component in Z_1 is at most $\frac{D}{\gamma^{i+1}}$, the diameter of the component in Z_2 is at most $\frac{D}{\gamma^i}\left(\frac{1}{4} + \frac{2}{\gamma}\right)$, and any component of Z_1 has been associated with a component of Z_2 at distance no greater than $\frac{D}{4\gamma^i}$. For $\gamma \geq 16$, $\frac{D}{\gamma^i}\left(\frac{3}{4} + \frac{4}{\gamma}\right)$ is no greater than $\frac{D}{\gamma^i}$.

By eliminating multiple occurrences of the same component at different levels (Fig. 2, (d)), we obtain a hierarchical decomposition/complete laminar set system on M. We conclude the proof by establishing that this decomposition has the desired properties. Let K' be any component different from the root ($K' \neq M$) and let $i+1$, $i \geq 0$, be the first level (i.e., the level with the smallest index) at which K' appears before the elimination of multiple occurrences. Since K' appears at level $i+1$, but it does not appear at level i, there must be a level i component K such that $K' \subset K$. Then, K' is a child of K in the hierarchical decomposition.

We claim that $D(K) \in (\frac{D}{\gamma^{i+1}}, \frac{D}{\gamma^i}]$. To prove the claim, we consider the level i group G_K containing K. If $D(G_K) > \frac{D}{\gamma^i}$, G_K is partitioned into level i components of diameter in $(\frac{D}{\gamma^{i+1}}, \frac{D}{\gamma^i}]$ and the claim follows. Otherwise, $G_K = K$, since we create a single level i component for each level i group of diameter at most $\frac{D}{\gamma^i}$. In addition, if $D(K) = D(G_K)$ were at most $\frac{D}{\gamma^{i+1}}$, G_K would also exist as a level $i+1$ group, and K would exist as a level $i+1$ component. This contradicts to the hypothesis that K', which is a proper subset of K, appears as a level $i+1$ component. Since $D(K) \in (\frac{D}{\gamma^{i+1}}, \frac{D}{\gamma^i}]$ and K' is a level $i+1$ component, we conclude that (i) $D(K') \leq \frac{D}{\gamma^{i+1}} < D(K)$, and (ii) either $D(K') > \frac{D}{\gamma^{i+2}} \geq \frac{D(K)}{\gamma^2}$ or $\mathrm{sep}(K') > \frac{D}{4\gamma^{i+1}} \geq \frac{D(K)}{4\gamma}$. $\square$

Lemma 6 states that every metric space has a hierarchical decomposition such that each component either is well-separated or has a large diameter. Well-separated components become isolated coalitions soon after they have become active, while large diameter components stop being active coalitions (become broken) soon after they have become active. Therefore, no coalition can become active long before it becomes either isolated or broken.

Lemma 7. *For every $\gamma \geq 12\rho$, there is a hierarchical decomposition $\mathcal{K}$ of F^* such that for any non-isolated active coalition $K \in \mathcal{K}$, $\rho D_N(K) \leq g(c_K) < (\rho+1)\gamma^2 D_N(K)$.*

Proof. We recall that $D_N(K) = \max\{D(K), \frac{1}{3\rho}\,\mathrm{sep}(K)\}$. For the lower bound, we observe that the coalition K becomes either isolated or broken as soon as $g(c_K) < \rho D_N(K)$. For the upper bound, we first observe that the root of $T_{\mathcal{K}}$, which is F^*, is an isolated coalition as long as it remains active. Hence, we can restrict our attention to the coalitions in $T_{\mathcal{K}}$ which are different from the root.

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 $g(c_K) \leq d(c_K, c_{\mathsf{p}_K}) + g(c_{\mathsf{p}_K}) < (\rho+1)D(\mathsf{p}_K)$, where the first inequality follows from the definition of the configuration distance.

Let $\mathcal{K}$ be the hierarchical decomposition of F^* implied by Lemma 6. Next, we show that for any $K \in \mathcal{K}$ different from the root F^*, $\gamma^2 D_N(K) = \gamma^2 \max\{D(K), \frac{1}{3\rho}\,\mathrm{sep}(K)\} > D(\mathsf{p}_K)$. If $D(K) > \frac{D(\mathsf{p}_K)}{\gamma^2}$, then $\gamma^2 D_N(K) \geq \gamma^2 D(K) > D(\mathsf{p}_K)$. Otherwise, by Lemma 6, it must be $\mathrm{sep}(K) > \frac{D(\mathsf{p}_K)}{4\gamma}$. Hence, $\gamma^2 D_N(K) \geq \frac{\gamma^2}{3\rho}\,\mathrm{sep}(K) > \frac{\gamma^2}{3\rho}\,\frac{D(\mathsf{p}_K)}{4\gamma} \geq D(\mathsf{p}_K)$, where the last inequality holds for any $\gamma \geq 12\rho$. □

A.5 Preliminaries

In this section, we prove several propositions which are repeatedly used in the analysis of isolated and non-isolated coalitions. In the following, we sometimes say that a facility w is mapped to an optimal center in a set of optimal centers $K \in \mathcal{K}$ instead of simply saying that w is mapped to the coalition K, because we want to also consider facilities which open either before K becomes active or after K has become broken.

Proposition 7. *Let $K \in \mathcal{K}$ be a set of optimal centers with representative c_K, and let w be a facility mapped to an optimal center in K. For every $x \geq 10$,* (A) $d(c_K, w) < \frac{2}{5}\max\{d(F_w, c_K), \lambda D(K)\}$*, and* (B) $g(c_K, w) < \max\{d(F_w, c_K), \lambda D(K)\}$.

Proof. The first inequality follows from Lemma 1 and $d(c_K, c_w) \leq D(K)$. For the second inequality, using $m(w) \leq \frac{3}{x}\,d(F_w, w)$, we show that $g(c_K, w) \leq \frac{x}{x-3}\,d(c_K, w) + \frac{3}{x-3}\,d(F_w, c_K)$. Then, the inequality follows from (A). □

Proposition 8. *Let $K \in \mathcal{K}$ be a set of optimal centers with representative c_K, and let w be an unsupported facility mapped to an optimal center in K. For every $x \geq 16$, if $d(c_K, w) \geq \lambda D(K)$, then w is merged with the first new facility which is also mapped to an optimal center in K and opens while w is still open.*

Proof. The proof is similar to the proof of Lemma 2. Using Proposition 4, Proposition 5, and $d(c_K, c_w) \leq D(K)$, we show that $m(w) \geq \frac{7}{5}d(w, c_K)$. Then, the proposition follows from Proposition 7.A. □

Proposition 9. *Let $K \in \mathcal{K}$ be a set of optimal centers with representative c_K, and let w be a facility whose neighborhood B_w includes a demand u such that $c_u \in K$ and $d^*_u < \frac{1}{\lambda}\max\{d(F_w, c_K), \lambda D(K)\}$. Then, $d(c_K, w) < \frac{3}{x}\max\{d(F_w, c_K), \lambda D(K)\}$.*

Proof. Immediate consequence of (i) $d(c_K, c_u) \leq D(K)$, (ii) $d(u, w) \leq \frac{d(F_w, w)}{x}$ because $u \in B_w$, and (iii) the upper bound on d^*_u required by the hypothesis of the proposition. □

Proposition 10. *Let $K \in \mathcal{K}$ be a set of optimal centers with representative c_K, and let w be a facility mapped to an optimal center not in K. If $d(F_w, c_K) < \frac{1}{3}\,\mathrm{sep}(K)$, then $d(c_K, w) > \frac{5}{3}\,d(F_w, c_K)$.*

Proof. Consequence of $d(c_K, c_w) \geq \mathrm{sep}(K)$, because $c_K \in K$ and $c_w \notin K$, and Lemma 1. □

Proposition 11. *Let $K \in \mathcal{K}$ be a set of optimal centers with representative c_K, and let w be a facility mapped to an optimal center not in K. Then, $d(c_K, w) \geq \frac{1}{2}\mathrm{sep}(K)$.*

Proof. It is $d(c_K, w) \geq d(c_w, w)$, because w is mapped to c_w instead of c_K, and $d(c_K, c_w) \geq \mathrm{sep}(K)$, because $c_K \in K$ and $c_w \notin K$. $\square$

Proposition 12. *Let $K \in \mathcal{K}$ be a set of optimal centers with representative c_K, and for some $\delta > 0$, let w be a facility such that $d(c_K, w) < \lambda\delta$. For every $x \geq 18$ and $\psi \leq 5$, if $m(w) < \psi(\lambda + 2)\delta$, then $g(c_K) < \rho\delta$.*

Proof. Immediate consequence of the definition of $g(c_K)$ and the choice of $\rho = (\psi + 2)(\lambda + 2)$. $\square$

Proposition 13. *For every facility w, $m(w) \leq \beta f$.*

Proof. The set $\mathrm{Init}(w)$ always contains the demand which opens the facility w and is located at the same point with w. Therefore, $|\mathrm{Init}(w) \cap \mathrm{Ball}(w, \frac{m(w)}{\psi})| \geq 1$. The proposition follows from Ineq. (1). $\square$

Proposition 14. *Let $x \geq 9$. For every demand u, $\overline{d}_u \leq 2.5\beta f$.*

Proof. The proposition follows from $m(w) \leq \beta f$, the fact that the initial assignment cost of every demand is less than βf, and the definition of the final assignment cost. $\square$

Proposition 15. *Let u be a demand currently assigned to a facility w, and let c be an optimal center in F^*. Then, $\overline{d}_u \leq d(c, u) + g(c, w)$.*

Proof. Immediate consequence of $\overline{d}_u \leq d(u, w) + \frac{x}{x-3}\, m(w)$ and the definition of $g(c, w)$. $\square$

Proposition 16. *Let u be a demand initially assigned to a facility w. For every $x \geq 9$, $\overline{d}_u \leq 4\, d(F'_u \setminus \{w\}, u)$.*

Proof. Recall that F'_u is the (updated) algorithm's facility configuration at u's assignment time. Wlog. we assume that $F'_u \setminus \{w\} \neq \emptyset$, since the bound is trivial otherwise. Let $w' \in F'_u$ be the second nearest facility to u, i.e., $d(u, w') = d(F'_u \setminus \{w\}, u)$. Since u is initially assigned to w instead of w' it must be $d(u, w) \leq d(u, w')$, which implies that $d(w, w') \leq 2\, d(u, w')$. The final assignment cost of u is $\overline{d}_u \leq d(u, w) + \frac{x}{x-3}\, m(w) \leq d(u, w') + \frac{x}{x-3}\, m(w)$.

By hypothesis, both w and w' are open at u's assignment time. If w opens before w', it must be $m(w) < d(w, w') \leq 2\, d(u, w')$, since w would have been merged with w' otherwise. Therefore, for every $x \geq 9$, $\overline{d}_u < d(u, w') + \frac{2x}{x-3}\, d(u, w') \leq 4\, d(u, w')$. If w opens after w', for every $x \geq 9$, $\frac{x}{x-3}\, m(w) \leq \frac{x}{x-3}\frac{3}{x}\, d(w', w) \leq \frac{6}{x-3}\, d(u, w') \leq d(u, w')$, because $d(F_w, w) \leq d(w', w)$, since w' opens before w and is still open at u's assignment time. Hence, $\overline{d}_u \leq d(u, w') + \frac{x}{x-3} m(w) \leq 2\, d(u, w')$. $\square$

Proposition 17. *Let K be a coalition with representative c_K, and let u be a demand mapped to K. Then, for every $x \geq 9$, $\overline{d}_u \leq 4\,(d(c_K, u) + g'_u(c_K)) \leq 4\,(d(c_K, u) + g_u(c_K))$.*

Proof. Recall that $g_u(c_K)$ denotes the configuration distance of c_K just before u arrives and $g'_u(c_K)$ denotes the configuration distance of c_K at u's assignment time (i.e., according to the updated algorithm's configuration). The second inequality follows from $g'_u(c_K) \leq g_u(c_K)$, because the configuration distance of c_K is non-increasing with time. For the first inequality, let w be the facility minimizing the configuration distance of c_K at u's assignment time (i.e., $g'_u(c_K) = g'_u(c_K, w)$). If u is initially assigned to w, using Proposition 15, we obtain that $\overline{d}_u \leq d(c_K, u) + g'_u(c_K, w) = d(c_K, u) + g'_u(c_K)$. If u is initially assigned to another facility w', Proposition 16 implies that $\overline{d}_u \leq 4\, d(u, w)$. Furthermore, $d(u, w) \leq d(c_K, u) + d(c_K, w) \leq d(c_K, u) + g'_u(c_K, w) = d(c_K, u) + g'_u(c_K)$. $\square$

A.6 Isolated Active Coalitions

Throughout this section, let K be an isolated active coalition with representative c_K. As before, we sometimes say that a facility w is mapped to an optimal center in K instead of simply saying that w is mapped to K, because we want to also consider facilities which open either before K becomes an isolated active coalition or after K has become broken/not active.

Basic Properties. We start by proving the main properties of a facility spending some time as the nearest facility to the representative c_K of an isolated active coalition K.

Proposition 18. *Let w be a facility which spends some time as the nearest facility to the representative c_K of an isolated active coalition K. Then,*

A. $d(c_K, w) < \frac{1}{3}\,\mathrm{sep}(K)$ *and w is mapped to an optimal center in K.*
B. After K has become isolated, $g(c_K, w) < \frac{1}{3}\,\mathrm{sep}(K)$*.*
C. w can be merged only with a new facility mapped to an optimal center in K.

Proof. We observe that it suffices to establish that each of the above claims holds as long as K is an isolated active coalition and w is the nearest facility to c_K. Then, (A) holds because the mapping of a facility to an optimal center does not depend on the algorithm's configuration, (B) holds because $g(c_K, w)$ is non-increasing, and (C) holds because $m(w)$ is non-increasing.
A. Let us consider any point in time such that K is an isolated coalition and w is the nearest facility to c_K. Then, it must be $d(c_K, w) = d(F, c_K) \leq g(c_K) < \frac{1}{3}\,\mathrm{sep}(K)$. The second claim follows from Proposition 11.
B. To reach a contradiction, let us assume that there is a point in time when K is an isolated active coalition, w is the nearest facility to c_K, and $g(c_K, w) \geq \frac{1}{3}\,\mathrm{sep}(K)$. Since K is an isolated active coalition, there must exist some other facility w' which is open and satisfies the following inequalities at that particular point in time.

$$d(c_K, w) \leq d(c_K, w') \leq g(c_K, w') < \tfrac{1}{3}\,\mathrm{sep}(K) \leq g(c_K, w)\,.$$

By Proposition 11, both w and w' are mapped to optimal centers in K. To establish the contradiction, we consider the moment that the latest of w, w' opens. If w' opens after w, either $d(c_K, w) \geq \lambda D(K)$, in which case $d(c_K, w') < \frac{2}{5} d(c_K, w)$ (w' is closer to c_K than w) by Proposition 7.A, or $d(c_K, w) < \lambda D(K)$, and $g(c_K) < \lambda D(K)$ (w' has made K broken) by Proposition 7.B. If w opens after w', then $g(c_K, w) < \max\{d(c_K, w'), \lambda D(K)\}$ by Proposition 7.B. Therefore, depending on whether $d(c_K, w') \geq \lambda D(K)$ or not, either $g(c_K, w) < d(c_K, w') < \frac{1}{3}\,\mathrm{sep}(K)$ or w has made K broken.
C. There is a point in time that w is open and K is an isolated active coalition. After that time, w can be merged only with a new facility at distance less than $\frac{1}{3}\,\mathrm{sep}(K)$ from c_K. The claim follows from Proposition 11. □

The following proposition establishes that we have correctly defined the final assignment cost of a demand u which is mapped to an isolated active coalition K and is currently assigned to the nearest facility to c_K.

Proposition 19. *Let u be a demand which is mapped to an isolated active coalition K and is currently assigned to a facility w. If w is currently the nearest facility to c_K, then u's actual assignment cost can never exceed* $(1 + \frac{1}{\lambda}) \max\{d(c_K, w), \lambda D(K)\} + d^*_u$*.*

Proof. As long as w is not merged with a new facility, u's actual assignment cost is $d(u, w)$ and the upper bound holds because $d(c_K, c_u) \leq D(K)$, since $c_u \in K$. If w is merged with a new facility w', by Proposition 18.C, w' must be mapped to an optimal center in K. Hence, by Proposition 7.B, $g(c_K, w') < \max\{d(c_K, w), \lambda D(K)\}$. Then, the upper bound holds because u is now assigned to w' and by Proposition 1, u's actual assignment cost can never exceed $d(u, w') + \frac{x}{x-3} m(w') \leq d^*_u + D(K) + g(c_K, w') \leq (1 + \frac{1}{\lambda}) \max\{d(c_K, w), \lambda D(K)\} + d^*_u$. □

Proposition 20. *Let w' be a new facility mapped to an isolated active coalition K. Then either w' makes K broken or* $d(c_K, w_K) \geq \lambda D(K)$, $d(c_K, w') < \frac{2}{5}\, d(c_K, w_K)$*, and the location of w_K changes to $w'_K = w'$.*

Proof. If $d(c_K, w_K) < \lambda D(K)$, Proposition 7.B implies that $g(c_K, w') < \lambda D(K)$ and w' makes K broken. Hence, if K remains active, it must be $d(c_K, w_K) \geq \lambda D(K)$. Then, by Proposition 7.A, $d(c_K, w') < \frac{2}{5}\, d(c_K, w_K)$. Therefore, w' is much closer to c_K than w_K and the location of w_K must change to $w'_K = w'$. □

Proposition 21. *Let K be an isolated active coalition, let $w_K = w$ be the nearest facility to c_K, and let w' be a new facility mapped to K. If w' does not make K broken, then w will never become again the nearest facility to any of the optimal centers in K.*

Proof. If w' does not make K broken, Proposition 20, implies that $d(c_K, w) \geq \lambda D(K)$. Moreover, $d(c_K, w') < \frac{2}{5}\, d(c_K, w)$, and the location of the nearest facility to c_K must change from $w_K = w$ to $w'_K = w'$. Similarly to the proof of Proposition 7.B, we can show that for every optimal center $c \in K$, $g(c, w') < (\frac{2}{5} + \frac{1}{\lambda} + \frac{7}{5}\,\frac{3}{x-3})\, d(c_K, w) \leq (\frac{2}{5} + \frac{1}{\lambda} + \frac{7}{5}\,\frac{3}{x-3})(1 + \frac{1}{\lambda})\, d(c, w) < d(c, w)$, where the last inequality holds for every $x \geq 11$ (recall that $\lambda = 3x + 2$). Hence, after w' opens, there will always exist a facility closer to c than w. $\square$

Lemma 8. *Let K be an isolated active coalition when a new demand u arrives. If u is not mapped to K, then neither the location of w_K nor the value of $g(c_K)$ can change.*

Proof. If a new facility w opens when u arrives, w is located at the same point with u and is also not mapped to K. Then, the value of $g(c_K)$ cannot decrease because $g(c_K) < \frac{1}{3}\,\mathrm{sep}(K)$, while $g'(c_K, w) \geq d(c_K, w) = d(c_K, u) \geq \frac{1}{2}\,\mathrm{sep}(K)$ (Proposition 11). In addition, the location of w_K cannot change, because w is not closer to c_K than w_K and the facility at the current location of w_K can only be merged with a new facility mapped to K (Proposition 18).

If no new facilities open when u arrives, the location of w_K cannot change. Next, we show that $g(c_K)$ cannot decrease because of u. Let w be the facility u is initially assigned to. Then, only the configuration distance between c_K and w can be affected by u. To reach a contradiction, we assume that after u's initial assignment to w, it becomes

$$g'(c_K, w) = d(c_K, w) + \tfrac{x}{x-3}\, m'(w) < g(c_K) < \tfrac{1}{3}\,\mathrm{sep}(K)\,.$$

Therefore, it must be $d(c_K, w) < \frac{1}{3}\,\mathrm{sep}(K)$ and $m'(w) < \frac{x-3}{3x}\,\mathrm{sep}(K)$. In addition, since $d(c_K, u) \geq \frac{1}{2}\,\mathrm{sep}(K)$, it must be $d(u, w) > 0$. The configuration distance $g(c_K, w)$ decreases only if the initial assignment of u to w causes Ineq. (1) to be violated. Then, the algorithm decreases $m(w)$ to restore the invariant. Ineq. (1) can be violated only if $d(u, w) \leq \frac{m(w)}{\psi}$ and u is included in $\mathrm{Ball}(w, \frac{m(w)}{\psi})$. The new merge radius $m'(w)$ cannot be less than $\frac{\psi(x-3)}{x} d(u, w)$, because if $m'(w) = \frac{\psi(x-3)}{x} d(u, w)$, then $\frac{m'(w)}{\psi} < d(u, w)$. Hence, u is no longer included in $\mathrm{Ball}(w, \frac{m'(w)}{\psi})$, and the invariant is restored. Therefore, $\frac{\psi(x-3)}{x} d(u, w) \leq m'(w) < \frac{x-3}{3x}\,\mathrm{sep}(K)$, which implies that $d(u, w) < \frac{1}{3\psi}\,\mathrm{sep}(K)$. Consequently, if $g'(c_K, w)$ could drop below $\frac{1}{3}\,\mathrm{sep}(K)$ because of u's initial assignment to w, it would be $d(c_K, w) < \frac{1}{3}\,\mathrm{sep}(K)$ and $d(u, w) < \frac{1}{3\psi}\,\mathrm{sep}(K)$. Therefore, for every $\psi \geq 3$, it would be $d(c_K, u) < \frac{1}{2}\,\mathrm{sep}(K)$, which is a contradiction. $\square$

Lemma 9. *Let K be an isolated active coalition, and let w_K denote the nearest facility to c_K.*

A. w_K is at least $\frac{5}{3}$ times closer to c_K than any other facility, i.e., $d(c_K, w_K) < \frac{3}{5} d(F \setminus \{w_K\}, c_K)$.
B. If $d(c_K, w_K) < \lambda D(K)$, then $d(F \setminus \{w_K\}, c_K) > \rho D(K)$.
C. Every inner demand of the isolated active coalition K which does not make K broken is initially assigned to w'_K, i.e., the nearest facility to c_K at the demand's assignment time.

Proof. Before we provide a formal proof, let us give the intuition behind this lemma. As long as K is an isolated active coalition, the location of w_K cannot change unless a new facility mapped to K opens (Lemma 8). Hence, each time the location of the nearest facility to c_K changes, either the new facility makes K broken or the new location w'_K is at least 2.5 times closer to c_K than the previous location w_K (Proposition 20). On the other hand, a facility not mapped to an optimal center in K must be at least $\frac{5}{3}$ times further from c_K than w_K (Proposition 10). Moreover, since inner demands of K are included in a small ball around c_K, every inner demand of K is initially assigned to the nearest facility to c_K.

A. Similarly to the proof of Proposition 18.B, let us assume that there is a point in time such that K is an isolated active coalition, $w = w_K$ is the nearest facility to c_K, and there is a facility $w' \in F \setminus \{w\}$ such that $d(c_K, w') \geq d(c_K, w) \geq \frac{3}{5} d(c_K, w')$. To establish the contradiction, we consider the moment that the latest of w, w' opens.

If w' opens after w and w' is mapped to an optimal center in K, then either $d(c_K, w) \geq \lambda D(K)$, in which case $d(c_K, w') < \frac{2}{5}\, d(c_K, w)$ (w' is closer to c_K than w) by Proposition 7.A, or $d(c_K, w) < \lambda D(K)$, and $g(c_K, w') < \lambda D(K)$ (w' has made K broken) by Proposition 7.B. If w' opens after w and w' is mapped to an optimal center not in K, then $d(c_K, w) < \frac{3}{5}\, d(c_K, w')$ by Proposition 10, since $d(c_K, w) < \frac{1}{3}\operatorname{sep}(K)$ by Proposition 18.A.

On the other hand, by Proposition 18.A, the facility w is mapped to an optimal center in K. Hence, if w opens after w', then either $d(c_K, w') \geq \lambda D(K)$, in which case $d(c_K, w) < \frac{2}{5}\, d(c_K, w')$ by Proposition 7.A, or $d(c_K, w') < \lambda D(K)$, and $g(c_K, w) < \lambda D(K)$ (w has made K broken) by Proposition 7.B.

B. The proof is similar to the proof of (A). The details can be found in [9].

C. Let u be an inner demand of K which does not make K broken. Let also w'_K be the nearest facility to c_K and F' be the algorithm's facility configuration at u's assignment time. A demand u mapped to an isolated active coalition K is inner if $d^*_u < \frac{1}{\lambda} \max\{d(c_K, w'_K), \lambda D(K)\}$.

If $d(c_K, w'_K) \geq \lambda D(K)$, then $d(u, w'_K) \leq d^*_u + D(K) + d(c_K, w'_K) < (1 + \frac{2}{\lambda}) d(c_K, w'_K)$, while for every other facility $w \in F' \setminus \{w'_K\}$, $d(u, w) \geq d(c_K, w) - D(K) - d^*_u > (\frac{5}{3} - \frac{2}{\lambda})\, d(c_K, w'_K) > d(u, w'_K)$, where the second inequality follows from (A), and the third inequality holds for every $x \geq 2$ and $\lambda = 3x + 2$.

If $d(c_K, w'_K) < \lambda D(K)$, then $d(u, w'_K) \leq d^*_u + D(K) + d(c_K, w'_K) < (\lambda + 2) D(K)$, while for every other facility $w \in F' \setminus \{w'_K\}$, $d(u, w) \geq d(c_K, w) - D(K) - d^*_u > ((\psi + 2)(\lambda + 2) - 2) D(K) > d(u, w'_K)$, where the second inequality follows from (B) (recall that $\rho = (\psi + 2)(\lambda + 2)$). In both cases, $d(u, w'_K) < d(u, F' \setminus \{w'_K\})$ and u is initially assigned to w'_K, i.e., the nearest facility to c_K at u's assignment time. $\square$

Lemma 10. *Let u be an outer demand of the isolated active coalition K which does not make K broken. Then, $\overline{d}_u \leq 4(\lambda + 2) d^*_u$.*

Proof. It is $d^*_u \geq \frac{1}{\lambda} \max\{d(c_K, w'_K), \lambda D(K)\}$ because u is an outer demand mapped to an isolated active coalition. If u is initially assigned to w'_K (i.e., the nearest facility to c_K at u's assignment time), u's final assignment cost is $\overline{d}_u \leq (1 + \frac{1}{\lambda}) \max\{d(c_K, w'_K), \lambda D(K)\} + d^*_u \leq (\lambda + 2)\, d^*_u$. Otherwise, let $\hat{w} \neq w'_K$ be the facility u is initially assigned to. Then, we apply Proposition 16 and obtain that $\overline{d}_u \leq 4\, d(F'_u \setminus \{\hat{w}\}, u) \leq 4\, d(u, w'_K) \leq 4(\lambda + 2)\, d^*_u$, where the second inequality follows from the fact that $w'_K \in F'_u \setminus \{\hat{w}\}$, because $\hat{w} \neq w'_K$. $\square$

Good Inner Demands. Let w be the facility which is currently the nearest facility to c_K, i.e., $w_K = w$. Let $\mathrm{G}_K(w)$ denote the set of good inner demands of K (i.e., the subset of G_K) whose initial assignment takes place while w is the nearest facility to c_K. Since $\mathrm{G}_K(w)$ is a subset of G_K, it is empty if K is either non-isolated or not active. In addition, $\mathrm{G}_K(w)$ is empty before w opens and after w has been merged with a new facility. As long as w is the nearest facility to c_K and K remains an isolated active coalition, each new inner demand mapped to K is initially assigned to w (Lemma 9.C) and is added to both G_K and $\mathrm{G}_K(w)$. After the location of the nearest facility to c_K has changed, w cannot become the nearest facility to c_K again (Proposition 21) and no new demands are added to $\mathrm{G}_K(w)$. Next, we bound the actual and the final assignment cost of the demands in $\mathrm{G}_K(w)$.

Lemma 11. *Let K be an isolated active coalition, and let w be the nearest facility to c_K (i.e., $w_K = w$). If $d(c_K, w) \geq \lambda D(K)$, then $\sum_{u \in \mathrm{G}_K(w)} d(u, w) < \beta f$ and $\sum_{u \in \mathrm{G}_K(w)} \overline{d}_u < \frac{\lambda+2}{\lambda-2}\, \beta f$.*

Proof. We consider the above sums just after a new demand is added to $\mathrm{G}_K(w)$. Let v be the demand in $\mathrm{G}_K(w)$ arriving last. Since v is added to $\mathrm{G}_K(w)$, K must be an isolated active coalition and w must be the nearest facility to c_K at v's assignment time. Therefore, from the moment the first demand is added to $\mathrm{G}_K(w)$ until the assignment time of v, no new facilities mapped to K have opened. Otherwise, by Proposition 20 and Proposition 21, either K would have become broken or the location of the nearest facility to c_K would have changed and w could not become the nearest facility to c_K again.

Since (i) all the demands in $\mathrm{G}_K(w)$ are inner demands mapped to the isolated active coalition K, (ii) their initial assignment takes place while w is the nearest facility to c_K and (iii) $d(c_K, w) \geq \lambda D(K)$, by the definition of inner demands mapped to an isolated active coalition, we obtain that for every $u \in \mathrm{G}_K(w)$, $d^*_u < \frac{1}{\lambda} d(c_K, w)$.

We first establish that just after v has been added to $\mathrm{G}_K(w)$, it is the case that $\mathrm{G}_K(w) \subseteq B_v = \mathrm{Ball}(v, r_v) \cap L$. Let u_w be the demand which is located at the same point with w and causes the algorithm to open the facility w. The demand u_w is initially assigned to w, but is does not belong to L. We first show that $u_w \notin \mathrm{G}_K(w)$ and the demands in $\mathrm{G}_K(w)$ arrive after w has opened. If u_w is not mapped to the isolated active coalition K, it does not belong to $\mathrm{G}_K(w)$ by definition. Otherwise, u_w cannot be an inner demand of K, because w is the nearest facility to c_K at u_w's assignment time and $d^*_{u_w} \geq d(c_K, w) - D(K) \geq (1 - \frac{1}{\lambda})\, d(c_K, w) \geq \frac{1}{\lambda}\, \max\{d(c_K, w), \lambda D(K)\}$.

Each new demand is added to the set of unsatisfied demands L when it arrives. Next, we show that none of the demands in $\mathrm{G}_K(w)$ can be removed from L before either K becomes broken or a new facility mapped to K opens. More specifically, for every new facility w' which opens after w and includes in its neighborhood $B_{w'}$ some demands from $\mathrm{G}_K(w)$, it must be $d(c_K, w') < \frac{3}{x} d(c_K, w) < \frac{1}{x}\, \mathrm{sep}(K)$, where the first inequality follows from Proposition 9 and the second from Proposition 18.A. Hence, w' must be mapped to an optimal center in K (Proposition 11). Since no new facilities mapped to K open until v's assignment time, it must be the case that $\mathrm{G}_K(w) \subseteq L$.

We should also prove that $\mathrm{G}_K(w) \subseteq \mathrm{Ball}(v, r_v)$. We first bound r_v from below.

$$r_v = \tfrac{d(F_v, v)}{x} = \tfrac{d(v, w)}{x} \geq \tfrac{d(c_K, w) - d(c_K, c_v) - d^*_v}{x} > \tfrac{\lambda - 2}{x} \tfrac{d(c_K, w)}{\lambda} = \tfrac{3}{\lambda}\, d(c_K, w)\,,$$

where the second equality holds because no new facility opens when v arrives and v is initially assigned to w (i.e., $d(F_v, v) = d(F'_v, v) = d(v, w)$), the fourth inequality holds because $d(c_K, c_v) \leq D(K) \leq \frac{1}{\lambda}\, d(c_K, w)$, since $c_v \in K$, and $d^*_v < \frac{1}{\lambda} d(c_K, w)$, and the last equality follows from $\lambda = 3x + 2$. On the other hand, since for every $u \in \mathrm{G}_K(w)$, $d^*_u < \frac{1}{\lambda}\, d(c_K, w)$, $\mathrm{G}_K(w)$ has a small diameter. In particular, for every $u \in \mathrm{G}_K(w)$, $d(u, v) \leq d(u, c_u) + d(c_u, c_v) + d(c_v, v) < \frac{3}{\lambda}\, d(c_K, w) < r_v$. Therefore, $\mathrm{Ball}(v, r_v)$ includes every demand in $\mathrm{G}_K(w)$, and $\mathrm{G}_K(w) \subseteq B_v$.

Every $u \in \mathrm{G}_K(w)$ is added to the set of unsatisfied demands L with a potential equal to its initial assignment cost $d(u, w)$. This potential cannot increase because w remains open at v's assignment time. Next, we prove that the potential of each $u \in \mathrm{G}_K(w)$ cannot decrease as long as K remains an isolated active coalition and w is the nearest facility to c_K. Let u be a demand in $\mathrm{G}_K(w)$ (including v). By Lemma 9.A, when v arrives, it is the case that $d(c_K, w) < \frac{3}{5}\, d(F_v \setminus \{w\}, c_K)$. Similarly to the proof of Lemma 9.C, case $d(c_K, w) \geq \lambda D(K)$, we can prove that $d(F_v \setminus \{w\}, u) \geq d(F_v \setminus \{w\}, c_K) - D(K) - d^*_u > (\frac{5}{3} - \frac{2}{\lambda})\, d(c_K, w) > d(u, w)$. Therefore, when the demand v arrives, for every $u \in \mathrm{G}_K(w)$, $d(F_v, u) = d(u, w)$.

The total potential in v's neighborhood B_v must be less than βf, because v is initially assigned to w instead of opening a new facility and being initially assigned there. Consequently,

$$\beta f > \mathrm{Pot}(B_v) = \textstyle\sum_{u \in B_v} d(F_v, u) \geq \sum_{u \in \mathrm{G}_K(w)} d(F_v, u) = \sum_{u \in \mathrm{G}_K(w)} d(u, w)\,,$$

where the third inequality follows from $\mathrm{G}_K(w) \subseteq B_v$, and the fourth inequality from the fact that the potential of each $u \in \mathrm{G}_K(w)$ remains equal to $d(u, w)$ as long as K remains an isolated active coalition and w remains the nearest facility to c_K. This concludes the proof of the first part of the lemma.

As for the final assignment cost, for every $u \in \mathrm{G}_K(w)$, we bound $\overline{d}_u$ using Proposition 19:

$$\overline{d}_u \leq (1 + \tfrac{1}{\lambda}) \max\{d(c_K, w), \lambda D(K)\} + d^*_u < \tfrac{\lambda + 2}{\lambda}\, d(c_K, w)\,.$$

We conclude the proof by observing that for every $u \in \mathrm{G}_K(w)$, $d(u, w) > \frac{\lambda}{\lambda - 2}\, d(c_K, w)$. □

Lemma 12. *Let K be an isolated active coalition, and let w be the nearest facility to c_K. If $d(c_K, w) < \lambda D(K)$, for every $x \geq 18$ and $\frac{10}{3} \leq \psi \leq 5$, $\sum_{u \in \mathrm{G}_K(w)} d(u, w) < \frac{\beta}{\psi}\, f$ and $\sum_{u \in \mathrm{G}_K(w)} \overline{d}_u < \frac{3}{2}\, \beta f$.*

Proof. By the definition of the set $\mathrm{G}_K(w)$ (see also the claims (i) and (ii) at the beginning of the proof of Lemma 11) and since $d(c_K, w) < \lambda D(K)$, for every $u \in \mathrm{G}_K(w)$, it must be $d^*_u < 2D(K)$. As in the proof of Lemma 11, we consider the above sums just after a new demand is added to $\mathrm{G}_K(w)$. Let v be the demand in $\mathrm{G}_K(w)$ arriving last. Since v is added to $\mathrm{G}_K(w)$, K must be an isolated active coalition at v's assignment time. Therefore, it must be $m(w) \geq \psi(\lambda+2)D(K)$, because K would have become broken otherwise (Proposition 12, for $\delta = D(K)$, $x \geq 18$, and $\psi \leq 5$). Using Ineq. (1), we obtain that:

$$|\mathrm{Init}(w) \cap \mathrm{Ball}(w, (\lambda+2)D(K))| \cdot \psi(\lambda+2)D(K) \leq \beta f \tag{8}$$

We show that as long as w is the nearest facility to c_K and K remains an isolated active coalition, $\mathrm{G}_K(w) \subseteq \mathrm{Init}(w) \cap \mathrm{Ball}(w, (\lambda+2)D(K))$. Since all the demands in $\mathrm{G}_K(w)$ are initially assigned to w (Lemma 9.C), $\mathrm{G}_K(w) \subseteq \mathrm{Init}(w)$. In addition, for every $u \in \mathrm{G}_K(w)$, $d(u, w) \leq d^*_u + D(K) + d(w, c_K) < (\lambda+2)D(K)$. Hence, $\mathrm{G}_K(w) \subseteq \mathrm{Ball}(w, (\lambda+2)D(K))$. Combining $\mathrm{G}_K(w) \subseteq \mathrm{Init}(w) \cap \mathrm{Ball}(w, (\lambda+2)D(K))$ with Ineq. (8), we conclude that

$$\textstyle\sum_{u \in \mathrm{G}_K(w)} d(u, w) < |\mathrm{G}_K(w)| \cdot (\lambda+2)D(K) \leq \frac{\beta}{\psi} f\,.$$

As for the final assignment cost, we have shown that as long as w is the nearest facility to c_K and K remains an isolated active coalition, $\mathrm{G}_K(w) \subseteq \mathrm{Init}(w) \cap \mathrm{Ball}(w, \frac{m(w)}{\psi})$. Hence, by Ineq. (1). , $|\mathrm{G}_K(w)| \cdot m(w) \leq \beta f$. By the definition of the final assignment cost, we obtain that

$$\sum_{u \in \mathrm{G}_K(w)} \overline{d}_u \leq \sum_{u \in \mathrm{G}_K(w)} d(u, w) + \tfrac{x}{x-3} \sum_{u \in \mathrm{G}_K(w)} m(w) \leq (\tfrac{1}{\psi} + \tfrac{x}{x-3})\,\beta f \leq \tfrac{3}{2}\beta f\,,$$

where the last inequality holds for every $\psi \geq \frac{10}{3}$ and $x \geq 18$. □

The proof of Lemma 3. We prove that the total actual assignment cost of the good demands of K is

$$\textstyle\sum_{u \in \mathrm{G}_K} d(u, w_K) < 2\beta f + 3 \sum_{u \in \mathrm{G}_K} d^*_u \tag{9}$$

and the total final assignment cost (according to the current algorithm's configuration) of the good demands of K is

$$\textstyle\sum_{u \in \mathrm{G}_K} \overline{d}_u < 4.5\beta f + 7 \sum_{u \in \mathrm{G}_K} d^*_u \tag{10}$$

We first prove Ineq. (9) by induction and then derive Ineq. (10) from Ineq. (9). Ineq. (9) is trivially true while $\mathrm{G}_K = \emptyset$. By definition, the set of good demands G_K is empty as long as K is either non-isolated or not active. In addition, G_K becomes empty every time the location of the nearest facility to c_K changes without the facility at the previous location w_K being merged with the new facility at the new location w'_K. We inductively assume that the inequality holds just before the current location w_K of the nearest facility to c_K changes. We show that the inequality remains valid until either the location of w_K changes again or K becomes broken.

Let w be the facility at the current location of w_K i.e. $w_K = w$. By Lemma 8, the location of w_K cannot change unless a new facility mapped to K opens. Let w' be the next facility mapped to K. Let also G_K be the set of good demands just before w' opens. We inductively assume that Ineq. (9) holds just before w' opens. Wlog. we assume that w' does not make K broken, since as long as K is not active/broken, $\mathrm{G}_K = \emptyset$ and Ineq. (9) holds trivially. Therefore, by Proposition 20, it must be (i) $d(c_K, w) \geq \lambda D(K)$, (ii) $d(c_K, w') < \frac{2}{5}\, d(c_K, w)$, and (iii) the location of the nearest facility to c_K changes from $w_K = w$ to $w'_K = w'$.

If w is not merged with w', the set of good demands becomes empty. Hence, $\sum_{u \in \mathrm{G}_K} d(u, w') = 0$ just after w' opens. If w is merged with w', we show that just after w has been merged with w', it is the case that $\sum_{u \in \mathrm{G}_K} d(u, w') < \beta f + 3 \sum_{u \in \mathrm{G}_K} d^*_u$. For every $u \in \mathrm{G}_K$, we bound $d(u, w')$ from above in terms of $d(u, w)$ and d^*_u.

$$\begin{aligned}
d(u, w') &\leq d^*_u + D(K) + d(c_K, w') && u \text{ is mapped to } c_u \in K \\
&\leq d^*_u + D(K) + \tfrac{2}{5}\, d(c_K, w) && \text{Proposit. 7.A, } w' \text{ is mapped to } K,\ d(c_K, w) \geq \lambda D(K) \\
&\leq d^*_u + (\tfrac{1}{\lambda} + \tfrac{2}{5}) d(c_K, w) && d(c_K, w) \geq \lambda D(K) \\
&\leq d^*_u + (\tfrac{1}{\lambda} + \tfrac{2}{5}) \tfrac{\lambda}{\lambda-1} [d^*_u + d(u, w)] && d(c_K, w) \geq \lambda D(K) \Rightarrow d(c_K, w) \leq \tfrac{\lambda}{\lambda-1} [d^*_u + d(u, w)] \\
&\leq \tfrac{1}{2}\, d(u, w) + \tfrac{3}{2}\, d^*_u && \text{for every } x \geq 5 \text{ and } \lambda = 3x+2
\end{aligned}$$

Before the initial assignment of the demand which causes w' to open, the set of good demands of K is G_K, i.e., exactly the same with the set of good demands of K just before w' opens. Using the previous bound on $d(u, w')$ and the inductive hypothesis, we conclude that just after w has been merged with w', it is the case that

$$\sum_{u \in \mathrm{G}_K} d(u, w') \leq \frac{1}{2} \sum_{u \in \mathrm{G}_K} d(u, w) + \frac{3}{2} \sum_{u \in \mathrm{G}_K} d_u^* < \frac{1}{2}\left[2\beta f + 3 \sum_{u \in \mathrm{G}_K} d_u^*\right] + \frac{3}{2} \sum_{u \in \mathrm{G}_K} d_u^* = \beta f + 3 \sum_{u \in \mathrm{G}_K} d_u^*$$

Let G'_K be the set of good demands of K just before either K becomes broken or a new facility mapped to K opens and the location of the nearest facility to c_K changes again. By the definition of the set $\mathrm{G}_K(w')$ (i.e., the set of good demands of K whose initial assignment takes place as long as w' is the nearest facility to c_K), every inner demand added to G'_K after w' has opened is also added to $\mathrm{G}_K(w')$. Therefore, $\mathrm{G}'_K = \mathrm{G}_K \cup \mathrm{G}_K(w')$. By Lemma 11 and Lemma 12, we know that $\sum_{u \in \mathrm{G}_K(w')} d(u, w') < \beta f$. Hence, as long as w' is the nearest facility to c_K and K remains an isolated active coalition, it is the case that

$$\sum_{u \in \mathrm{G}'_K} d(u, w') = \sum_{u \in \mathrm{G}_K} d(u, w') + \sum_{u \in \mathrm{G}_K(w')} d(u, w') < 2\beta f + 3 \sum_{u \in \mathrm{G}_K} d_u^* \,.$$

This concludes the proof of Ineq. (9). We proceed to establish Ineq. (10).

As before, let w' be the nearest facility to c_K (i.e., $w_K = w'$), and let G'_K be the current set of good inner demands of K. We first consider the case that $d(c_K, w') \geq \lambda D(K)$. Since every demand $u \in \mathrm{G}'_K$ is mapped to the isolated active coalition K and is currently assigned to the nearest facility to c_K, we can bound the final assignment cost of u using the upper bound of Proposition 19:

$$\begin{array}{ll}
\overline{d}_u \leq (1 + \frac{1}{\lambda}) \max\{d(c_K, w'), \lambda D(K)\} + d_u^* & \\
\quad \leq \frac{\lambda+1}{\lambda}\, d(c_K, w') + d_u^* & \text{by the assumption that } d(c_K, w') \geq \lambda D(K) \\
\quad \leq \frac{\lambda+1}{\lambda} \frac{\lambda}{\lambda-1}\, [d_u^* + d(u, w')] + d_u^* & d(c_K, w') \leq \frac{\lambda}{\lambda-1}[d_u^* + d(u, w')] \\
\quad \leq \frac{3}{2}\, d(u, w') + \frac{5}{2}\, d_u^* & \text{for every } x \geq 1 \text{ and } \lambda = 3x + 2
\end{array}$$

Using Ineq. (9), we conclude that

$$\textstyle\sum_{u \in \mathrm{G}'_K} \overline{d}_u < 3\beta f + 7 \sum_{u \in \mathrm{G}'_K} d_u^* \tag{11}$$

We have also to consider the case that $d(c_K, w') < \lambda D(K)$. As before, let G'_K denote the current set of good demands of K, and let G_K be the set of good demands of K just after w' opens. By the definition of the set $\mathrm{G}_K(w')$ (i.e., the set of good demands of K whose initial assignment takes place as long as w' is the nearest facility to c_K), $\mathrm{G}'_K = \mathrm{G}_K \cup \mathrm{G}_K(w')$. By Lemma 12, it is the case that $\sum_{u \in \mathrm{G}_K(w')} \overline{d}_u \leq \frac{3}{2}\beta f$.

We should also bound the final assignment cost of the demands in G_K according to the current algorithm's configuration. Wlog. we can assume that the set of good demands of K is non-empty just after w' opens ($\mathrm{G}_K \neq \emptyset$), since there is nothing to bound otherwise. Let w be the nearest facility to c_K just before w' opens. By Proposition 18.A, both w and w' are mapped to optimal centers in K. Hence, it must be $d(c_K, w) \geq \lambda D(K)$, since w' would have made K broken otherwise (Proposition 7.B). Since we have assumed that $\mathrm{G}_K \neq \emptyset$, the facility w must have been merged with w' and every demand $u \in \mathrm{G}_K$ was assigned to w before w' opens. Consequently, by Ineq. (9), $\sum_{u \in \mathrm{G}_K} d(u, w) < 2\beta f + 3 \sum_{u \in \mathrm{G}_K} d_u^*$.

By the upper bound of Proposition 19, the final assignment cost of every demand $u \in \mathrm{G}_K$ according to the current algorithm's configuration is $\overline{d}_u < (\lambda + 1) D(K) + d_u^*$, because u is currently assigned to w' and $d(c_K, w') < \lambda D(K)$. Since $d(c_K, w) \geq \lambda D(K)$, similarly to the proof of Ineq. (11), we obtain that for every $u \in \mathrm{G}_K$,

$$\overline{d}_u < (\lambda + 1) D(K) + d_u^* \leq (1 + \tfrac{1}{\lambda}) \max\{d(c_K, w), \lambda D(K)\} + d_u^* \leq \tfrac{3}{2}\, d(u, w) + \tfrac{5}{2}\, d_u^* \,.$$

Therefore, Ineq. (11) also holds for the final assignment cost of the demands in G_K according to the current algorithm's configuration. We conclude the proof of the lemma by applying Lemma 12 for the demands in $\mathrm{G}_K(w')$ and Inequality (11) for the demands in G_K:

$$\sum_{u\in \mathrm{G}'_K} \bar{d}_u = \sum_{u\in \mathrm{G}_K} \bar{d}_u + \sum_{u\in \mathrm{G}_K(w')} \bar{d}_u < 3\beta f + 7\sum_{u\in \mathrm{G}_K} d^*_u + 1.5\beta f \leq 4.5\beta f + 7\sum_{u\in \mathrm{G}'_K} d^*_u\,.$$

A.7 Non-Isolated Active Coalitions

Throughout this section, let K be a non-isolated active coalition with representative c_K.

The proof of Lemma 4. Let $x \geq 18$ and $3 \leq \psi \leq 5$. We also recall that $D_N(K) = \max\{D(K), \frac{1}{3\rho}\,\mathrm{sep}(K)\}$ and $\mathrm{N}_K = \mathrm{in}_N(K) \cap L$. We want to prove that $|\mathrm{N}_K| \cdot g(c_K) \leq (\psi+4)\gamma^2\beta f$. Since $g(c_K)$ is non-increasing with time, the product $|\mathrm{N}_K| \cdot g(c_K)$ can increase only if a new demand is added to N_K. Therefore, it suffices to establish the inequality just after a new demand is added to N_K.

A new demand v is added to N_K if (i) v is an inner demand mapped to the non-isolated active coalition K (i.e., $v \in \mathrm{in}_N(K)$), (ii) no new facilities open when v arrives, and v is not removed from the set of unsatisfied demands, and (iii) v does not make the coalition K either isolated or broken (i.e., $g'(c_K) \geq \rho D_N(K)$). We recall that a demand v mapped to a non-isolated active coalition is inner if $d^*_v < D_N(K)$. Since $\mathrm{N}_K \subseteq \mathrm{in}_N(K)$, for every $u \in \mathrm{N}_K$, $d(c_K, u) < 2D_N(K)$ and the diameter of the sets N_K and $\mathrm{in}_N(K)$ is less than $3D_N(K)$.

Let v be the last demand added to N_K. Let also $\mathrm{N}_K/\mathrm{N}'_K = \mathrm{N}_K \cup \{v\}$ be the set of unsatisfied inner demands of K before/after v. As usual, we use plain symbols to refer to the algorithm's configuration at v's arrival time and primed symbols to refer to the updated algorithm's configuration at v's assignment time. We first consider the case that when v arrives, $d(F_v, c_K) \geq \lambda D_N(K)$. Consequently, $d(F_v, v) > (\lambda-2)D_N(K)$, and since $\lambda = 3x+2$, $r_v = \frac{d(F_v,v)}{x} > 3D_N(K)$. Hence, $\mathrm{Ball}(v, r_v)$ includes every demand in $\mathrm{in}_N(K)$, and $B_v = \mathrm{Ball}(v, r_v) \cap L$ includes every demand in N'_K. Since no new facilities open when v arrives, it must be

$$\beta f > \mathrm{Pot}(B_v) = \textstyle\sum_{u\in B_v} d(F_v, u) \geq \sum_{u\in \mathrm{N}'_K} d(F_v, u) > |N'_K|(\lambda-2)D_N(K)\,,$$

where the last inequality holds because for every $u \in \mathrm{N}'_K$, $d(F_v, u) > (\lambda-2)D_N(K)$, since $d(c_K, u) < 2D_N(K)$ and $d(F_v, c_K) \geq \lambda D_N(K)$. Since K is an active coalition when v arrives, it must be $g'(c_K) \leq g(c_K) < (\rho+1)\gamma^2 D_N(K)$ (Lemma 7). Therefore,

$$|\mathrm{N}'_K| \cdot g'(c_K) < \tfrac{((\psi+2)(\lambda+2)+1)\gamma^2}{\lambda-2}\,\beta f \leq (\psi+4)\gamma^2\beta f\,,$$

where the last inequality holds for $\lambda = 3x+2$, $\rho = (\psi+2)(\lambda+2)$, and $\psi \leq \frac{3x-4}{2}$.

We have also to consider the case that $d(F_v, c_K) < \lambda D_N(K)$. Let w be the nearest facility to c_K when v arrives. We will show that

$$\mathrm{N}'_K \subseteq \mathrm{Init}'(w) \cap \mathrm{Ball}(w, \tfrac{m'(w)}{\psi}) \tag{12}$$

Before establishing (12), we show that it indeed implies the lemma. We first observe that $m'(w) \geq \psi(\lambda+2)D_N(K)$ (Proposition 12 for $\delta = D_N(K)$, $x \geq 18$, and $\psi \leq 5$), because K remains a non-isolated active coalition after v. Therefore, $d(c_K, w) < \lambda D_N(K) < \frac{1}{\psi}m'(w)$. Since we have assumed that Ineq. (12) holds, by Ineq. (1), it must be $|\mathrm{N}'_K| \cdot m'(w) \leq \beta f$. Therefore,

$$\begin{aligned}|\mathrm{N}'_K| \cdot g'(c_K) &\leq |\mathrm{N}'_K| \cdot g'(c_K, w) \leq |\mathrm{N}'_K|(d(c_K, w) + \tfrac{x}{x-3}\,m'(w))\\ &< |\mathrm{N}'_K|(\tfrac{1}{\psi} + \tfrac{x}{x-3})m'(w) \leq \tfrac{3}{2}\,\beta f < (\psi+4)\gamma^2\beta f\,.\end{aligned}$$

We should also prove Ineq. (12). For all $u \in \mathrm{in}_N(K)$, $d(u, w) < (\lambda+2)D_N(K)$, because $d(c_K, w) < \lambda D_N(K)$. Since $m'(w) \geq \psi(\lambda+2)D_N(K)$, we obtain that every demand in N'_K is also included in $\mathrm{Ball}(w, \frac{m'(w)}{\psi})$.

Next, we show that every demand in N'_K is initially assigned to w. We first observe that when w opens, it must be $d(F_w, w) \geq x(\lambda+2)D_N(K)$. Otherwise, it would be

$$m(w) \leq \tfrac{3}{x} d(F_w, w) < 3(\lambda+2)D_N(K) \leq \psi(\lambda+2)D_N(K)\,,$$

and w would have made K either isolated or broken. Since $d(F_w, w) \geq x(\lambda+2)D_N(K)$ and $d(c_K, w) < \lambda D_N(K)$, $\mathrm{Ball}(w, \frac{d(F_w,w)}{x})$ includes every demand in $\mathrm{in}_N(K)$. Thus, w's neighborhood $B_w = \mathrm{Ball}(w, \frac{d(F_w,w)}{x}) \cap L$ includes every inner demand of K which is unsatisfied when w opens. Therefore, the set of unsatisfied inner demands of K becomes empty when w opens, and the demands currently in N'_K arrive after w's opening.

In addition, for every facility w' opening after w and before v, it must be $d(c_K, w') \geq 2(\lambda+2)D_N(K)$. Otherwise, w' would have made K either isolated or broken, because

$$\begin{aligned} g(c_K, w') &= d(c_K, w') + \tfrac{x}{x-3} m(w') < 2(\lambda+2)D_N(K) + \tfrac{x}{x-3}\tfrac{3}{x}\, d(w, w') \\ &\leq 2(\lambda+2)D_N(K) + \tfrac{3}{x-3}\left(d(c_K, w') + d(c_K, w)\right) < (2 + \tfrac{9}{x-3})(\lambda+2)D_N(K) \\ &= \tfrac{2x+3}{x-3}(\lambda+2)D_N(K) \leq (\psi+2)(\lambda+2)D_N(K)\,, \end{aligned}$$

Hence, between w's opening time and v's assignment time, it is the case that $d(F \setminus \{w\}, c_K) \geq 2(\lambda+2)D_N(K)$. Therefore, for every inner demand u which is mapped to K and arrives after w's opening and before v's assignment (including v) it must be $d(u, w) < (\lambda+2)D_N(K) \leq (2\lambda+2)D_N(K) \leq d(F \setminus \{w\}, u)$. In other words, every inner demand which is mapped to K and arrives after w's opening and before v's assignment is initially assigned to w and is added to $\mathrm{Init}(w)$. Since the set of unsatisfied inner demands of K becomes empty when w opens, at v's assignment time, it is the case that $\mathrm{N}'_K \subseteq \mathrm{Init}'(w)$. Combining this with $\mathrm{N}'_K \subseteq \mathrm{Ball}(w, \frac{m'(w)}{\psi})$, we obtain (12). □

Proposition 22. *Let K be a non-isolated active coalition and let w be a new facility such that $B_w \cap \mathrm{in}_N(K) \neq \emptyset$. Then, for every $x \geq 18$, $g'(c_K, w) < \frac{1}{3}\, g(c_K)$.*

Proof. Let u be a demand which belongs to both B_w and $\mathrm{in}_N(K)$. Therefore, $d(u, w) \leq \frac{d(F_w,w)}{x}$ and $d(c_K, u) < 2D_N(K)$. Using these inequalities, we obtain that

$$d(F_w, w) \leq d(F_w, c_K) + d(c_K, u) + d(u, w) < d(F_w, c_K) + 2D_N(K) + \tfrac{d(F_w,w)}{x}\,,$$

which implies that $d(F_w, w) < \frac{x}{x-1} d(F_w, c_K) + \frac{2x}{x-1} D_N(K)$. After w opens, the configuration distance between c_K and w becomes

$$\begin{aligned} g'(c_K, w) &= d(c_K, w) + \tfrac{x}{x-3} m(w) \\ &\leq d(c_K, u) + d(u, w) + \tfrac{x}{x-3}\tfrac{3}{x} d(F_w, w) \\ &< 2D_N(K) + (\tfrac{1}{x} + \tfrac{3}{x-3}) d(F_w, w) && d(u, w) \leq \tfrac{d(F_w,w)}{x} \\ &\leq 2D_N(K) + \tfrac{5}{x} d(F_w, w) && \text{for every } x \geq 12 \\ &< (2 + \tfrac{10}{x-1})D_N(K) + \tfrac{5}{x-1} d(F_w, c_K) && d(F_w, w) < \tfrac{x}{x-1} d(F_w, c_K) + \tfrac{2x}{x-1} D_N(K) \\ &\leq 3D_N(K) + \tfrac{5}{x-1} d(F_w, c_K) && \text{for every } x \geq 11 \\ &\leq (\tfrac{3}{\rho} + \tfrac{5}{x-1}) g(c_K) && g(c_K) \geq \rho D_N(K) \text{ and } g(c_K) \geq d(F_w, c_K) \\ &\leq \tfrac{1}{3}\, g(c_K) && x \geq 18,\ \lambda = 3x+2, \text{ and } \rho = (\psi+2)(\lambda+2) \end{aligned}$$

Lemma 13. *Let u be an outer demand mapped to a non-isolated active coalition K. Then, $\overline{d}_u \leq 4[(\rho+1)\gamma^2 + 2]\, d^*_u$.*

Proof. Applying Proposition 17, we obtain that

$$\overline{d}_u \leq 4\left[d(c_K, u) + g_u(c_K)\right] < 4\left[d^*_u + D(K) + (\rho+1)\gamma^2 D_N(K)\right] \leq 4((\rho+1)\gamma^2 + 2)\, d^*_u\,,$$

where the second inequality follows from $g_u(c_K) < (\rho+1)\gamma^2 D_N(K)$, since K is a non-isolated active coalition when u arrives (see also Lemma 7, Section A.4), and the third inequality from $d^*_u \geq D_N(K) \geq D(K)$, because u is an outer demand mapped to the non-isolated active coalition K. □

Lemma 14. *Let u be an inner demand mapped to a non-isolated active coalition K. If u does not make K either isolated or broken, then $\overline{d}_u \leq 5\, g'_u(c_K)$.*

Proof. We recall that $g'_u(c_K)$ denotes the configuration distance of c_K at u's assignment time (i.e., according to the updated algorithm's configuration). Applying Proposition 17, we obtain that

$$\overline{d}_u \leq 4\,[d^*_u + D(K) + g'_u(c_K)] \leq 4(\tfrac{2}{\rho} + 1)\, g'_u(c_K) < 5\, g'_u(c_K)\,.$$

The second inequality follows from (i) $d^*_u < D_N(K)$, because u is an inner demand mapped to the non-isolated active coalition K, (ii) $\rho D_N(K) \leq g'_u(c_K)$, since K remains a non-isolated active coalition after u, and (iii) $D(K) \leq D_N(K)$. The last inequality holds because $\rho = (\psi+2)(\lambda+2) > 8$. ◻

A.8 The Potential Function Argument

In this section, we develop a formal potential function argument establishing that the actual assignment cost of IFL is within a constant factor from the optimal cost. Let $\mathcal{K}$ be the hierarchical decomposition of F^* implied by Lemma 7. Then, for every non-isolated active coalition $K \in \mathcal{K}$, $\rho D_N(K) \leq g(c_K) < (\rho+1)\gamma^2 D_N(K)$.

We recall that with the exception of good demands, each new demand is *irrevocably* charged with its final assignment cost at its assignment time. A good demand is charged with its actual assignment cost, which is always equal to its distance from the nearest facility to the representative of the isolated active coalition the demand is mapped to. Good demands are *irrevocably* charged with their final assignment cost at the moment they become bad. The assignment cost of a demand is always charged to the active coalition the demand is mapped to. We use Asg_K to denote the assignment cost the algorithm has been charged with for the demands mapped to the coalition K.

We use the following potential function to bound the algorithm's assignment cost.

$$\Phi = \textstyle\sum_{K\in\mathcal{K}} \Phi_K\,, \text{ where } \Phi_K = \Xi^{(1)}_K + \Xi^{(2)}_K - \Upsilon^{(N)}_K - \Upsilon^{(I)}_K\,.$$

The functions $\Xi^{(1)}_K$ and $\Xi^{(2)}_K$ are defined as:

$$\Xi^{(1)}_K = \begin{cases} [5(\psi+4)\gamma^2 + 9.5]\,\beta f & \text{if } g(c_K) \geq \rho D_N(K) = \rho\max\{D(K), \frac{1}{3\rho}\,\mathrm{sep}(K)\} \\ & \text{(K is a non-isolated coalition).} \\ 7\beta f & \text{if } \rho D(K) \leq g(c_K) < \frac{1}{3}\,\mathrm{sep}(K) \text{ (K is an isolated coalition).} \\ 0 & \text{if } g(c_K) < \rho D(K) \text{ (K is not active/broken).} \end{cases}$$

$$\Xi^{(2)}_K = [5(\psi+4)\gamma^2 + 2.5]\,\beta f \max\left\{\ln\left(\frac{\min\{g(c_K), (\rho+1)\gamma^2 D_N(K)\}}{\rho D_N(K)}\right), 0\right\}$$

In addition, the functions $\Upsilon^{(N)}_K$ and $\Upsilon^{(I)}_K$ are defined as:

$$\Upsilon^{(N)}_K = 5\,|\mathrm{N}_K| \cdot g(c_K) \quad\text{and}\quad \Upsilon^{(I)}_K = \textstyle\sum_{u\in \mathrm{G}_K} (d(u, w_K) - 7\, d^*_u)\,.$$

The functions $\Xi^{(1)}_K$ and $\Xi^{(2)}_K$ hold the credit given to the coalition $K \in \mathcal{K}$. The credit held by $\Xi^{(1)}_K + \Xi^{(2)}_K$ compensates for the final assignment cost of the inner demands of K which arrive before K becomes isolated. When K becomes isolated, there is a credit of $7\beta f$ remaining in $\Xi^{(1)}_K$. This credit absorbs the actual assignment cost of good demands and eventually compensates for the final assignment cost of the good demands becoming bad when K becomes broken. The function $\Upsilon^{(N)}_K$ accounts for the part of the final assignment cost of the demands in N_K which has not been charged to $\Xi^{(2)}_K$ yet. By the definition of N_K, $\Upsilon^{(N)}_K = 0$ while K is either isolated or not active. As for the function $\Upsilon^{(I)}_K$, the quantity $\sum_{u\in \mathrm{G}_K} d(u, w_K)$ is always equal to the actual assignment cost of the demands in G_K, i.e., the

Isolated Active Coalition K:

- Outer demand u: $\overline{d}_u \le 4(\lambda+2)d_u^*$ (Lemma 10).
- Inner demands:
  1. They initially become good (Lemma 9) and $\mathrm{Asg}(\mathrm{G}_K) \le 2\beta f + 3\mathrm{Asg}^*(\mathrm{G}_K)$ (Lemma 3).
  2. When they turn into bad demands, $\sum_{u\in \mathrm{G}_K} \overline{d}_u \le 4.5\beta f + 7\mathrm{Asg}^*(\mathrm{G}_K)$ (Lemma 3).
  3. Lemma 8 and Proposition 18: Good demands turn into bad demands only if
     (a) K becomes broken: Charge K with $4.5\beta f + \overline{d}_u \le 7\beta f$ (Proposition 14).
     (b) The current nearest facility to c_K, denoted by w, is not merged with a new facility w' mapped to K (w' becomes the nearest facility to c_K, Proposition 20).
     w must be a supported facility (Proposition 20 and Proposition 8).
     $\mathrm{Asg}^*(B_w) \ge \frac{1}{3x}\beta f$ is charged with $4.5\beta f$. Each B_w is charged at most once (Proposition 21).

Non-Isolated Active Coalition K:

- K's credit $= (5(\psi+4)\gamma^2+2.5)(\ln(\frac{g(c_K)}{\rho D_N(K)})+1)\beta f$.
  Initially, K's credit $\le (5(\psi+4)\gamma^2+2.5)(\ln(\frac{(\rho+1)\gamma^2}{\rho})+1)\beta f$ (Lemma 7).
- Unsatisfied inner demands $\mathrm{N}_K = \mathrm{in}_N(K) \cap L$.
  Function $-\Upsilon_K^{(N)} = -5\,|\mathrm{N}_K| \cdot g(c_K)$ accounts for the final assignment cost of N_K.
  $\Upsilon_K^{(N)} \le 5(\psi+4)\gamma^2\beta f$ (Lemma 4).
- $g(c_K)$ decreases by a factor of $\alpha > 1$: the decrease in K's credit compensates for the increase in $-\Upsilon_K^{(N)}$ (for every $\alpha \ge 1$, $\ln(\alpha) \ge (1-\frac{1}{\alpha})$).
- Demand u does not make K isolated or broken and if a new facility w' opens, then $B_{w'} \cap \mathrm{in}_N(K) = \emptyset$.
  1. Outer demand u: $\overline{d}_u \le 4[(\rho+1)\gamma^2+2]\,d_u^*$ (Lemma 13).
  2. Inner demand u: $\overline{d}_u \le 5\,g'_u(c_K)$ (Lemma 14). Function $-\Upsilon_K^{(N)}$ compensates for $\overline{d}_u$.
- Demand u either makes K isolated or broken or opens a new facility w' such that $B_{w'} \cap \mathrm{in}_N(K) \ne \emptyset$.
  1. $\overline{d}_u \le 2.5\beta f$ (Proposition 14).
  2. $-\Upsilon_K^{(N)}$ increases by at most $5(\psi+4)\gamma^2\beta f$.
  3. K's credit decreases by at least $(5(\psi+4)\gamma^2+2.5)\beta f$ (if $B_{w'} \cap \mathrm{in}_N(K) \ne \emptyset$, then $g'(c_K) < \frac{1}{3}g(c_K)$, Proposition 22).

Fig. 3. A sketch of the potential function argument.

set of good inner demands of K. By the definition of G_K, $\Upsilon_K^{(I)} = 0$ while K is either non-isolated or not active. Fig. 3 provides a brief sketch of the potential function argument.

In the following, we use plain symbols to denote the value of the potential function and its components at the arrival time of a new demand and primed symbols to denote the value of the potential function at the assignment time of the demand. In addition, for a coalition K, let $\Delta\Phi_K = \Phi'_K - \Phi_K$ denote the change in the value of the potential function Φ_K, and let $\Delta\mathrm{Asg}_K = \mathrm{Asg}'_K - \mathrm{Asg}_K$ denote the difference in the assignment cost charged to K.

Lemma 15. *For every $K \in \mathcal{K}$, Φ_K is always non-negative.*

Proof. As long as K is a non-isolated coalition, $\Xi_K^{(1)} = [5(\psi+4)\gamma^2+9.5]\,\beta f$, while $\Upsilon_K^{(I)} = 0$, and $\Upsilon_K^{(N)}$ does not exceed $5(\psi+4)\gamma^2\beta f$ (Lemma 4). As long as K is an isolated coalition, $\Xi_K^{(1)} = 7\beta f$, while $\Upsilon_K^{(N)} = 0$, and $\Upsilon_K^{(I)} < 3\beta f$ (Lemma 3). Finally, after K has become broken, $\Phi_K = 0$. In any case, $\Phi_K \ge 0$. □

For every $K \in \mathcal{K}$, the functions $\Xi_K^{(1)}$ and $\Xi_K^{(2)}$ are non-increasing, because $g(c_K)$ is non-increasing. Thus, the potential function Φ_K can increase only if K is an active coalition and either G_K or N_K is non-empty. In addition, a new demand cannot affect the cost charged to the algorithm for the demands mapped to a non-active coalition K. In particular, if K has not become active yet, there are no such demands, while if K has become active and then broken, K has been charged irrevocably charged with the final assignment cost of all the demands mapped to it. Hence, we can restrict our attention to the coalitions which are active when a new demand arrives.

Isolated Active Coalitions. Let K be an isolated active coalition with representative c_K. Then, $\Phi_K = \Xi_K^{(1)} + 7\sum_{u\in \mathrm{G}_K} d_u^* - \sum_{u\in \mathrm{G}_K} d(u, w_K)$, because $\Xi_K^{(2)} = \Upsilon_K^{(N)} = 0$. In addition, $\Xi_K^{(1)}$ is equal to $7\beta f$ as long as K is an active coalition ($g(c_K) \geq \rho D(K)$), and it becomes zero as soon as K becomes broken ($g(c_K) < \rho D(K)$).

The algorithm is *irrevocably* charged with the final assignment cost of each new outer demand mapped to the isolated active coalition K. Each new inner demand of K which does not make K broken is initially assigned to the nearest facility to c_K (Lemma 9.C) and becomes a good demand. As long as an inner demand of K remains good, the algorithm is charged with its actual assignment cost. The algorithm is *irrevocably* charged with the final assignment cost of the inner demands which have become bad and the demand making K broken. In addition, a new demand mapped to K may change the location of the nearest facility to c_K and/or make the set of good demands of K empty. Thus, we should account for the case that the actual assignment cost of the good demands of K changes because the location of the nearest facility to c_K has changed and the case that all the good demands of K become bad, and from now on, the algorithm is charged with their final instead of their actual assignment cost.

More specifically, let u be a new demand mapped to the isolated active coalition K, and let $\Delta\mathrm{Asg}_K$ be the difference in the assignment cost charged to the algorithm for the demands mapped to K before and after u. Since the final assignment cost charged to the algorithm for outer and bad demands is irrevocable, $\Delta\mathrm{Asg}_K$ is equal to the assignment cost charged to the algorithm for the demands in $\mathrm{G}_K \cup \{u\}$ after u minus the actual assignment cost of the demands in G_K before u. The exact value of $\Delta\mathrm{Asg}_K$ depends on G'_K (i.e., the set of good demands after u).

$$\Delta\mathrm{Asg}_K = \begin{cases} \bar{d}_u + \sum_{v\in \mathrm{G}_K} \bar{d}_v - \sum_{v\in \mathrm{G}_K} d(v, w_K) & \text{if } \mathrm{G}'_K = \emptyset \\ \quad \text{The demands in } \mathrm{G}_K \text{ become bad and are charged with their final assignment cost.} \\ \quad u \text{ either is an outer demand or makes } K \text{ broken.} \\ \bar{d}_u + \sum_{v\in \mathrm{G}_K} d(v, w'_K) - \sum_{v\in \mathrm{G}_K} d(v, w_K) & \text{if } \mathrm{G}'_K = \mathrm{G}_K. \\ \quad \text{The actual assignment cost of the demands in } \mathrm{G}_K \text{ is updated. } u \text{ is an outer demand.} \\ d(u, w'_K) + \sum_{v\in \mathrm{G}_K} \bar{d}_v - \sum_{v\in \mathrm{G}_K} d(v, w_K) & \text{if } \mathrm{G}'_K = \{u\}. \\ \quad \text{The demands in } \mathrm{G}_K \text{ become bad and are charged with their final assignment cost.} \\ \quad u \text{ opens a new facility, is an inner demand and becomes good.} \\ d(u, w'_K) + \sum_{v\in \mathrm{G}_K} d(v, w'_K) - \sum_{v\in \mathrm{G}_K} d(v, w_K) & \text{if } \mathrm{G}'_K = \mathrm{G}_K \cup \{u\}. \\ \quad \text{The actual assignment cost of the demands in } \mathrm{G}_K \text{ is updated.} \\ \quad u \text{ is an inner demand and becomes good.} \end{cases} \tag{13}$$

In the above definition, it may be the case that $w'_K = w_K$, i.e., u opens no new facility and the location of the nearest facility to c_K does not change. By the definitions of the good inner demands of K and the final assignment cost, it should be clear that the actual assignment cost of the demands mapped to the isolated active coalition K can never exceed the cost charged to the algorithm for them. The following lemma establishes that we can ignore the demands not mapped to an isolated active coalition K in the analysis of K.

Lemma 16. *Let K be an isolated active coalition when a new demand u arrives. If u is not mapped to K, then $\Delta\mathrm{Asg}_K = 0$ and $\Delta\Phi_K = 0$.*

Proof. Since u is not mapped to K, its assignment cost is not charged to K. In addition, neither the value of $g(c_K)$ nor the location of the nearest facility to c_K can change because of u (Lemma 8). Therefore, the set of good demands G_K does not change and the cost charged to the algorithm for the demands mapped to K is not affected by u. Hence, $\Delta\mathrm{Asg}_K = 0$ and $\Delta\Upsilon_K^{(I)} = 0$. In addition, since the value of $g(c_K)$ does not change, the function $\Xi_K^{(1)}$ remains equal to $7\beta f$. Consequently, $\Delta\Phi_K = 0$. □

Lemma 17. *Let u be a new demand mapped to the isolated active coalition K. If u makes K broken, then $\Delta\Phi_K + \Delta\mathrm{Asg}_K \leq 0$.*

Proof. Since u makes K broken, the function $\Xi_K^{(1)}$ decreases by $7\beta f$ and G'_K becomes empty. Hence,

$$\Delta\mathrm{Asg}_K = \overline{d}_u + \sum_{v\in\mathrm{G}_K} \overline{d}_v - \sum_{v\in\mathrm{G}_K} d(v,w_K) \le 2.5\beta f + 4.5\beta f + 7\sum_{v\in\mathrm{G}_K} d_v^* - \sum_{v\in\mathrm{G}_K} d(v,w_K)\,,$$

where the second inequality follows from Proposition 14 and Lemma 3. On the other hand, since the set of good demands becomes empty and the function $\Upsilon_K^{(I)}$ becomes zero, we obtain that $-\Delta\Upsilon_K^{(I)} = \sum_{v\in\mathrm{G}_K} d(v,w_K) - 7\sum_{v\in\mathrm{G}_K} d_v^*$. Putting everything together, we conclude that $\Delta\Phi_K + \Delta\mathrm{Asg}_K \le 0$. $\square$

Lemma 18. *Let u be a new demand mapped to the isolated active coalition K. If no new facilities open and K remains an isolated active coalition after u, then $\Delta\Phi_K + \Delta\mathrm{Asg}_K \le 4(\lambda+2)d_u^*$.*

Proof. Since K does not become broken and no new facilities open, the location of the nearest facility to c_K remains the same (i.e., $w'_K = w_K$) and no good demands become bad (i.e., $\mathrm{G}_K \subseteq \mathrm{G}'_K$). In addition, $\Delta\Xi_K^{(1)} = 0$.

If u is an outer demand, then $\mathrm{G}'_K = \mathrm{G}_K$. Since $w'_K = w_K$, and for every $v \in \mathrm{G}_K$, $d(v,w'_K) = d(v,w_K)$, we obtain that $\Delta\mathrm{Asg}_K = \overline{d}_u \le 4(\lambda+2)d_u^*$ (Lemma 10). On the other hand, $\Delta\Upsilon_K^{(I)} = 0$ because $\mathrm{G}'_K = \mathrm{G}_K$ and $w'_K = w_K$. We conclude that $\Delta\Phi_K + \Delta\mathrm{Asg}_K \le 4(\lambda+2)d_u^*$.

If u is an inner demand, by Lemma 9.C, u is initially assigned to w'_K and becomes a good inner demand. Hence, $\mathrm{G}'_K = \mathrm{G}_K \cup \{u\}$ and $-\Delta\Upsilon_K^{(I)} = -d(u,w'_K) + 7\,d_u^*$ (recall that $w'_K = w_K$). On the other hand, $\Delta\mathrm{Asg}_K = d(u,w'_K)$. Therefore, $\Delta\Phi_K + \Delta\mathrm{Asg}_K = 7\,d_u^* \le 4(\lambda+2)d_u^*$. $\square$

Lemma 19. *Let u be a new demand mapped to the isolated active coalition K, and let w be the nearest facility to c_K at u's arrival time (i.e., $w_K = w$). If a new facility w' opens and K is an isolated active coalition after u, then*

$$\Delta\Phi_K + \Delta\mathrm{Asg}_K \le 4(\lambda+2)d_u^* + 14x\mathrm{Asg}^*(B_w) \tag{14}$$

In addition, the neighborhood B_w of each facility w is charged by Inequality (14) at most once.

Proof. We start by observing that $\Delta\Xi_K^{(1)} = 0$, because K remains active after u. Let $w = w_K$ be the nearest facility to c_K at u's arrival time. By Proposition 18.A, w is mapped to an optimal center in K. The new facility w' is mapped to the isolated active coalition K because w' is located at the same point with u. Since w' does not make K broken, it must be $d(c_K,w) \ge \lambda D(K)$ and $d(c_K,w') < \frac{2}{5}\,d(c_K,w)$ (Proposition 20). Moreover, the location of the nearest facility to c_K changes from $w_K = w$ to $w'_K = w'$.

Case A. If w is merged with w', every demand $v \in \mathrm{G}_K$, which was assigned to w before u, is now assigned to w'. Hence, every demand $v \in \mathrm{G}_K$ remains assigned to the nearest facility to c_K, which is now w', and no good demands become bad. Therefore, $\mathrm{G}_K \subseteq \mathrm{G}'_K$. If u is an outer demand, then $\mathrm{G}'_K = \mathrm{G}_K$, and

$$\Delta\mathrm{Asg}_K = \overline{d}_u + \sum_{v\in\mathrm{G}_K} d(v,w'_K) - \sum_{v\in\mathrm{G}_K} d(v,w_K) \quad\text{and}\quad -\Delta\Upsilon_K^{(I)} = -\sum_{v\in\mathrm{G}_K} d(v,w'_K) + \sum_{v\in\mathrm{G}_K} d(v,w_K)\,.$$

Using $\overline{d}_u \le 4(\lambda+2)d_u^*$ (Lemma 10), we conclude that $\Delta\Phi_K + \Delta\mathrm{Asg}_K \le 4(\lambda+2)d_u^*$.

If u is an inner demand, then $\mathrm{G}'_K = \mathrm{G}_K \cup \{u\}$ and $\Delta\mathrm{Asg}_K = d(u,w'_K) + \sum_{v\in\mathrm{G}_K} d(v,w'_K) - \sum_{v\in\mathrm{G}_K} d(v,w_K)$. On the other hand, $-\Delta\Upsilon_K^{(I)} = -d(u,w'_K) + 7\,d_u^* - \sum_{v\in\mathrm{G}_K} d(v,w'_K) + \sum_{v\in\mathrm{G}_K} d(v,w_K)$. Therefore, $\Delta\Phi_K + \Delta\mathrm{Asg}_K \le 7\,d_u^* < 4(\lambda+2)d_u^*$. In any case, if w is merged with w', no good demand of K becomes bad and w's neighborhood B_w (i.e., the set of demands contributing to the opening cost of w) is not charged with any assignment cost.

Case B. If w is not merged with w', then w must be a supported facility (Proposition 8). In this case, all the demands in G_K become bad and are not included in G'_K, since they are no longer assigned to the nearest facility to c_K, which is now w'. If u is an outer demand, then $\mathrm{G}'_K = \emptyset$. Therefore,

$$\Delta\mathrm{Asg}_K = \overline{d}_u + \sum_{v\in\mathrm{G}_K} \overline{d}_v - \sum_{v\in\mathrm{G}_K} d(v,w_K) \quad\text{and}\quad -\Delta\Upsilon_K^{(I)} = \sum_{v\in\mathrm{G}_K} d(v,w_K) - 7\sum_{v\in\mathrm{G}_K} d_v^*\,.$$

If u is an inner demand, then $\mathrm{G}'_K = \{u\}$. Hence, $\Delta\mathrm{Asg}_K = d(u, w'_K) + \sum_{v\in\mathrm{G}_K} \overline{d}_v - \sum_{v\in\mathrm{G}_K} d(v, w_K)$, and $-\Delta\Upsilon_K^{(I)} = -d(u, w'_K) + 7\,d^*_u + \sum_{v\in\mathrm{G}_K} d(v, w_K) - 7\sum_{v\in\mathrm{G}_K} d^*_v$. In addition, if u is an inner demand, then $7\,d^*_u < 4(\lambda+2)d^*_u$, while if u is an outer demand, then $\overline{d}_u \le 4(\lambda+2)d^*_u$ (Lemma 10). Therefore, in both cases,

$$\Delta\Phi_K + \Delta\mathrm{Asg}_K \le 4(\lambda+2)d^*_u + \sum_{v\in\mathrm{G}_K} \overline{d}_v - 7\sum_{v\in\mathrm{G}_K} d^*_v < 4(\lambda+2)d^*_u + 4.5\beta f\,,$$

where the second inequality follows from Lemma 3. Since w is a supported facility, $3x\mathrm{Asg}^*(B_w) = 3x\sum_{v\in B_w} d^*_v \ge \beta f$. Hence, we can charge the final assignment cost of the inner demands which become bad to the optimal assignment cost of the demands in w's neighborhood B_w (i.e., the demands contributing to the opening cost of w). Thus,

$$\Delta\Phi_K + \Delta\mathrm{Asg}_K \le 4(\lambda+2)d^*_u + 14\,x\,\mathrm{Asg}^*(B_w)\,.$$

To conclude the proof, we should also establish that the neighborhood of each facility is charged with the final assignment cost of some demands which become bad at most once. For simplicity, if a facility w is charged with the final assignment cost of some inner demands of an isolated active coalition K which become bad, we say that w is charged by K.

A facility w is charged by an isolated active coalition K only if (i) w is the nearest facility to c_K, (ii) a new facility w' mapped to K opens, and (iii) w' does not make K broken. By Proposition 18.A, the nearest facility to the representative of an isolated active coalition is mapped to an optimal center in the coalition. Therefore, w is mapped to an optimal center in K and cannot be the nearest facility to the representative of any other isolated coalition K' which is disjoint from K (i.e., $K' \subseteq F^* \setminus K$). Consequently, the facility w cannot be charged by any coalition $K' \subseteq F^* \setminus K$. Moreover, by Proposition 21, if (i) w is the nearest facility to c_K, (ii) a new facility w' mapped to K opens and (iii) w' does not make K broken, then w can never become again the nearest facility to any of the optimal centers in K. Hence, once w has been charged by the isolated coalition K, it cannot be charged again by K or any subset / descendant of K in the hierarchical decomposition $\mathcal{K}$, because w will never become again the nearest facility to the representative of any coalition $K' \subseteq K$. □

Non-Isolated Active Coalitions. Each new demand u mapped to a non-isolated active coalition K is *irrevocably* charged with its final assignment cost at its assignment time. Hence, $\Delta\mathrm{Asg}_K = \overline{d}_u$ and for every active coalition K', $K \neq K'$, $\Delta\mathrm{Asg}_{K'} = 0$. In particular, if K' is a non-isolated active coalition, then u cannot affect the irrevocable final assignment cost which has been charged to the algorithm for the demands mapped to K', while if K' is an isolated active coalition, the claim follows from Lemma 16.

Lemma 20. *Let u be a new demand mapped to the non-isolated active coalition K. Then,*

$$\Delta\Phi_K + \Delta\mathrm{Asg}_K = \Delta\Phi_K + \overline{d}_u \le 4((\rho+1)\gamma^2 + 2)d^*_u\,.$$

Proof. The demand u cannot affect $\Upsilon_K^{(I)}$ because it arrives while K is a non-isolated active coalition. In the following, let N_K be the set of unsatisfied inner demands of K at u's arrival time, and let N'_K be the set of unsatisfied inner demands of K at u's assignment time. We distinguish between the following three cases:

Case A. Either $\mathrm{N}'_K \subset \mathrm{N}_K$ or K becomes isolated or broken. In this case, either $g'(c_K) < \rho D_N(K)$ or a new facility w' opens and $B_{w'} \cap \mathrm{N}_K \neq \emptyset$, in which case some of the demands in N_K become satisfied and are no longer included in N'_K.

The function $\Upsilon_K^{(N)}$ does not exceed $5(\psi+4)\gamma^2\beta f$ at u's arrival time (Lemma 4), and is non-negative at u's assignment time. Hence, the increase in the function $-\Upsilon_K^{(N)}$ is $-\Delta\Upsilon_K^{(N)} \le 5(\psi+4)\gamma^2\beta f$. In addition, $\overline{d}_u$ cannot exceed $2.5\beta f$ (Proposition 14). On the other hand, if u makes K either isolated or broken, the function $\Xi_K^{(1)}$ decreases by $[5(\psi+4)\gamma^2 + 2.5]\beta f$. Otherwise, a new facility w' opens and its neighborhood $B_{w'}$ intersects $\mathrm{N}_K \subseteq \mathrm{in}_N(K)$. Then, by Proposition 22, the configuration distance of c_K decreases by a factor greater than 3, i.e., $g'(c_K) < \frac{1}{3}\,g(c_K)$. Since $g(c_K) < (\rho+1)\gamma^2 D_N(K)$, because K is an active coalition before u, and $g'(c_K) \geq \rho D_N(K)$, because K remains a non-isolated active coalition after u, the function $\Xi_K^{(2)}$ decreases by more than $[5(\psi+4)\gamma^2+2.5]\beta f$. In any case, $\Delta\Xi_K^{(1)} + \Delta\Xi_K^{(2)} \leq -[5(\psi+4)\gamma^2+2.5]\beta f$. Hence,

$$\Delta\Phi_K + \overline{d}_u = \Delta\Xi_K^{(1)} + \Delta\Xi_K^{(2)} - \Delta\Upsilon_K^{(N)} + \overline{d}_u \leq -[5(\psi+4)\gamma^2+2.5]\beta f + 5(\psi+4)\gamma^2\beta f + 2.5\beta f \leq 0\,.$$

Case B. $\mathrm{N}'_K = \mathrm{N}_K$ and K remains a non-isolated active coalition. We distinguish between the case that u is an outer demand and the case that u is an inner demand.

If u is an outer demand, $\overline{d}_u \leq 4\,[(\rho+1)\gamma^2+2]\,d_u^*$ (Lemma 13). Let also $\alpha = \frac{g(c_K)}{g'(c_K)} \geq 1$ be the factor by which $g(c_K)$ decreases because of u. Since $\Upsilon_K^{(N)} \leq 5(\psi+4)\gamma^2\beta f$ (Lemma 4), the increase in the function $-\Upsilon_K^{(N)}$ is bounded by $5(1-\frac{1}{\alpha})(\psi+4)\gamma^2\beta f$. On the other hand, the function $\Xi_K^{(2)}$ decreases by $\ln(\alpha)(5(\psi+4)\gamma^2+2.5)\beta f$, because $g(c_K) < (\rho+1)\gamma^2 D_N(K)$, since K is an active coalition before u, and $g'(c_K) \geq \rho D_N(K)$, since K remains a non-isolated active coalition after u. Using $\ln(\alpha) \geq (1-\frac{1}{\alpha})$, for every $\alpha \geq 1$, we conclude that $\Delta\Phi_K + \overline{d}_u \leq 4\,[(\rho+1)\gamma^2+2]\,d_u^*$.

If u is an inner demand of K, since u is not added to N_K (i.e., u has become satisfied), a new facility w' located at the same point with u must have opened. By Proposition 22, the configuration distance of c_K decreases by a factor greater than 3 (i.e., $g'(c_K) < \frac{1}{3}\,g(c_K)$), because the neighborhood $B_{w'}$ of the new facility w' intersects $\mathrm{in}_N(K)$ at u. Similarly to Case A, $\overline{d}_u \leq 2.5\beta f$ (Proposition 14), $-\Delta\Upsilon_K^{(N)} \leq 5(\psi+4)\gamma^2\beta f$ (Lemma 4), and $\Delta\Xi_K^{(2)} \leq -[5(\psi+4)\gamma^2+2.5]\beta f$. Therefore, $\Delta\Phi_K + \overline{d}_u \leq 0$.

Case C. $\mathrm{N}'_K = \mathrm{N}_K \cup \{u\}$ and K remains a non-isolated active coalition. Then, u must be an inner demand which becomes unsatisfied. By Lemma 14, $\overline{d}_u \leq 5\,g'(c_K)$.

Let $\alpha = \frac{g(c_K)}{g'(c_K)} \geq 1$ be the factor by which $g(c_K)$ decreases because of u. Similarly to Case B, $\Delta\Xi_K^{(2)} \leq -(1-\frac{1}{\alpha})[5(\psi+4)\gamma^2+2.5]\beta f$. On the other hand, the function $-\Upsilon_K^{(N)}$ increases by at most $5(1-\frac{1}{\alpha})(\psi+4)\gamma^2\beta f$, because $g(c_K)$ decreases by a factor of α, and decreases by $5\,g'(c_K)$, because u is added to the set of unsatisfied inner demands of K (i.e., $\mathrm{N}'_K = \mathrm{N}_K \cup \{u\}$). Putting everything together, we obtain that

$$\Delta\Phi_K + \overline{d}_u = \Delta\Xi_K^{(2)} - \Delta\Upsilon_K^{(N)} + \overline{d}_u \leq -(1-\tfrac{1}{\alpha})[5(\psi+4)\gamma^2+2.5]\beta f + 5\,(1-\tfrac{1}{\alpha})(\psi+4)\gamma^2\beta f - 5\,g'(c_K) + 5\,g'(c_K)\,.$$

Therefore, $\Delta\Phi_K + \overline{d}_u \leq 0$. □

If a new demand u is not mapped to the non-isolated active coalition K, then $\Delta\mathrm{Asg}_K = 0$, because u's assignment cost is not charged to K and u cannot affect the irrevocable final assignment cost charged to the algorithm for the demands mapped to K.

Lemma 21. *Let K be a non-isolated active coalition when a new demand u arrives. If u is not mapped to K, then $\Delta\Phi_K + \Delta\mathrm{Asg}_K = 0$.*

Proof. The proof is essentially identical to the proof of Lemma 20. There are some differences which only make the proof simpler. In particular, since $\Delta\mathrm{Asg}_K = 0$, we have to bound $\Delta\Phi_K$ instead of $\Delta\Phi_K + \overline{d}_u$, and since u is not mapped to K, we do not have to consider the case that $\mathrm{N}'_K = \mathrm{N}_K \cup \{u\}$. Furthermore, in Case B, we do not have to consider the possibility that u could have been added to N_K (i.e., u is an inner demand). □

The potential function argument implies that for every j, $1 \leq j \leq n$, the assignment cost incurred by the algorithm just after the demand u_j has been considered is at most $2\beta\,[5(\psi+4)\gamma^2+3]\ln(3\gamma^2)\,\mathrm{Fac}^* + [4\,((\rho+1)\gamma^2+2) + 14\,x]\sum_{i=1}^{j} d_{u_i}^*$.

phase(0)
  $\delta = \frac{a_2}{b_2}$; $\Lambda_0 \leftarrow 0$; $f_0 \leftarrow 0$; $\mathrm{A}_0 \leftarrow 0$;
  $R(0) \leftarrow \emptyset$; $F_0 \leftarrow \emptyset$; $\overline{F}_{-1} \leftarrow \emptyset$; $\mathrm{Asg}_0 \leftarrow 0$;
  For each new demand u:
    $R(0) \leftarrow R(0) \cup \{u\}$;
    $\mathrm{IFL}_0(u)$; `/* Updates` F_0 `*/`
    if $|F_0| > \nu k$ then go to phase(1);
IFL_i-Initialization(Λ_i, k)
  $f_i \leftarrow \frac{\Lambda_i}{\delta k}$; $R(i) \leftarrow \emptyset$; $F_i \leftarrow \emptyset$; $\mathrm{A}_i \leftarrow 0$;
  $\mathrm{Asg}_i \leftarrow \mathrm{Asg}_{i-1} + \mathrm{M}_{i-1}$;
complete_phase(i, u)
  Let w be the location of u;
  if $w \notin F_i$ then open(w);
  initial_assignment(u, w);
  $\mathrm{Asg}_i \leftarrow \mathrm{Asg}_{i-1} + \mathrm{M}_{i-1} + \mathrm{A}_i$;

phase(i)
  Merge $\overline{F}_{i-2} \cup F_{i-1}$ into $c_2\, k$ weighted medians using a bi-criteria (c_1, c_2)-approximation algorithm for k-Median.
  Let $\overline{F}_{i-1}$ be the resulting set of (weighted) medians and let M_{i-1} be the cost of assigning the weighted medians in $\overline{F}_{i-2} \cup F_{i-1}$ to $\overline{F}_{i-1}$.
  $\Lambda_i \leftarrow \max\{\alpha\Lambda_{i-1}, \mathrm{M}_{i-1}\}$;
  IFL_i-Initialization(Λ_i, k);
  For each new demand u:
    $R(i) \leftarrow R(i) \cup \{u\}$;
    $\mathrm{IFL}_i(u)$; `/* Updates` F_i `and` A_i `*/`
    if ($|F_i| > \nu k$ or $\mathrm{A}_i > \mu\Lambda_i$) then
      restore IFL_i's configuration before u;
      complete_phase(i, u); go to phase($i+1$);
    else $\mathrm{Asg}_i \leftarrow \mathrm{Asg}_{i-1} + \mathrm{M}_{i-1} + \mathrm{A}_i$;

Fig. 4. The algorithm Incremental k-Median – $\mathrm{IM}(k)$.

A.9 The Proof of Lemma 5

Proof. We recall that given an add-optimal facility configuration of facility cost $\mathrm{Fac}^*_\mathrm{o}$ and assignment cost $\mathrm{Asg}^*_\mathrm{o}$, IFL maintains a solution of facility cost $a_1\mathrm{Fac}^*_\mathrm{o} + b_1\mathrm{Asg}^*_\mathrm{o}$ and assignment cost $a_2\mathrm{Fac}^*_\mathrm{o} + b_2\mathrm{Asg}^*_\mathrm{o}$, where $a_1 = 1$, $a_2 = 2\beta\ln(3\gamma^2)(5(\psi+4)\gamma^2+3)$, $b_1 = \frac{3x}{\beta}$, and $b_2 = 4((\rho+1)\gamma^2+2)+14x$.

Let F^* be a k-Median configuration of cost Asg^* (the medians in F^* are not restricted to the demand locations). The k-Median instance can be regarded as an instance of Facility Location with facility cost $f = \frac{\Lambda}{\delta k}$, where $\delta = \frac{a_2}{b_2}$. Then, F^* is a facility configuration of facility cost $\mathrm{Fac}^* = \frac{\Lambda}{\delta}$ and assignment cost Asg^*. If F^* is not add-optimal, there must be a set of facilities whose addition to F^* makes it add-optimal without increasing its cost. Let F^*_o, $F^* \subseteq F^*_\mathrm{o}$, be the corresponding add-optimal facility configuration. Let also $\mathrm{Fac}^*_\mathrm{o}$ be the facility cost, and let $\mathrm{Asg}^*_\mathrm{o}$ be the assignment cost of F^*_o. It must be the case that (i) $\mathrm{Fac}^*_\mathrm{o} + \mathrm{Asg}^*_\mathrm{o} \le \mathrm{Fac}^* + \mathrm{Asg}^*$, (ii) $\mathrm{Asg}^*_\mathrm{o} \le \mathrm{Asg}^*$, (iii) $\mathrm{Fac}^*_\mathrm{o} \le \mathrm{Fac}^* + \mathrm{Asg}^*$, and (iv) for every $0 \le a \le b$, $a\mathrm{Fac}^*_\mathrm{o} + b\mathrm{Asg}^*_\mathrm{o} \le a\mathrm{Fac}^* + b\mathrm{Asg}^*$, where the last claim follows from (i), (ii), and $F^* \subseteq F^*_\mathrm{o}$.

Let Fac be the facility cost and Asg be the assignment cost of the solution maintained by IFL. Since $a_1 \le b_1$,

$$\mathrm{Fac} \le a_1\mathrm{Fac}^*_\mathrm{o} + b_1\mathrm{Asg}^*_\mathrm{o} \le a_1\mathrm{Fac}^* + b_1\mathrm{Asg}^* \le a_1\frac{\Lambda}{\delta} + b_1\,\mathrm{Asg}^*\,.$$

Using $f = \frac{\Lambda}{\delta k}$ and $\delta = \frac{a_2}{b_2}$, we obtain that IFL's solution consists of no more than $(a_1 + a_2\frac{b_1}{b_2}\frac{\mathrm{Asg}^*}{\Lambda})\,k$ medians. As for the assignment cost,

$$\mathrm{Asg} \le a_2\,\mathrm{Fac}^*_\mathrm{o} + b_2\,\mathrm{Asg}^*_\mathrm{o} \le a_2\,(\tfrac{\Lambda}{\delta} + \mathrm{Asg}^*) + b_2\,\mathrm{Asg}^* \le (a_2+b_2)\,\mathrm{Asg}^* + b_2\,\Lambda\,.$$

A.10 A Deterministic Incremental Algorithm for k-Median

The algorithm $\mathrm{IM}(k)$ (Fig. 4) starts in phase 0, also called the initialization phase, with $\Lambda_0 = 0$ and $f_0 = 0$. An invocation of IFL with facility cost 0 simply opens a new facility/median at each different demand location. Hence, phase 0 ends as soon as IFL_0 has considered exactly $\nu k + 1$ different demand locations. Since, there is a median at each of these locations, the algorithm incurs no assignment cost during the initialization phase.

Phase i, $i \ge 1$, starts with merging the medians produced by the last phase with the medians produced by the previous phases. Thus, we ensure that the total number of medians in the current solution does not depend on the number of phases. More specifically, for each median w in the current solution, we maintain its weight $|C(w)|$, which is equal to the number of demands currently assigned to w. At the beginning of phase i, the set $\overline{F}_{i-2}$ containing the weighted medians produced by phases $0, \ldots, i-2$ is merged with the set F_{i-1} containing the weighted medians produced by phase $i-1$. We can use any bi-criteria (c_1, c_2)-approximation algorithm for k-Median (e.g., the algorithm of [18] for $c_1 = 32$ and $c_2 = 1$ in $\mathrm{O}(k^2 \log n)$ time) to merge $\overline{F}_{i-2}$ with F_{i-1}. The resulting set $\overline{F}_{i-1}$ consists of no more than $c_2\, k$ weighted medians, which are the medians produced by phases $0, \ldots, i-1$. M_{i-1} denotes the cost of assigning the weighted medians in $\overline{F}_{i-2} \cup F_{i-1}$ to $\overline{F}_{i-1}$. The demands considered up to the end of phase $i-1$ are currently assigned to the medians in $\overline{F}_{i-1}$.

The upper bound Λ_i which characterizes the phase i is set to the maximum of $\alpha\Lambda_{i-1}$ and M_{i-1}, where α is a constant chosen sufficiently large. Hence, we ensure that the cost incurred by the algorithm up to the end of phase $i-1$, denoted by Asg_{i-1}, does not exceed Λ_i (Lemma 22). After initializing the invocation of IFL corresponding to phase i, denoted by IFL_i, $\mathrm{IM}(k)$ starts considering new demands. IFL_i incorporates each new demand into the current solution and updates its median configuration, denoted by F_i, and its assignment cost for the demands considered in the current phase i, denoted by A_i. If either F_i contains more than νk medians or the assignment cost A_i exceeds $\mu\Lambda_i$, phase i ends. The algorithm places a new median at the location of the last demand of each phase instead of letting IFL_i incorporate it into the current solution. Hence, the algorithm maintains the invariant that $|F_i| \leq \nu k + 1$ and $\mathrm{A}_i \leq \mu\Lambda_i$. It is straight-forward to modify $\mathrm{IM}(k)$ so as to ensure that no phase ends before it considers at least $\nu k + 1$ new demands. Hence, we can assume that the number of phases is $\mathrm{O}(\frac{n}{k})$. To establish the algorithm's performance ratio, we prove that for every complete phase i, the optimal cost for the demands considered up to the end of phase i is at least $\max\{\Lambda_i, \frac{\mathrm{M}_i}{2c_1(\alpha+1)}\}$ (Lemma 23 and Lemma 24).

Notation. Let $R(i)$ denote the set of demands considered in phase i. If phase i is the current phase, $R(i)$ is the set of demands considered from the beginning of the phase up to the present time. Let $\overline{R}(i) = \cup_{\ell=0}^{i} R(\ell)$. If i is a complete phase, $\overline{R}(i)$ is the set of demands considered up to the end of phase i, while if i is the current phase, $\overline{R}(i)$ includes all the demands considered by the algorithm so far. Let also OPT_i denote the cost of the optimal solution on $\overline{R}(i)$.

Asg_i denotes the cost of the solution $\overline{F}_{i-1} \cup F_i$ on $\overline{R}(i)$. More specifically, if i is a complete phase, Asg_i denotes the cost incurred by the algorithm up to the end of phase i, while if i is the current phase, Asg_i denotes the cost of the current solution on the demands considered so far. Asg_i is always equal to Asg_{i-1}, namely, the cost of $\overline{F}_{i-2} \cup F_{i-1}$ on $\overline{R}(i-1)$, plus M_{i-1}, namely, the cost of assigning $\overline{F}_{i-2} \cup F_{i-1}$ to $\overline{F}_{i-1}$, plus A_i, namely, the cost of F_i on $R(i)$.

Analysis. The algorithm $\mathrm{IM}(k)$ maintains a solution consisting of no more than $(\nu + c_2)\, k + 1$ medians. The following proposition establishes that we have correctly defined Asg_i as the algorithm's cost on the demands considered so far.

Proposition 23. *For every phase i, Asg_i is equal to the cost of $\overline{F}_{i-1} \cup F_i$ on $\overline{R}(i)$.*

Proof. We prove the proposition by induction on i. For the initialization phase ($i = 0$), the proposition holds because IFL_0 is invoked with facility cost $f_0 = 0$. Hence, F_0 contains a median at each different demand location, and the total algorithm's cost is $0 = \mathrm{Asg}_0$. We inductively assume that Asg_{i-1} is equal to the cost of $\overline{F}_{i-2} \cup F_{i-1}$ on $\overline{R}(i-1)$. Then, the cost of assigning $\overline{R}(i-1)$ to $\overline{F}_{i-1}$ is at most $\mathrm{Asg}_{i-1} + \mathrm{M}_{i-1}$, i.e., the cost of first moving the demands in $\overline{R}(i-1)$ from their original locations to $\overline{F}_{i-2} \cup F_{i-1}$ and then to $\overline{F}_{i-1}$. In addition, the cost of assigning the demands in $R(i)$ to F_i is A_i. Therefore, Asg_i, which is always equal to $\mathrm{Asg}_{i-1} + \mathrm{M}_{i-1} + \mathrm{A}_i$, is indeed equal to the cost of $\overline{F}_{i-1} \cup F_i$ on $\overline{R}(i)$. □

Lemma 22. *Let $\alpha \geq \mu + 2$. Then, for every phase i, $\mathrm{Asg}_i \leq \alpha\Lambda_i \leq \Lambda_{i+1}$.*

Proof. We prove the lemma by induction on i. For the initialization phase ($i = 0$), the lemma holds because $\mathrm{Asg}_0 = 0 \leq \alpha\Lambda_0 \leq \Lambda_1$. We inductively assume that the lemma holds until the end of phase i, $i \geq 0$. Then, until the end of phase $i+1$, it is the case that

$$\mathrm{Asg}_{i+1} = \mathrm{Asg}_i + \mathrm{M}_i + \mathrm{A}_{i+1} \leq \alpha\Lambda_i + \mathrm{M}_i + \mu\Lambda_{i+1} \leq (\mu+2)\Lambda_{i+1} \leq \alpha\Lambda_{i+1} \leq \Lambda_{i+2}\,,$$

where the first inequality follows from the inductive hypothesis and the invariant $\mathrm{A}_{i+1} \leq \mu\Lambda_{i+1}$ maintained in phase $i+1$, the second inequality from $\Lambda_{i+1} = \max\{\alpha\Lambda_i, \mathrm{M}_i\}$, and the third inequality from $\alpha \geq \mu + 2$. $\square$

Lemma 23. *Let $\nu \geq a_1 + a_2\frac{b_1}{b_2}$ and $\mu \geq a_2 + 2b_2$. For every complete phase i, $\mathrm{OPT}_i > \Lambda_i$.*

Proof. For the initialization phase ($i = 0$), the lemma holds because $\mathrm{OPT}_0 > \Lambda_0 = 0$. Let us assume that for some complete phase $i \geq 1$, $\mathrm{OPT}_i \leq \Lambda_i$. Let OPT'_i be the optimal cost for the demands considered in phase i. It must be $\mathrm{OPT}'_i \leq \mathrm{OPT}_i \leq \Lambda_i$. Therefore, by Lemma 5, IFL_i must maintain a solution consisting of no more than $(a_1 + a_2\frac{b_1}{b_2})k \leq \nu k$ medians and costing at most $(a_2 + 2b_2)\Lambda_i \leq \mu\Lambda_i$. This contradicts to the hypothesis that phase i is complete. $\square$

Lemma 24. *For every complete phase i, $\mathrm{M}_i \leq 2c_1(\alpha+1)\,\mathrm{OPT}_i$.*

Proof. The optimal solution on $\overline{R}(i)$ suggests a way of merging $\overline{F}_{i-1} \cup F_i$ into k medians. In particular, we can assign each weighted median in $\overline{F}_{i-1} \cup F_i$ to the nearest optimal median. Similarly to the proof of [12, Theorem 2.3], we can show that this assignment costs no more than $\mathrm{Asg}_i + \mathrm{OPT}_i$, i.e., the cost of first moving the demands back to their original locations and then to the optimal medians. Consequently, there is a way of merging $\overline{F}_{i-1} \cup F_i$ into k medians at a cost no greater than $\mathrm{Asg}_i + \mathrm{OPT}_i \leq \alpha\Lambda_i + \mathrm{OPT}_i \leq (\alpha+1)\,\mathrm{OPT}_i$, where the first inequality follows from Lemma 22 and the second inequality from Lemma 23. The above solution can be transformed to a solution using medians only in $\overline{F}_{i-1} \cup F_i$ and costing at most $2\,(\alpha+1)\,\mathrm{OPT}_i$ (e.g., [12, Theorem 2.1]). Since $\overline{F}_i$ is computed by a bi-criteria (c_1, c_2)-approximation algorithm for k-Median, M_i, i.e., the cost of assigning $\overline{F}_{i-1} \cup F_i$ to $\overline{F}_i$, cannot exceed $2c_1(\alpha+1)\,\mathrm{OPT}_i$. $\square$

The Proof of Theorem 2. The number of medians in the current solution can never exceed $(a_1 + a_2\frac{b_1}{b_2} + c_2)\,k + 1$. In the initialization phase, $\mathrm{OPT}_0 > 0 = \mathrm{Asg}_0$. Let $i \geq 0$ be the last complete phase. By Lemma 23 and Lemma 24, it must be $\mathrm{OPT}_i \geq \max\{\Lambda_i, \frac{\mathrm{M}_i}{2c_1(\alpha+1)}\}$. On the other hand, the current algorithm's assignment cost is $\mathrm{Asg}_{i+1} \leq \alpha\Lambda_{i+1}$ (Lemma 22). If $\Lambda_{i+1} = \alpha\Lambda_i$, then $\mathrm{Asg}_{i+1} \leq \alpha^2\,\mathrm{OPT}_i$. If $\Lambda_{i+1} = \mathrm{M}_i$, then $\mathrm{Asg}_{i+1} \leq 2c_1\alpha(\alpha+1)\,\mathrm{OPT}_i$ (Lemma 24). Since $\alpha \geq \mu + 2$ (Lemma 22) and $\mu \geq a_2 + 2b_2$ (Lemma 23), the performance ratio of $\mathrm{IM}(k)$ is less than $2c_1(a_2 + 2b_2 + 3)^2$.

The algorithm $\mathrm{IM}(k)$ runs in $\mathrm{O}(n^2k)$ time and $\mathrm{O}(n)$ space. More specifically, computing $\overline{F}_i$ from $\overline{F}_{i-1} \cup F_i$ at the beginning of phase i takes $\mathrm{O}(k^2\log n)$ time (e.g., [18]) and there are $\mathrm{O}(\frac{n}{k})$ phases. In addition, IFL needs $\mathrm{O}(nk)$ time to incorporate each new demand into the current solution. The bound on the space complexity is trivial, since it implies that every demand is stored in main memory. $\square$

A.11 A Randomized Incremental Algorithm for *k*-Median

The algorithm $\mathrm{RIM}(k)$ (Fig. 5) also operates in phases, where phase i is characterized by an upper bound Λ_i on the optimal cost of the demands considered in the current phase. In phase i, $\mathrm{RIM}(k)$ invokes Gather_i with upper bound Λ_i and IFL_i with facility cost $f_i = \frac{\Lambda_i}{\delta k}$. Each new demand u is first given to Gather_i, which returns a demand to the nearest gathering point $\hat{u}$. Then, $\hat{u}$ is given to IFL_i, which assigns it to a median in F^*_i. Apart from the use of Gather, the description and the analysis of $\mathrm{RIM}(k)$ are similar to those of $\mathrm{IM}(k)$, Section A.10. In the following, we use the notation introduced in the previous section with exactly the same meaning.

We should emphasize that IFL_i still treats different demands moved to the same gathering point by Gather_i as different unit demands and may be put them in different clusters. In other words, the output of Gather_i should be thought of as just a sample taken from the points of the underlying metric space and not as a first-level clustering of the demand sequence. $\mathrm{RIM}(k)$ uses this sample to generate a modified instance which can be represented in a space efficient manner. Then, IFL_i is solely responsible for maintaining a good hierarchical clustering of the modified instance, which can be directly translated into a good hierarchical clustering of the original instance.

phase(0)
  $\delta = \frac{a_2}{b_2}$; $\Lambda_0 \leftarrow 0$; $f_0 \leftarrow 0$; $\mathrm{A}_0 \leftarrow 0$;
  $R(0) \leftarrow \emptyset$; $F_0 \leftarrow \emptyset$; $\overline{F}_{-1} \leftarrow \emptyset$; $\mathrm{Asg}_0 \leftarrow 0$;
  For each new demand u:
    $R(0) \leftarrow R(0) \cup \{u\}$;
    $\mathrm{IFL}_0(u)$; `/* Updates` F_0 `*/`
    if $|F_0| > \nu k$ then go to phase(1);

IFL_i-Initialization(Λ_i, k)
  $f_i \leftarrow \frac{\Lambda_i}{\delta k}$; $R(i) \leftarrow \emptyset$; $F_i \leftarrow \emptyset$; $\mathrm{A}_i \leftarrow 0$;
  $\mathrm{Asg}_i \leftarrow \mathrm{Asg}_{i-1} + \mathrm{M}_{i-1}$;

complete_phase($i, \hat{u}$)
  Let w be the location of $\hat{u}$;
  if $w \notin F_i$ then open(w);
  initial_assignment($\hat{u}, w$);
  $\mathrm{Asg}_i \leftarrow \mathrm{Asg}_{i-1} + \mathrm{M}_{i-1} + \mathrm{A}_i^G + \mathrm{A}_i$;

phase(i)
  Merge $\overline{F}_{i-2} \cup F_{i-1}$ into $c_2\,k$ weighted medians using a bi-criteria (c_1, c_2)-approximation algorithm for k-Median.
  Let $\overline{F}_{i-1}$ be the resulting set of (weighted) medians and let M_{i-1} be the cost of assigning the weighted medians in $\overline{F}_{i-2} \cup F_{i-1}$ to $\overline{F}_{i-1}$.
  $\Lambda_i \leftarrow \max\{\alpha\Lambda_{i-1}, \mathrm{M}_{i-1}\}$;
  $\mathrm{Gather}_i(\Lambda_i, k, \log n, t)$-Initialization;
  IFL_i-Initialization(Λ_i, k);
  For each new demand u:
    $R(i) \leftarrow R(i) \cup \{u\}$;
    $\hat{u} \leftarrow \mathrm{Gather}_i(u)$; `/* Updates` A_i^G `*/`
    if $\mathrm{Gather}_i(u)$ failed then
      complete_phase(i, u); go to phase($i+1$);
    $\mathrm{IFL}_i(\hat{u})$; `/* Updates` F_i `and` A_i `*/`
    if ($|F_i| > \nu k$ or $\mathrm{A}_i > \mu\Lambda_i$) then
      restore IFL_i's configuration before $\hat{u}$;
      complete_phase($i, \hat{u}$); go to phase($i+1$);
    else $\mathrm{Asg}_i \leftarrow \mathrm{Asg}_{i-1} + \mathrm{M}_{i-1} + \mathrm{A}_i^G + \mathrm{A}_i$;

Fig. 5. The algorithm Randomized Incremental k-Median – $\mathrm{RIM}(k)$.

The current phase of $\mathrm{RIM}(k)$ ends if either Gather_i fails to maintain the invariants on the number of gathering points and the gathering cost or IFL_i fails to maintain the invariants on the number of medians and the assignment cost. To establish the performance ratio of $\mathrm{RIM}(k)$, we prove that the total algorithm's cost up to the end of phase i cannot exceed $\alpha\Lambda_i$ (Lemma 28), while for every complete phase i, OPT_i is at least $\max\{\Lambda_i, \frac{\mathrm{M}_i}{2c_1(\alpha+1)}\}$ whp. (Lemma 29) and Lemma 30).

The algorithm Gather (Fig. 6) can be thought of as the incremental version of PARA_CLUSTER [6]. It is made up of $\mathrm{O}(\log n)$ independent invocations of Meyerson's randomized algorithm for Online Facility Location [19], denoted by ROFL. In phase i, Gather_i invokes ROFL_i with facility cost $\hat{f}_i = \frac{\Lambda_i}{k(\log n+1)}$. The j-th invocation of ROFL_i, denoted by $\mathrm{ROFL}_i(j)$, maintains its own set of gathering points, denoted by $G_i(j)$, and its individual cost, denoted by $\mathrm{A}_i^G(j)$. When a new demand u is considered, with probability $\min\{\frac{d(G_i(j),u)}{\hat{f}_i}, 1\}$, $\mathrm{ROFL}_i(j)$ places a new gathering point at u's location. Then, it moves u to the nearest gathering point in $G_i(j)$. $\mathrm{ROFL}_i(j)$ fails as soon as either its number of gathering points exceeds $20k\,(\log n + 1)$ or its individual cost exceeds $20\Lambda_i$. After $\mathrm{ROFL}_i(j)$ has failed, it stops considering new demands.

Gather_i maintains the union of the sets of gathering points, denoted by G_i, and the gathering cost, denoted by A_i^G. When a new demand u is considered, Gather_i places a new gathering point at u's location if at least one of the invocations $\mathrm{ROFL}_i(j)$ does so. Then, it moves u to the nearest gathering point currently in G_i, denoted by $\hat{u}$. Gather_i fails as soon as all the invocations $\mathrm{ROFL}_i(j)$ have failed.

Lemma 25. $\mathrm{Gather}_i(\Lambda_i, k, \log n, t)$ *maintains a collection of no more than* $20kt\,(\log n+1)^2$ *gathering points at a cost not exceeding* $20\Lambda_i$.

Proof. The set of gathering points G_i is equal to the union of the sets $G_i(j)$ maintained by the invocations $\mathrm{ROFL}_i(j)$, $j = 1, \ldots, t\log n$. The cardinality of each $G_i(j)$ cannot exceed $20k\,(\log n+1)+1$, because as soon as $|G_i(j)|$ becomes greater than $20k\,(\log n + 1)$, $\mathrm{ROFL}_i(j)$ fails and stops considering new demands. In addition, the location of the last demand, namely, the demand making Gather_i fail, is also added to G_i. Hence, the number of gathering points maintained by $\mathrm{Gather}_i(\Lambda_i, k, \log n, t)$ is upper bounded by $[20k\,(\log n + 1) + 1]\,t\log n + 1 \leq 20kt\,(\log n + 1)^2$.

As long as Gather_i does not fail, the gathering cost A_i^G is upper bounded by the individual cost $\mathrm{A}_i^G(j)$ of any invocation $\mathrm{ROFL}_i(j)$ which has not failed yet. This is true because Gather_i moves each

$\mathrm{Gather}_i(\Lambda_i, k, \log n, t)$-Initialization $\quad \hat{f}_i \leftarrow \frac{\Lambda_i}{k(\log n+1)}$; $G_i \leftarrow \emptyset$; $\mathrm{A}_i^G \leftarrow 0$; $\quad$ for $j \leftarrow 1$ to $t \log n$ do $\qquad G_i(j) \leftarrow \emptyset$; $\mathrm{A}_i^G(j) \leftarrow 0$; $\qquad \mathrm{fail}_i(j) \leftarrow$ FALSE; $\mathrm{ROFL}_i(j)(u)$ $\quad$ with probability $\min\{d(G_i(j), u)/\hat{f}_i, 1\}$ do $\qquad G_i(j) \leftarrow G_i(j) \cup \{u\}$; $\quad \mathrm{A}_i^G(j) \leftarrow \mathrm{A}_i^G(j) + d(G_i(j), u)$;	$\mathrm{Gather}_i(u)$ $\quad$ for $j \leftarrow 1$ to $t \log n$ do $\qquad$ if (not $\mathrm{fail}_i(j)$) then $\qquad\quad \mathrm{ROFL}_i(j)(u)$; $\qquad\quad$ if ($\lvert G_i(j)\rvert > 20k(\log n + 1)$ or $\mathrm{A}_i^G(j) > 20\Lambda_i$) then $\qquad\qquad \mathrm{fail}_i(j) \leftarrow$ TRUE; $\quad$ if ($\bigwedge_{j=1}^t \mathrm{fail}_i(j)$) then $\qquad G_i \leftarrow G_i \cup \{u\}$; return(FAILURE); $\quad G_i \leftarrow \bigcup_{j=1}^t G_i(j)$; $\quad$ let $\hat{u}$ be the nearest gathering point to u; $\quad$ move u to $\hat{u}$; $\mathrm{A}_i^G \leftarrow \mathrm{A}_i^G + d(G_i, u)$; $\quad$ return($\hat{u}$);

Fig. 6. A randomized algorithm for gathering the original demands in $\mathrm{O}(k \log^2 n)$ points.

new demand to the nearest gathering point in G_i and $G_i(j) \subseteq G_i$. Hence, as long as there exists an invocation $\mathrm{ROFL}_i(j)$ which has not failed yet, it is the case that $\mathrm{A}_i^G \leq \mathrm{A}_i^G(j) \leq 20\Lambda_i$. In addition, Gather_i places a gathering point at the location of the last demand and incurs no gathering cost for it. □

The following lemma is proven in [6], Lemma 1 and Corollary 1. Its proof follows from the analysis of [19].

Lemma 26. *Let* Asg^* *be the cost of a feasible solution for an instance of* k*-Median consisting of no more than* n *unit demands, and let* Λ *be an estimation of* Asg^**. With probability at least* $\frac{1}{2}$, ROFL *with facility cost* $f = \frac{\Lambda}{k(\log n+1)}$ *maintains a solution consisting of no more than* $4k\,(\log n + 1)(1 + \frac{4\,\mathrm{Asg}^*}{\Lambda})$ *medians and costing at most* $4(\Lambda + 4\mathrm{Asg}^*)$.

Lemma 27. *Let* Asg^* *be the cost of a feasible solution for an instance of* k*-Median consisting of no more than* n *unit demands, let* $\Lambda \geq \mathrm{Asg}^*$ *be an upper bound on* Asg^**, and let* t *be a positive constant. Then, with probability at least* $1 - n^{-t}$, $\mathrm{Gather}(\Lambda, k, \log n, t)$ *does not fail on this instance.*

Proof. The algorithm Gather fails only if all independent invocations $\mathrm{ROFL}(j)$ fail. For every j, $j = 1, \ldots, t\log n$, $\mathrm{ROFL}(j)$ fails only if either $\lvert G(j)\rvert > 20k(\log n + 1)$ or $\mathrm{A}^G(j) > 20\Lambda$. Since $\Lambda \geq \mathrm{Asg}^*$, by Lemma 26, the probability that $\mathrm{ROFL}(j)$ fails on this instance is at most $\frac{1}{2}$. Since the invocations of ROFL are independent from each other, the probability that all of them fail on this instance is at most n^{-t}. □

As before, $\mathrm{RIM}(k)$ operates in $\mathrm{O}(\frac{n}{k})$ phases and always maintains a solution consisting of no more than $(\nu + c_2)\,k + 1$ medians. Similarly to Proposition 23, we can prove that for every phase i, $\mathrm{Asg}_i = \mathrm{Asg}_{i-1} + \mathrm{M}_{i-1} + \mathrm{A}_i^G + A_i$ is equal to the cost of $\overline{F}_{i-1} \cup F_i$ on $\overline{R}(i)$ (i.e., the set of original demands considered up to the end of phase i).

Lemma 28. *Let* $\alpha \geq \mu + 22$*. Then, for every phase* i, $\mathrm{Asg}_i \leq \alpha\Lambda_i \leq \Lambda_{i+1}$.

Proof. We prove the lemma by induction on i. For the initialization phase ($i = 0$), the lemma holds because no gathering takes place and $\mathrm{Asg}_0 = 0 \leq \alpha\Lambda_0 \leq \Lambda_1$. We inductively assume that the lemma holds until the end of phase i, $i \geq 0$. Then, until the end of phase $i + 1$, it is the case that

$$\mathrm{Asg}_{i+1} = \mathrm{Asg}_i + \mathrm{M}_i + \mathrm{A}_{i+1}^G + \mathrm{A}_{i+1} \leq \alpha\Lambda_i + \mathrm{M}_i + 20\Lambda_{i+1} + \mu\Lambda_{i+1} \leq (\mu + 22)\Lambda_{i+1} \leq \alpha\Lambda_{i+1} \leq \Lambda_{i+2}\,.$$

The first inequality follows from the inductive hypothesis, the invariant $\mathrm{A}_{i+1}^G \leq 20\Lambda_{i+1}$ maintained by Gather_i (Lemma 25), and the invariant $\mathrm{A}_{i+1} \leq \mu\Lambda_{i+1}$ maintained by $\mathrm{RIM}(k)$ in phase i. The second inequality follows from $\Lambda_{i+1} = \max\{\alpha\Lambda_i, \mathrm{M}_i\}$, and the third inequality from $\alpha \geq \mu + 22$. □

Lemma 29. *Let* $\nu \geq a_1 + 21a_2\frac{b_1}{b_2}$, $\mu \geq 21a_2 + 22b_2$*, and let* i *be a complete phase. With probability at least* $1 - n^{-t}$, $\mathrm{OPT}_i > \Lambda_i$.

Proof. For the initialization phase ($i = 0$), the lemma holds with certainty because $\mathrm{OPT}_0 > \Lambda_0 = 0$. Let us assume that for some complete phase $i \geq 1$, $\mathrm{OPT}_i \leq \Lambda_i$. Let OPT'_i be the optimal cost for the demands considered in phase i. It must be $\mathrm{OPT}'_i \leq \mathrm{OPT}_i \leq \Lambda_i$. By Lemma 27, the probability that phase i ends because Gather_i fails is at most n^{-t}. On the other hand, if $\mathrm{OPT}'_i \leq \Lambda_i$ and Gather_i does not fail, phase i cannot end because of IFL_i (see also Lemma 23). In particular, let us assume that Gather_i does not fail and phase i ends because either $|F_i| > \nu k$ or $\mathrm{A}_i > \mu\Lambda_i$. By Lemma 25, the gathering cost A_i^G is at most $20\Lambda_i$. Hence, for the modified instance considered by IFL_i, there exists a k-Median solution of cost no greater than $\mathrm{A}_i^G + \mathrm{OPT}'_i \leq 21\Lambda_i$, namely, the solution obtained by first moving the demands in $R(i)$ from the gathering points to their original locations and then to the optimal medians. We also recall that IFL_i still treats different demands moved to the same gathering point by Gather_i as different unit demands. By Lemma 5, the solution produced by IFL_i on the modified instance consists of no more than $(a_1 + 21a_2\frac{b_1}{b_2})k \leq \nu k$ medians and costs at most $(21a_2 + 22b_2)\Lambda_i \leq \mu\Lambda_i$. Therefore, if $\mathrm{OPT}'_i \leq \Lambda_i$ and Gather_i does not fail, phase i cannot end because of IFL_i. Consequently, the probability that $\mathrm{OPT}_i \leq \Lambda_i$ and phase i ends is at most n^{-t}. □

Lemma 30. *Let i be a complete phase. If $\mathrm{OPT}_i \geq \Lambda_i$, then $\mathrm{M}_i \leq 2c_1(\alpha + 1)\,\mathrm{OPT}_i$.*

Proof. The proof is essentially identical to the proof of Lemma 24. □

The Proof of Theorem 3. The number of medians in the current solution (i.e., $\overline{F}_{i-1} \cup F_i$) can never exceed $(a_1 + 21a_2\frac{b_1}{b_2} + c_2)\,k + 1$. In the initialization phase, $\mathrm{OPT}_0 > 0 = \mathrm{Asg}_0$. Let $i+1$, $i \geq 0$, be the current phase, and let i be the last complete phase. The current algorithm's cost is $\mathrm{Asg}_{i+1} \leq \alpha\Lambda_{i+1}$ (Lemma 28). Given that $\mathrm{OPT}_i \geq \Lambda_i$, we distinguish between $\Lambda_{i+1} = \alpha\Lambda_i$ and $\Lambda_{i+1} = \mathrm{M}_i$. In the first case, $\mathrm{Asg}_{i+1} \leq \alpha^2\,\mathrm{OPT}_i$, while in the second case, $\mathrm{Asg}_{i+1} \leq 2c_1\alpha(\alpha+1)\,\mathrm{OPT}_i$ (Lemma 30). Let $t \geq 2$. By Lemma 29 and since there are $\mathrm{O}(\frac{n}{k})$ complete phases, the probability that there exist a complete phase i such that $\mathrm{OPT}_i < \Lambda_i$ is at most n^{-t+1}. Since $\alpha \geq \mu + 22$ (Lemma 28) and $\mu \geq 21a_2 + 22b_2$ (Lemma 29), the performance ratio of $\mathrm{RIM}(k)$ is less than $2c_1[22(a_2+b_2+1)]^2$ with probability at least $1 - n^{-t+1}$.

The algorithm $\mathrm{RIM}(k)$ runs in $\mathrm{O}(nk^2\log^2 n)$ time and $\mathrm{O}(k^2\log^2 n)$ space. More specifically, computing $\overline{F}_i$ from $\overline{F}_{i-1} \cup F_i$ at the beginning of phase i takes $\mathrm{O}(k^2\log n)$ time (e.g., [18]) and there are $\mathrm{O}(\frac{n}{k})$ phases. In addition, Gather needs $\mathrm{O}(k\log^2 n)$ time to move each new demand to the nearest gathering point and IFL needs $\mathrm{O}(k^2\log^2 n)$ time to incorporate each new demand (of the modified instance) into the current solution. As for the space complexity, Gather can be implemented in $\mathrm{O}(k\log^2 n)$ space and IFL can be implemented in $\mathrm{O}(k^2\log^2 n)$ space. □