Incremental Consistent k-Center Clustering

Mara Grilnberger

Department of Computer Science, University of Salzburg, Austria

Antonis Skarlatos

Department of Computer Science, University of Warwick, Coventry, UK

Abstract

In the incremental consistent k-center clustering problem, we are given a sequence of adversarial point insertions and aim to maintain a k-center solution with small approximation ratio and small recourse. In this work, we explore the following question: what is the best approximation ratio of a polynomial-time algorithm with a worst-case recourse of 1? Our result improves upon the 6-approximation algorithm of Forster and Skarlatos [SODA '25], which itself improved over the 8-approximation algorithm of Charikar, Chekuri, Feder, and Motwani [STOC '97]. Moreover, we show that any incremental k-center algorithm that achieves an approximation ratio strictly less than $\sqrt{2}$ requires a worst-case recourse of k.

2012 ACM Subject Classification Theory of computation → Facility location and clustering

Keywords and phrases Consistent Clustering, k-Center, Dynamic Algorithms

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2026.29

Category APPROX

Funding *Mara Grilnberger*: This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 947702). This publication has been supported by the EXDIGIT (Excellence in Digital Sciences and Interdisciplinary Technologies) project, funded by Land Salzburg under grant number 20204-WISS/263/6-6022.

Antonis Skarlatos: Antonis Skarlatos is funded by the European Union (ERC grant, DYNALP, 101170133). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

1 Introduction

Clustering is a fundamental problem in data analysis and classification. One of the most classical k-clustering problems is the k-center problem, which has been studied in various settings [15, 9, 2]. Given a metric space $(\mathcal{X}, \text{dist})$, a set of points $P \subseteq \mathcal{X}$ with $|P| = n$, and an integer $k \leq n$, the goal of the k-center problem is to output a subset $S \subseteq P$ of at most k points, called *centers*, such that the maximum distance from every point in P to its closest center is minimized. The k-center problem admits polynomial-time 2-approximation algorithms [16, 12], and it is known to be NP-hard to approximate the k-center objective within a factor of $(2 - \epsilon)$ for any constant $\epsilon > 0$ [17].

In the *fully dynamic* setting, the adversary can insert and delete points. In the *partially dynamic* setting, only one type of adversarial update – either insertions or deletions – is allowed. For example, in the *incremental* setting, the set of centers S is maintained under adversarial point insertions into P. The k-clustering problems in the dynamic setting with recourse have attracted a lot of attention [7, 19, 10, 11, 6, 5, 14].

Worst-case recourse

The *worst-case recourse* is the number of replacements made to the set of centers after any adversarial point update. In general, worst-case guarantees are crucial in real-time systems, and extensive research efforts have focused on achieving them [8, 3, 20, 4, 1, 14, 13]. From a theoretical perspective, minimizing worst-case recourse is a fundamental combinatorial problem, as it captures how smoothly solutions adapt in response to adversarial updates. From a practical standpoint, consider a scenario where the centers are used for an expensive task, and any modification to the solution introduces latency. Amortized recourse guarantees may still allow occasional long delays, whereas worst-case recourse guarantees ensure that no significant delays occur at any step.

To understand the limitations of recourse, an interesting direction is to study the best approximation ratio achievable by a polynomial-time algorithm with optimal recourse bounds. Fully dynamic algorithms with small recourse tend to have large approximation ratios. Hence, it is natural to study the incremental setting when the recourse is optimal – namely, a worst-case recourse of 1.

1.1 Upper Bounds

The area of research focusing on optimizing the recourse is commonly referred to as *consistent clustering*, a concept introduced by Lattanzi and Vassilvitskii [19]. In the incremental consistent setting, the "doubling algorithm" of Charikar, Chekuri, Feder, and Motwani [7] achieves an 8-approximation; their incremental algorithm maintains k clusters and only permits merging two clusters. The same work also presents a 6-approximation algorithm that requires solving the NP-hard clique partition problem as a subroutine. More recently, Forster and Skarlatos [11] obtained a 6-approximation algorithm in the incremental setting with a worst-case recourse of 1 and a worst-case update time of $O(k^2)$.[1] To compare their incremental algorithm with the "doubling algorithm", the authors in [11] showed that a slightly modified version of the latter [7] admits a worst-case recourse of 1 as well. In this direction, we improve upon prior work by reducing the approximation ratio to $3 + \sqrt{5} \approx 5.23$.

▶ **Theorem 1.** *There is a deterministic incremental algorithm for the k-center clustering problem that, given a point set P in a metric space subject to point insertions and an integer $k \geq 1$, maintains a subset of points $S \subseteq P$ with $|S| = k$ such that:*

- *The set S is a $(3 + \sqrt{5})$-approximate k-center solution.*
- *The worst-case recourse is at most* 1.
- *The worst-case update time is $O(k^2)$.*

Interestingly, our improved approximation ratio beats even the randomized algorithm of [7], which yields a $2e \approx 5.43$ approximation. As a result, our deterministic algorithm improves on a longstanding bound from more than 25 years ago. To obtain this improvement, motivated by the idea of [11] to combine the two standard static k-center algorithms – Gonzalez's algorithm and the Hochbaum–Shmoys algorithm [16] – we introduce new ideas and provide a more fine-grained and sophisticated analysis of the approximation ratio.

[1] The main result of [11] is, however, their fully dynamic algorithm.

1.2 Lower Bounds

Another interesting direction is to find the best approximation ratio for some fixed worst-case recourse bound. While there are results giving lower bounds on amortized recourse [19], namely a lower bound of $\Omega(\frac{k \log n}{n})$ changes, nothing is known about lower bounds for worst-case recourse. We initiate this line of work by proving the following theorem for the partially dynamic setting.

▶ **Theorem 2.** *Consider a point set P in a metric space subject to point insertions, an integer $k \geq 1$, and an approximation factor $\alpha \in [1, \sqrt{2})$. Then any partially dynamic α-approximate k-center algorithm has a worst-case recourse of k.*

Recently, Kozma [18] investigates the *nested* setting, where the k-center solution has to be built center by center, giving a good solution at each step (where the number of centers is at most k). While this differs from the incremental setting that we consider, two of the presented lower bounds transfer to "our" incremental setting, where we allow adversarial point insertions. This is because the adversary can add $k-1$ points to their examples that are at a distance "infinity" from all other points, leaving less centers for the original points at each insertion. Once the number of centers on the original points becomes small enough, by their arguments, the same bounds follow. As a result, the actual lower bound for an approximation factor is at least $2\cos\left(\frac{\pi}{7}\right)$ (see Section 2.2 in [18]).

2 Preliminaries

Consider a point set P in a metric space $(\mathcal{X}, \mathrm{dist})$ with $|P| = n$. We omit the notation $(\mathcal{X}, \mathrm{dist})$ when the context is clear. For each point $p \in P$ and each subset $S \subseteq P$, the *distance* from p to S is equal to $\mathrm{dist}(p, S) \coloneqq \min_{q \in S} \mathrm{dist}(p, q)$.

▶ **Definition 3** (k-center clustering problem)**.** *Given a point set P in a metric space and an integer $k \geq 1$, the goal of the k-center clustering problem is to output a subset $S \subseteq P$ of at most k points such that the value $\max_{p \in P} \mathrm{dist}(p, S)$ is minimized.*

The points in S are referred to as *centers*, and the (k-center) *cost* of S is equal to $\mathrm{cost}(S) \coloneqq \max_{p \in P} \mathrm{dist}(p, S)$. The *optimal* k-center cost is denoted by $R^* \coloneqq \min_{|S| \leq k} \mathrm{cost}(S)$. In the *improper* k-center clustering problem, the k-center solution S is allowed to include points from the metric space that are not contained in the point set P.

▶ **Definition 4** (improper k-center clustering problem)**.** *Given a point set P in a metric space $(\mathcal{X}, \mathrm{dist})$ and an integer $k \geq 1$, the goal of the improper k-center clustering problem is to output a subset $S \subseteq \mathcal{X}$ of at most k points such that the value $\max_{p \in P} \mathrm{dist}(p, S)$ is minimized.*

The improper k-center clustering problem is defined for the following technical reason. The optimal k-center cost $\hat{R}^*$ of the improper version is non-decreasing under adversarial point insertions, whereas the optimal k-center cost R^* of Definition 3 may decrease after an adversarial point insertion. Moreover, by definition it holds that $\hat{R}^* \leq R^*$. Therefore, we can upper bound the cost of our k-center solution with respect to the monotonic $\hat{R}^*$, which in turn is upper bounded by the desired optimal k-center cost R^*.

In the *incremental* setting, the set of centers S is maintained under adversarial point insertions into P. The *worst-case recourse* is the number of replacements made to the set of centers after any adversarial point insertion. Note that deterministic dynamic algorithms work against an adaptive adversary.

3 Technical Overview

3.1 Incremental Algorithm

Our incremental k-center algorithm is given an incremental point set P that is subject to adversarial point insertions. The algorithm unfolds in *phases*, where each phase can last for multiple adversarial point insertions. During a phase the radius r remains unchanged, and when the subsequent phase begins the radius r is increased by at least a factor of $\alpha > 1$.

Each center is labeled as either *leader* or *regular*. Each point is assigned to its responsible center, and the corresponding distance provides an upper bound on the cost of the k-center solution. A leader center is responsible for points within distance $\alpha r + \beta r$, while a regular center (at the beginning of a phase) is responsible for points within distance βr. Whenever the radius r is increased, some centers become regular. For the updated radius r, every point lies within distance βr from its nearest center. Therefore, the parameters α and β should satisfy $\alpha r + \beta r = \alpha\beta r$, which means that $\alpha + \beta = \alpha\beta$.

Gonzalez Operation

At the beginning of each phase, the centers are ordered using Gonzalez's algorithm [12], which iteratively selects the point farthest from the previously chosen centers as the next center. Any center within distance αr of another center with a lower index is labeled as regular, and the rest are labeled as leaders.

3.1.1 Approximation Ratio Within a Phase

At the last adversarial point insertion of a phase, all pairwise distances between centers are greater than $\alpha^i \cdot r$, all centers are leaders, and the newly inserted point is at a distance greater than $\alpha^i \cdot r$ from every center, where $i \geq 1$ is an integer. The radius r is then increased to at least $\alpha^i \cdot r$, which implies that $r \leq 2\hat{R}^* \leq 2R^*$ for the increased radius r. As new points are inserted by the adversary, the *primary goal* during a phase is to adjust the centers so that the minimum pairwise distance between them is at least αr. The leader centers remain fixed, while the regular centers are allowed to *move* in order to achieve the primary goal. Whenever a regular center c moves, it becomes a leader center and the incremental algorithm ensures that the distance between c and all other centers is greater than βr. Furthermore, the algorithm ensures that c is at a distance greater than αr from any other leader center. Once the primary goal is fulfilled, the phase ends and the subsequent phase begins.

Subphases

Notice that when a regular center moves, some points may lose their assigned center. In turn, another center c becomes responsible for these points, potentially increasing the distance from c to points assigned to it. This center c might also move as a result of another adversarial point insertion, causing a chain of changes in which some points are repeatedly reassigned to different centers. Since these adjustments affect the approximation ratio, for the purpose of analyzing the approximation ratio, we divide each phase into two subphases.

The first subphase is active while the pairwise distances of centers are at least r. The second subphase begins when the minimum distance between two centers becomes greater than βr. Since the incremental algorithm ensures that the distance between a center that moves and all other centers is greater than βr, the minimum distance between any two centers stays greater than βr until the end of the phase.

First subphase. During the first subphase, the pairwise distances of centers are at least r, which implies that $r \leq 2R^*$. Moreover during the first subphase, the minimum distance between two centers involves only centers that have not moved. In turn due to the Gonzalez Operation, the old position in the metric space of every *replaced* center is within distance $2R^*$ from its closest center. As a result, all points are within distance $\max(2R^* + \beta r, \alpha r)$ from their closest center.

Second subphase. During the second subphase, the pairwise distances of centers are greater than βr, which implies that $\beta r \leq 2R^*$. Since every regular center is within distance αr from a leader center, the old position in the metric space of every *replaced* center is within distance αr from its closest center. As a result, all points are within distance $\alpha r + \beta r$ from their nearest center. Observe that as $\beta r \leq 2R^*$, we have $\alpha r \leq 2R^* \cdot \frac{\alpha}{\beta}$. Consequently, it follows that all points are within distance $2R^* \cdot (\frac{\alpha}{\beta} + 1)$ from their closest center.
Finally, the values for the parameters α and β that satisfy the requirements (i.e., $1 \leq \beta < \alpha$ and $\alpha + \beta \leq \alpha\beta$) and minimize the approximation ratio are $\alpha = \frac{3+\sqrt{5}}{2}$ and $\beta = \frac{1+\sqrt{5}}{2}$. The resulting approximation ratio of the incremental algorithm is $3 + \sqrt{5}$.

4 Incremental Consistent k-Center Clustering

In the incremental setting, the point set P undergoes point insertions sent by the adaptive adversary. Our goal is to develop an incremental algorithm for the k-center clustering problem with small approximation ratio, low update time, and worst-case recourse of 1. In this section, we develop a deterministic incremental algorithm whose guarantees are demonstrated in Theorem 1. Recall that R^* denotes the current optimal k-center cost of the given instance.

▶ **Theorem 1.** *There is a deterministic incremental algorithm for the k-center clustering problem that, given a point set P in a metric space subject to point insertions and an integer $k \geq 1$, maintains a subset of points $S \subseteq P$ with $|S| = k$ such that:*
- *The set S is a $(3 + \sqrt{5})$-approximate k-center solution.*
- *The worst-case recourse is at most 1.*
- *The worst-case update time is $O(k^2)$.*

State of the incremental algorithm

Our incremental algorithm maintains an ordered set $S = \{c_1, \ldots, c_k\}$ of k centers, a *radius* r, and the minimum pairwise distance between centers $minDistS := \min_{x,y \in S,\, x \neq y} \mathrm{dist}(x, y)$. Each center is labeled as either *regular* or *leader*, and we denote by S_{reg} and S_{ldr} the sets of regular and leader centers respectively. Two parameters α and β are used in the algorithm satisfying $1 \leq \beta < \alpha$ and $\alpha + \beta \leq \alpha\beta$; the specific values of these parameters are $\alpha = \frac{3+\sqrt{5}}{2}$ and $\beta = \frac{1+\sqrt{5}}{2}$ (i.e., β is the golden ratio).

Invariants for the approximation ratio

The following four invariants are satisfied throughout the incremental algorithm. By enforcing Invariant 1 and Invariant 2, we deduce that $r \leq 2R^*$. Using Invariant 3 and a more fine-grained version of Invariant 4, we prove that the approximation ratio is $3 + \sqrt{5}$. These invariants are responsible solely for the approximation ratio, so our aim is to satisfy them with a worst-case recourse of 1.

Invariant 1: $|S| = k$.
Invariant 2: There is a point $p_S \in P \setminus S$ such that the pairwise distances of points in $S \cup \{p_S\}$ are at least r.
Invariant 3: For every regular center $c_1 \in S_{\text{reg}}$, there is a leader center $c_2 \in S_{\text{ldr}}$ such that $\text{dist}(c_1, c_2) \le \alpha r$.
Invariant 4: Every point $p \in P$ is within distance $\alpha r + \beta r$ from its closest center in S.

4.1 Incremental Algorithm

We begin by describing the Gonzalez Operation, which is central to the incremental algorithm and induces an ordering of the centers. Next, we present the preprocessing phase of the incremental algorithm, and then explain how the algorithm handles adversarial point insertions. A pseudocode of the incremental algorithm is provided in Algorithm 1. Recall that α and β satisfy $1 \le \beta < \alpha$ and $\alpha + \beta \le \alpha\beta$.

Gonzalez Operation

During the Gonzalez Operation, the algorithm sets c_1 to be an arbitrary point from S and labels it as a leader. In each subsequent iteration $i \ge 2$, the algorithm sets c_i to be a point from $S \setminus \{c_1, \dots, c_{i-1}\}$ that is farthest from the already assigned centers $\{c_1, \dots, c_{i-1}\}$. If c_i is within distance αr from a leader, it is labeled as regular; otherwise, c_i is labeled as a leader. Observe that this procedure results in an index $j \in [1, k]$ such that $S_{\text{ldr}} = \{c_1, \dots, c_j\}$ and $S_{\text{reg}} = \{c_{j+1}, \dots, c_k\}$.

4.1.1 Preprocessing Phase

The preprocessing phase is similar to that of the incremental algorithm in [11]. It utilizes the following method by Hochbaum and Shmoys [16] that repeatedly computes a *distance-r independent set*.

▶ **Definition 5** (Distance-r Independent Set)**.** *Given a point set P in a metric space and a parameter $r > 0$, a distance-r independent set $S \subseteq P$ is a subset of points such that $\min_{p,q \in S} \text{dist}(p, q) > r$.*

▶ **Lemma 6** ([16])**.** *Consider a k-center instance (P, k), where P is a point set in a metric space, and let D be the sorted set of pairwise distances among points in P. By applying binary search on the set D and calling a maximal distance-r independent set algorithm $\mathcal{A}$ at each iteration, we can find a 2-approximate solution for the k-center instance. The running time of this algorithm is $O(\mathcal{T}_{\mathcal{A}} \cdot \log |P|)$, where $\mathcal{T}_{\mathcal{A}}$ is the running time of $\mathcal{A}$.*

In the preprocessing phase, we run the method of Lemma 6 to obtain a 2-approximate solution S. Let $r := \max_{p \in P} \text{dist}(p, S)$ be the radius of S. The next task is to ensure that all four invariants are satisfied. To achieve this, the following steps are executed while $|S| < k$:
- If there is a point at a distance of at least r from S, then it is added to S.
- Otherwise, r is updated to $\frac{r}{\beta}$.

Once $|S| = k$, the radius r is repeatedly divided by β until there is a point $p \in P \setminus S$ such that $\text{dist}(p, S) \ge r$. Then the Gonzalez Operation is run, which assigns indices to the centers so that each subsequent center is the farthest from the previously selected centers. Additionally, the Gonzalez Operation assigns labels such that there is an index j for which all centers with index at most j are leaders, and all centers with index greater than j are regular.

Note that the point set P may initially contain at most k points. In this special case, the preprocessing phase continues while $|P| \leq k$, and all points in P become centers. Some invariants might be violated in this special case, but this is not an issue as every point is a center. Once $|P| = k+1$, the computation of the set S and the radius r is performed as described before. Notice that the worst-case recourse is 1, since all points except one belong to S after the computation. Finally, the Gonzalez Operation is executed to determine the sets S_{reg} and S_{ldr}.

Analysis of the preprocessing phase

In this paragraph, we show that all four invariants are satisfied after the preprocessing phase. At the end of the preprocessing phase, we have $|S| = k$ and so Invariant 1 is satisfied. By the choice of S and r given by the method in Lemma 6 and of the preprocessing phase, there is a point $p_S \in P \setminus S$ such that the pairwise distances of points in $S \cup \{p_S\}$ are at least r, as required for Invariant 2. The validity of Invariant 3 follows by the construction of the Gonzalez Operation. Regarding Invariant 4, observe that the radius r is divided by β only when there is no point $p \in P \setminus S$ with $\text{dist}(p, S) \geq r$. Consequently, no point lies at a distance of at least βr from the k-center solution S. Since $1 \leq \beta < \alpha$, for all points $p \in P$ it holds that $\text{dist}(p, S) < \beta r \leq \alpha r + \beta r$. Therefore, once the preprocessing phase has ended, all four invariants hold. Furthermore, the set S is a 2-approximate k-center solution because we run the method of Lemma 6, and adding more points to S can only reduce its approximation ratio.

4.1.2 Adversarial Point Insertion

When the adversary inserts a new point p^+ into the point set P (i.e., $P \coloneqq P \cup \{p^+\}$), the incremental algorithm proceeds as follows. If the new point p^+ is within distance βr from a regular center or within distance αr from a leader center, then the algorithm does nothing. Otherwise, the new point p^+ is at distance greater than βr from all centers and at distance greater than αr from all leader centers. In this scenario, the incremental algorithm continues based on whether the minimum pairwise distance between centers is at most αr.

Case 1: If $\mathit{minDistS} \leq \alpha r$, let c_i be the center that satisfies the following three conditions:
- $\text{dist}(c_i, S \setminus \{c_i\}) = \mathit{minDistS}$.
- c_i is a regular center.
- i is the largest possible index in $S = \{c_1, \ldots, c_k\}$.

The algorithm then replaces c_i in the set of centers S with the new point p^+, and sets $c_i \coloneqq p^+$.[2] Hence, the new point p^+ becomes the i-th center of the ordered set S. Finally, the algorithm labels $c_i = p^+$ as a leader.

Case 2: Otherwise if $\mathit{minDistS} > \alpha r$, let j be the smallest power such that:

$$\min\left(\text{dist}(p^+, S), \mathit{minDistS}\right) \leq \alpha^j r.$$

The algorithm then increases the radius r to $\alpha^{j-1} r$ and invokes the previously described Gonzalez Operation. Once the Gonzalez Operation finishes, the incremental algorithm repeats the initial steps for p^+ (see Line 33 in Algorithm 1), and notice that the algorithm either does nothing or proceeds as in Case 1.

[2] Note that c_i is a variable referring to a point in P. Hence, when we say "replace c_i in S with p^+", we mean replace the value of c_i in S with p^+. Following that, the value of c_i is updated to p^+.

Algorithm 1 INCREMENTAL CONSISTENT k-CENTER.

▷ Let $S = \{c_1, \ldots, c_k\}$ be the ordered set of k centers
▷ Let r be the radius
▷ Let $\mathit{minDistS} \coloneqq \min_{x,y \in S,\, x \neq y} \mathrm{dist}(x, y)$
▷ The algorithm has global access to all its variables and data structures

function GONZALEZOP()
  Let c_1 be an arbitrary center in S
  $j \leftarrow 1$
  for $i \in \{2, \ldots, k\}$ **do**
    $c_i \leftarrow \mathrm{argmax}_{c \in S \setminus \{c_1, \ldots, c_{i-1}\}} \mathrm{dist}(c, \{c_1, \ldots, c_{i-1}\})$
    if $\mathrm{dist}(c_i, \{c_1, \ldots, c_{i-1}\}) > \alpha r$ **then**
      $j \leftarrow i$
    end if
  end for
  $S \leftarrow \{c_1, \ldots, c_k\}$ ▷ Recall that S is an ordered set
  $S_{\mathrm{ldr}} \leftarrow \{c_1, \ldots, c_j\}$
  $S_{\mathrm{reg}} \leftarrow \{c_{j+1}, \ldots, c_k\}$
end function

function INSERTPOINT(p^+)
  Insert p^+ into P (if $p^+ \notin P$)
  if $\mathrm{dist}(p^+, S_{\mathrm{reg}}) \leq \beta r$ **or** $\mathrm{dist}(p^+, S_{\mathrm{ldr}}) \leq \alpha r$ **then**
    exit from INSERTPOINT()
  end if
  if $\mathit{minDistS} \leq \alpha r$ **then** ▷ case C1
    Find $c_i \in S_{\mathrm{reg}}$ that minimizes $\mathrm{dist}(c_i, S \setminus \{c_i\})$ and i is the largest possible index in $S = \{c_1, \ldots, c_k\}$
    Remove c_i from S and S_{reg}
    Add p^+ to S and S_{ldr}
    $c_i \leftarrow p^+$
  else ▷ case C2
    Find the smallest power $j+1$ such that $\min\big(\mathrm{dist}(p^+, S), \mathit{minDistS}\big) \leq \alpha^{j+1} r$
    $r \leftarrow \alpha^j r$
    ▷ New phase begins
    GONZALEZOP()
    INSERTPOINT(p^+)
  end if
end function

4.2 Analysis of the Incremental Algorithm

In this section, we prove Theorem 1 by analyzing the incremental Algorithm 1. We denote by $S^{(\mathrm{old})}$ and $r^{(\mathrm{old})}$ the set of centers and the radius before the adversarial point insertion; the set of centers and the radius after the incremental algorithm processes the adversarial point insertion are denoted by S and r. Hence throughout the analysis, we use (old) as a superscript for the states before the adversarial point insertion, and omit the (old) superscript for the states after the incremental algorithm processes the adversarial point insertion. Recall that α and β satisfy $1 \leq \beta < \alpha$ and $\alpha + \beta \leq \alpha\beta$.

A contiguous block of executions of `InsertPoint()` during which the radius remains unchanged is called a *phase* of Algorithm 1. A phase ends when, upon an adversarial point insertion, the radius increases (i.e., r is updated to at least $\alpha r^{(\mathrm{old})}$); the next phase then begins with the next invocation of `InsertPoint()`.

▶ **Definition 7** (phase of the incremental algorithm). *A new phase begins after the invocation of `InsertPoint()` whenever the radius r has increased (i.e., $r > r^{(old)}$).*

4.2.1 Proof of Invariant 3 and Auxiliary Results

▶ **Lemma 8.** *After an adversarial point insertion, for every regular center* $c_1 \in S_{reg}$ *there is a leader center* $c_2 \in S_{ldr}$ *such that* $\mathrm{dist}(c_1, c_2) \le \alpha r$ *(i.e., Invariant 3 is satisfied).*

Proof. Notice that immediately after an execution of the Gonzalez Operation, a center is regular if and only if it is within distance αr from a leader center. Hence, it suffices to analyze independently each sequence of adversarial point insertions between two consecutive executions of the Gonzalez Operation. The proof proceeds by induction on the adversarial point insertions within a single such sequence. The base case is the beginning of such a sequence, where the statement holds due to the Gonzalez Operation. Whenever a recursive call to `InsertPoint`() after the Gonzalez Operation changes the solution S, the same argument as in the induction step of Case 1 applies.

For the induction step, consider an adversarial point insertion p^+. Observe that the incremental Algorithm 1 could not have proceeded with Case 2, as otherwise the Gonzalez Operation would be invoked, signaling the end of the current sequence. During Case 1, a regular center becomes a leader (see Lines 25 and 26 in Algorithm 1). Hence by the induction hypothesis, for every regular center c_1 there is a leader center c_2 such that $\mathrm{dist}(c_1, c_2) \le \alpha r^{(\mathrm{old})}$. Since in Case 1 the radius does not increase, we have $r = r^{(\mathrm{old})}$, and thus the claim follows. ◀

The following lemma shows that during Case 2 of the incremental Algorithm 1 (in which the Gonzalez Operation is executed), a new phase begins (i.e., the radius increases).

▶ **Lemma 9.** *After an adversarial point insertion, a new phase begins if and only if the Gonzalez Operation is executed.*

Proof. By the construction of the incremental Algorithm 1, the Gonzalez Operation is invoked only during Case 2. Since during Case 1 the radius remains unchanged, it suffices to prove that during Case 2 a new phase always begins. Let p^+ be the new adversarially inserted point. By construction, the incremental algorithm proceeds with Case 2 if and only if $\mathrm{dist}(p^+, S^{(\mathrm{old})}_{\mathrm{reg}}) > \beta r^{(\mathrm{old})}$, $\mathrm{dist}(p^+, S^{(\mathrm{old})}_{\mathrm{ldr}}) > \alpha r^{(\mathrm{old})}$ and $\mathit{minDistS}^{(\mathrm{old})} > \alpha r^{(\mathrm{old})}$.

Because $\mathit{minDistS}^{(\mathrm{old})} > \alpha r^{(\mathrm{old})}$, no two centers are within distance $\alpha r^{(\mathrm{old})}$. Based on Lemma 8 (Invariant 3), every regular center is within distance $\alpha r^{(\mathrm{old})}$ from a leader center, and thus it holds that $S^{(\mathrm{old})}_{\mathrm{reg}} = \emptyset$. Since $S^{(\mathrm{old})} = S^{(\mathrm{old})}_{\mathrm{ldr}} \cup S^{(\mathrm{old})}_{\mathrm{reg}}$, it follows that $\mathrm{dist}(p^+, S^{(\mathrm{old})}) = \mathrm{dist}(p^+, S^{(\mathrm{old})}_{\mathrm{ldr}}) > \alpha r^{(\mathrm{old})}$. Therefore, it holds that $\min\left(\mathrm{dist}(p^+, S^{(\mathrm{old})}), \mathit{minDistS}^{(\mathrm{old})}\right) > \alpha r^{(\mathrm{old})}$, which implies that the power $j+1$ in Line 29 of Algorithm 1 is at least 2. As a result, the radius increases by at least a factor of α in Line 30 (i.e., $r \ge \alpha r^{(\mathrm{old})}$). Since $\alpha > 1$, a new phase begins according to Definition 7, and thus the claim follows. ◀

The next two observations establish the correctness of Line 24 in Algorithm 1, which involves identifying a proper regular center to be sacrificed.

▶ **Observation 10.** *After an adversarial point insertion, for every two leader centers* $c_1, c_2 \in S_{ldr}$ *it holds that* $\mathrm{dist}(c_1, c_2) > \alpha r$.

Proof. Notice that immediately after an execution of the Gonzalez Operation, any two leader centers are at distance greater than αr. Hence, it suffices to analyze each phase independently according to Lemma 9. The proof proceeds by induction on the adversarial point insertions within a single phase. The base case is the beginning of the phase, where the Gonzalez Operation is executed, after which `InsertPoint`() is called recursively. Since the statement holds immediately after the Gonzalez Operation, the argument for the recursive call of `InsertPoint`() is identical to the induction step, where $S^{(\text{old})}$ corresponds to the set of centers immediately after the Gonzalez Operation.

For the induction step, consider an adversarial point insertion p^+. Observe that the set of leader centers S_{ldr} changes only if $\text{dist}(p^+, S^{(\text{old})}_{\text{ldr}}) > \alpha r^{(\text{old})}$. Moreover, the incremental Algorithm 1 must have proceeded with Case 1, as otherwise the next phase begins according to Lemma 9. During Case 1, a center c becomes a leader (i.e., c is added to S_{ldr}) only in Line 26 where $c = p^+$, and also $r = r^{(\text{old})}$. Therefore, together with the induction hypothesis, any two leader centers are at distance greater than αr within a single phase. ◀

The next observation indicates that whenever the incremental Algorithm 1 proceeds with Case 1, a regular center is successfully found to be replaced in the set of centers. This implies the correctness of Line 24 in Algorithm 1.

▶ **Observation 11.** *During Case 1 of the incremental Algorithm 1, there is a regular center $c_i \in S^{(old)}_{reg}$ such that $\text{dist}(c_i, S^{(old)} \setminus \{c_i\}) = minDistS^{(old)}$.*

Proof. Suppose towards a contradiction that $\mathit{minDistS}^{(\text{old})} \le \alpha r^{(\text{old})}$ but for all regular centers $c \in S^{(\text{old})}_{\text{reg}}$ we have $\text{dist}(c, S^{(\text{old})} \setminus \{c\}) > \mathit{minDistS}^{(\text{old})}$. Since $S^{(\text{old})}_{\text{ldr}}$ and $S^{(\text{old})}_{\text{reg}}$ partition $S^{(\text{old})}$, there must be two leader centers $c_1, c_2 \in S^{(\text{old})}_{\text{ldr}}$ such that $\text{dist}(c_1, c_2) = \mathit{minDistS}^{(\text{old})} \le \alpha r^{(\text{old})}$. However, this yields a contradiction based on Observation 10, and thus the claim follows. ◀

4.2.2 Worst-Case Update Time

Observe that the incremental Algorithm 1 in Case 2 recursively calls `InsertPoint`(p^+). The next lemma shows that there is at most one recursive call.

▶ **Lemma 12.** *The recursion depth of `InsertPoint`() is at most two.*

Proof. When Line 33 in Algorithm 1 is executed, then either $\alpha r \ge \text{dist}(p^+, S^{(\text{old})})$ or $\alpha r \ge \mathit{minDistS}^{(\text{old})}$, where r is the updated radius. Observe that $S^{(\text{old})}$ is the set of centers after the Gonzalez Operation in Line 32 and $\mathit{minDistS}^{(\text{old})}$ is the corresponding minimum pairwise distance among centers in $S^{(\text{old})}$. Consider the two cases separately after the recursive call to `InsertPoint`() in Line 33:

- Assume that $\mathit{minDistS}^{(\text{old})} \le \alpha r$. In this case, the incremental algorithm either does nothing due to Line 20, or it proceeds with Case 1, in which no further recursive calls to `InsertPoint`() occur.
- Otherwise, assume that $\mathit{minDistS}^{(\text{old})} > \alpha r$ and $\text{dist}(p^+, S^{(\text{old})}) \le \alpha r$. Let $S^{(\text{interm})}_{\text{ldr}}$ denote the set of leader centers after the execution of the Gonzalez Operation with the updated radius r. By the construction of the Gonzalez Operation, it holds that $S^{(\text{interm})}_{\text{ldr}} = S^{(\text{old})}$. In turn, this implies that $\text{dist}(p^+, S^{(\text{interm})}_{\text{ldr}}) = \text{dist}(p^+, S^{(\text{old})}) \le \alpha r$. Hence, the condition of the if-statement in Line 20 is true, and the incremental algorithm terminates.

Consequently, there is at most one recursive call to `InsertPoint`(), and so the recursion depth of `InsertPoint`() is at most two. ◀

Therefore, the update time is dominated by the Gonzalez Operation that runs Gonzalez's algorithm [12] on the set S, which is of size k by Lemma 14 (Invariant 1).

▶ **Corollary 13.** *The worst-case update time of the incremental Algorithm 1 is $O(k^2)$.*

4.2.3 Proofs of Invariant 1 and Invariant 2

▶ **Lemma 14.** *After an adversarial point insertion, the size of S is k (i.e., Invariant 1 is satisfied). Moreover, the worst-case recourse is* 1.

Proof. In the preprocessing phase, the incremental algorithm constructs a solution S of k centers. During an adversarial point insertion, the set of centers S is altered only when the algorithm proceeds with Case 1. In this case, exactly one point is removed from S and the new adversarially inserted point is added to S (see Lines 25 and 26 in Algorithm 1). As a consequence, we have $|S| = k$ after every adversarial point insertion.

Notice that Case 1 occurs at most once per adversarial point insertion (see also Lemma 12). Thus, the set of centers changes by at most one point per adversarial update. ◀

The following lemma can be credited to Hochbaum and Shmoys [16], who presented a reduction from the k-center clustering problem to a variant of the maximal independent set problem. We denote by R^* the optimal k-center cost of Definition 3 and by $\hat{R}^*$ the optimal improper k-center cost of Definition 4.

▶ **Lemma 15** **([16]).** *Consider a k-center instance (P, k), and let $S \subseteq P$ be a subset of points with $|S| \geq k+1$ such that for every pair of distinct points $p, q \in S$, we have $\operatorname{dist}(p, q) \geq r$. Then it holds that $r \leq 2\hat{R}^*$.*

Proof. Let S^* be an optimal improper k-center solution. Since the size of S^* is at most k, by the pigeonhole principle there exists a center $c^* \in S^*$ and two distinct points $p, q \in S$ such that $\operatorname{dist}(p, c^*) \leq \hat{R}^*$ and $\operatorname{dist}(q, c^*) \leq \hat{R}^*$. By the triangle inequality, we have $\operatorname{dist}(p, c^*) + \operatorname{dist}(q, c^*) \geq \operatorname{dist}(p, q) \geq r$. Hence, without loss of generality we can assume that $\operatorname{dist}(p, c^*) \geq \frac{r}{2}$. Since $\operatorname{dist}(p, c^*) \leq \hat{R}^*$, it follows that $\hat{R}^* \geq \frac{r}{2}$, as needed. ◀

▶ **Observation 16.** *Consider a k-center instance (P, k). Then it holds that $\hat{R}^* \leq R^*$.*

Proof. Observe that every k-center solution $S \subseteq P$ with $|S| \leq k$ is also a feasible solution with the same cost for the improper k-center clustering problem. Hence, comparing the respective optimal k-center solutions, it follows that $\hat{R}^* \leq R^*$. ◀

▶ **Lemma 17.** *After an adversarial point insertion, there is a point $p_S \in P \setminus S$ such that the pairwise distances of points in $S \cup \{p_S\}$ are at least r (i.e., Invariant 2 is satisfied). Moreover, the value of the radius r is at most $2\hat{R}^* \leq 2R^*$.*

Proof. The proof is by induction on the adversarial point insertions. The base case is immediately after the preprocessing phase, where Invariant 2 is satisfied. For the induction step, consider an adversarial point insertion p^+, and let P be the updated point set including p^+. If $\operatorname{dist}(p^+, S^{(\text{old})}_{\text{reg}}) \leq \beta r$ or $\operatorname{dist}(p^+, S^{(\text{old})}_{\text{ldr}}) \leq \alpha r$, then neither the set of centers nor the radius is modified (i.e., $S = S^{(\text{old})}$ and $r = r^{(\text{old})}$). By the induction hypothesis, there is a point $p_S \in (P \setminus \{p^+\}) \setminus S^{(\text{old})}$ such that the pairwise distances of points in $S^{(\text{old})} \cup \{p_S\}$ are at least $r^{(\text{old})}$. Since $S = S^{(\text{old})}$ and $r = r^{(\text{old})}$ in this scenario, there is a point $p_S \in P \setminus S$ such that the pairwise distances of points in $S \cup \{p_S\}$ are at least r, as needed. Otherwise the incremental algorithm proceeds either with Case 1 or Case 2:

- Assume that the incremental algorithm proceeds with Case 1. In this case, it holds that $\operatorname{dist}(p^+, S^{(\text{old})}) > \beta r^{(\text{old})} \geq r^{(\text{old})}$, as $\beta \geq 1$. By the construction of the algorithm, we have $S = (S^{(\text{old})} \setminus \{c\}) \cup \{p^+\}$ for some $c \in S^{(\text{old})}$ and the radius remains unchanged (i.e., $r = r^{(\text{old})}$). By the induction hypothesis, the pairwise distances of points in $S^{(\text{old})}$ are at least $r^{(\text{old})} = r$. Thus by setting $p_S \coloneqq c$ and since $\operatorname{dist}(p^+, S^{(\text{old})}) > r$, we infer that there is a point $p_S \in P \setminus S$ such that the pairwise distances of points in $S \cup \{p_S\}$ are at least r, as required.

- Assume that the incremental algorithm proceeds with Case 2. In this case, the radius r is updated to $\alpha^j r^{(\text{old})}$, where j is the largest power such that $\mathit{minDistS}^{(\text{old})} > \alpha^j r^{(\text{old})}$ and $\text{dist}(p^+, S^{(\text{old})}) > \alpha^j r^{(\text{old})}$. Since $S = S^{(\text{old})}$ at that moment, Invariant 2 holds after the radius r is assigned in Line 30 of Algorithm 1. Next, the incremental algorithm continues recursively with `InsertPoint`(p^+), where $r = \alpha^j r^{(\text{old})}$. Let $S_{\text{reg}}^{(\text{interm})}$ and $S_{\text{ldr}}^{(\text{interm})}$ denote the sets of regular and leader centers immediately after the execution of the Gonzalez Operation.
  - If $\text{dist}(p^+, S_{\text{reg}}^{(\text{interm})}) \leq \beta r$ or $\text{dist}(p^+, S_{\text{ldr}}^{(\text{interm})}) \leq \alpha r$, then the set of centers does not change (i.e., $S = S^{(\text{old})}$) and so Invariant 2 remains satisfied.
  - Otherwise, we have $S = (S^{(\text{old})} \setminus \{c\}) \cup \{p^+\}$ for some $c \in S^{(\text{old})}$. Thus by setting again $p_S := c$ and since $\text{dist}(p^+, S^{(\text{old})}) > r$, we deduce that Invariant 2 remains satisfied. In other words, there is a point $p_S \in P \setminus S$ such that the pairwise distances of points in $S \cup \{p_S\}$ are at least r, as required.

Upper bound on r

By construction, the maintained k-center solution S is a subset of the point set P. Since Invariant 2 and Invariant 1 (see Lemma 14) are satisfied, there are $k + 1$ points in P whose pairwise distances are at least r. Thus by Lemma 15 and Observation 16, it follows that $r \leq 2\hat{R}^* \leq 2R^*$. ◀

4.2.4 Approximation Ratio

In this subsection, we analyze the approximation ratio of our incremental algorithm, using a more fine-grained version of Invariant 4. The set S_{init} denotes the ordered set of centers at the beginning of the current phase; according to Lemma 9, the set S_{init} is the result after the last execution of the Gonzalez Operation. Observe that the set S_{init} may differ from the current set of centers $S^{(\text{old})}$ just before the incremental algorithm processes an adversarial point insertion. We note that S_{init} is an auxiliary set used in the following lemma.

▶ **Lemma 18.** *After an adversarial point insertion, for every point $p \in P$ it holds that either*

$$\text{dist}(p, S_{init}) \leq \beta r \quad or \quad \text{dist}(p, S_{ldr}) \leq \alpha r.$$

Proof. The proof is by induction on the adversarial points insertions. The base case is immediately after the preprocessing phase, where $S_{\text{init}} = S$ and every point $p \in P$ is within distance βr from a center. For the induction step, consider an adversarial point insertion p^+, and let P be the updated point set including p^+. Since $S_{\text{reg}} \subseteq S_{\text{init}}$ at the beginning of each phase, and no new centers are added to S_{reg} during a phase, it follows that $S_{\text{reg}} \subseteq S_{\text{init}}$ throughout the algorithm.

In the scenario where $\text{dist}(p^+, S_{\text{reg}}^{(\text{old})}) \leq \beta r^{(\text{old})}$ or $\text{dist}(p^+, S_{\text{ldr}}^{(\text{old})}) \leq \alpha r^{(\text{old})}$, it holds that $\text{dist}(p^+, S_{\text{init}}) \leq \text{dist}(p^+, S_{\text{reg}}^{(\text{old})})$, $S_{\text{init}} = S_{\text{init}}^{(\text{old})}$, $r = r^{(\text{old})}$, $S_{\text{reg}} = S_{\text{reg}}^{(\text{old})}$, and $S_{\text{ldr}} = S_{\text{ldr}}^{(\text{old})}$, and so together with the induction hypothesis the correctness of the statement follows. Otherwise, the incremental algorithm proceeds either with Case 1 or Case 2:

- Assume that $\mathit{minDistS}^{(\text{old})} \leq \alpha r^{(\text{old})}$ (Case 1). Let $c_i \in S_{\text{reg}}^{(\text{old})}$ and $c_j \in S^{(\text{old})}$ such that $\text{dist}(c_i, c_j) = \mathit{minDistS}^{(\text{old})}$ and i is the largest possible index. Based on Observation 11, such centers c_i and c_j must exist. The algorithm replaces c_i in the set of centers with p^+ and labels p^+ as a leader. Since $S_{\text{reg}}^{(\text{old})} \subseteq S_{\text{init}}$, the old regular center c_i belongs to S_{init}. Hence, as in Case 1 we have $r = r^{(\text{old})}$ and $S_{\text{init}} = S_{\text{init}}^{(\text{old})}$, for all points in $P \setminus \{p^+\}$ the claim holds by the induction hypothesis. Additionally, for the inserted point p^+ we have $\text{dist}(p^+, S_{\text{ldr}}) = 0 \leq \alpha r$, as needed.

- Assume that $minDistS^{(\mathrm{old})} > \alpha r^{(\mathrm{old})}$ (Case 2); the incremental algorithm continues recursively with `InsertPoint`(p^+) in Line 33. According to Lemma 9, a new phases begins which means that $r \geq \alpha r^{(\mathrm{old})}$. For the updated radius r, the inserted point p^+ becomes a center if and only if (i) $\mathrm{dist}(p^+, S_{\mathrm{reg}}^{(\mathrm{old})}) > \beta r$, (ii) $\mathrm{dist}(p^+, S_{\mathrm{ldr}}^{(\mathrm{old})}) > \alpha r$, and (iii) $minDistS^{(\mathrm{old})} \leq \alpha r$. First, we consider the newly inserted point p^+. If p^+ becomes a center during Case 1, then the claim holds for p^+. Otherwise based on Lemma 12, the algorithm does not continue with Case 2 again. Consequently, it holds that either $\mathrm{dist}(p^+, S_{\mathrm{reg}}) \leq \beta r$ or $\mathrm{dist}(p^+, S_{\mathrm{ldr}}) \leq \alpha r$, and since $S_{\mathrm{reg}} \subseteq S_{\mathrm{init}}$ the claim follows for p^+. It remains to consider the remaining points in $P \setminus \{p^+\}$. Since a new phase begins during Case 2, we have $r \geq \alpha r^{(\mathrm{old})}$ and the set S_{init} changes. Let $S_{\mathrm{init}}^{(\mathrm{new})}$ be the set of centers after the Gonzalez Operation is invoked (i.e., $S_{\mathrm{init}}^{(\mathrm{new})} = S^{(\mathrm{old})}$). Based on Lemma 8 (Invariant 3) and the fact that $minDistS^{(\mathrm{old})} > \alpha r^{(\mathrm{old})}$, we can infer that $S_{\mathrm{ldr}}^{(\mathrm{old})} = S^{(\mathrm{old})}$, which means that $S_{\mathrm{ldr}}^{(\mathrm{old})} = S^{(\mathrm{old})} = S_{\mathrm{init}}^{(\mathrm{new})}$. For an arbitrary point $q \in P \setminus \{p^+\}$, by the induction hypothesis either $\mathrm{dist}(q, S_{\mathrm{init}}^{(\mathrm{old})}) \leq \beta r^{(\mathrm{old})}$ or $\mathrm{dist}(q, S_{\mathrm{ldr}}^{(\mathrm{old})}) \leq \alpha r^{(\mathrm{old})}$; we analyze the two cases separately:
  - Consider a point $q \in P \setminus \{p^+\}$ such that $\mathrm{dist}(q, S^{(\mathrm{old})}) \leq \alpha r^{(\mathrm{old})}$. Then as $\beta \geq 1$, it holds that:
  $$\mathrm{dist}(q, S_{\mathrm{init}}^{(\mathrm{new})}) \leq \alpha r^{(\mathrm{old})} \leq r \leq \beta r.$$
  - Consider a point $q \in P \setminus \{p^+\}$ such that $\mathrm{dist}(q, S_{\mathrm{init}}^{(\mathrm{old})}) \leq \beta r^{(\mathrm{old})}$ and $\mathrm{dist}(q, S^{(\mathrm{old})}) > \alpha r^{(\mathrm{old})}$. Observe that leader centers are never removed from the set of centers during a (fixed) phase and that the distance from q to the set of centers increased during the (previous) phase. Hence, the closest center $c_q \in S_{\mathrm{init}}^{(\mathrm{old})}$ to q that was removed from the set of centers was a regular center (when the radius was $r^{(\mathrm{old})}$). By Lemma 8 (Invariant 3), there must be another leader center $c \in S_{\mathrm{ldr}}^{(\mathrm{old})}$ such that $\mathrm{dist}(c_q, c) \leq \alpha r^{(\mathrm{old})}$. By the triangle inequality and since $\alpha + \beta \leq \alpha\beta$, we obtain that:
  $$\mathrm{dist}(q, S_{\mathrm{init}}^{(\mathrm{new})}) \leq \alpha r^{(\mathrm{old})} + \beta r^{(\mathrm{old})} \leq \alpha\beta r^{(\mathrm{old})} \leq \beta r.$$

  Since the algorithm is currently in Case 2, based on Lemma 12, the next recursive call of `InsertPoint`(p^+) does not enter Case 2 again. This implies that the radius is no longer modified by this adversarial point insertion and the set S_{init} remains unchanged. If the solution changes due to Case 1 of the recursive `InsertPoint`(p^+), the claim follows using the same arguments as before. ◀

To analyze the approximation ratio, we split every phase into two *subphases* and analyze each subphase separately; the subphases are defined as follows.

▶ **Definition 19** (subphases of a phase)**.** *During a phase of the incremental algorithm, the first subphase ends when the minimum pairwise distance between centers is greater than βr; the second subphase then begins.*

▶ **Observation 20.** *During a phase of the incremental algorithm, any $c \in S \setminus S_{init}$ that becomes a center after the beginning of the phase is at a distance greater than βr from all other centers until the end of the phase.*

Proof. According to Lemma 9, whenever the Gonzalez Operation is invoked during Case 2 of the incremental Algorithm 1, a new phase begins. Hence, it suffices to analyze only Case 1 of the incremental Algorithm 1. The proof is by induction on the adversarial point insertions within a single phase. The base case is at the beginning of the phase, where $S = S_{\mathrm{init}}$ and the statement trivially holds. For the induction step, consider an adversarial point insertion p^+. The set of centers is modified only in Case 1, which occurs only when the if-statement in Line 20 is false. This implies that for the new center p^+ (see Line 26 in Algorithm 1), it holds that $\mathrm{dist}(p^+, S_{\mathrm{reg}}^{(\mathrm{old})}) > \beta r$ and $\mathrm{dist}(p^+, S_{\mathrm{ldr}}^{(\mathrm{old})}) > \alpha r$, since $r = r^{(\mathrm{old})}$ during a phase. In turn, it holds that $\mathrm{dist}(p^+, S^{(\mathrm{old})}) = \mathrm{dist}(p^+, S_{\mathrm{reg}}^{(\mathrm{old})} \cup S_{\mathrm{ldr}}^{(\mathrm{old})}) > \beta r$. For any other center $c \in S^{(\mathrm{old})} \setminus S_{\mathrm{init}}$, the statement holds by the induction hypothesis, as $r = r^{(\mathrm{old})}$. ◀

Hence, during the first subphase, the value *minDistS* is defined by two centers in $S \cap S_{\mathrm{init}}$. From the start of the second subphase onward, *minDistS* exceeds βr until the end of the phase. This is because Observation 20 applies for any new center that is added until the end of the phase. Let $S_i \coloneqq \{c_j \in S_{\mathrm{init}} \mid j \leq i\}$ denote the ordered subset of S_{init} up to center c_i, for any $1 \leq i \leq k$.

▶ **Lemma 21.** *Consider a fixed phase of the incremental algorithm. After an adversarial point insertion during the first subphase, assume that the incremental algorithm removes the i-th center c_i from the set of centers. Then the distance $\mathrm{dist}(c_i, S_{i-1})$ equals the minimum pairwise distance between centers at that moment.*

Proof. Based on Observation 20 and Definition 19, the minimum pairwise distance between centers is attained between two centers in S_{init}. By the construction of the Gonzalez Operation, all leader centers have smaller index than regular centers in S_{init}. Since i is the largest such index in Line 24 of Algorithm 1 (during Case 1), it holds that $\mathrm{dist}(c_i, S_{i-1}) \leq \mathrm{dist}(c_i, S^{(\mathrm{old})} \setminus \{c_i\})$, where c_i refers to the old location determined in Line 24. Based on Observation 11, we have $\mathrm{dist}(c_i, S^{(\mathrm{old})} \setminus \{c_i\}) = \mathit{minDistS}^{(\mathrm{old})}$. Thus at the moment that c_i is removed from the set of centers, the minimum pairwise distance between centers is attained between c_i and S_{i-1}. ◀

▶ **Lemma 22.** *Consider a fixed phase of the incremental algorithm. After an adversarial point insertion during the first subphase, assume that the incremental algorithm removes the i-th center c_i from the set of centers. Then it holds that $S_{i-1} \subseteq S^{(old)}$ and $S_{i-1} \subseteq S$.*

Proof. After an adversarial insertion of a point p^+, the center c_i is replaced by p^+ in the set of centers (i.e., $S = (S^{(\mathrm{old})} \setminus \{c_i\}) \cup \{p^+\}$). Since $p^+ \notin S_{i-1}$, it suffices to show that $S_{i-1} \subseteq S$. Suppose to the contrary that there is an adversarial point insertion such that c_i (referring to the old location) is removed from the set of centers and $S_{i-1} \not\subseteq S$, and consider the first such adversarial point insertion. Let c_j be the center (referring to the old location) with the smallest index j in S_{init} which has been removed from the set of centers in a previous adversarial point insertion (i.e., $c_j \in S_i \setminus S$ with the smallest possible index j). By the minimality of the counter-example, it holds that $S_{j-1} \subseteq S$ and $j < i$, which implies that c_i does not belong to S_{j-1}. By the construction of the Gonzalez Operation, we have $\mathrm{dist}(c_i, S_{j-1}) \leq \mathrm{dist}(c_j, S_{j-1})$.

Let τ be the moment when c_j was removed from the set of centers. Since the phase is fixed, the τ-th adversarial point insertion occurs during the same first subphase. Based on Lemma 21, the minimum pairwise distance in Line 24 of Algorithm 1 (during Case 1) at time τ is achieved between c_j and S_{j-1}. However, as $\mathrm{dist}(c_i, S_{j-1}) \leq \mathrm{dist}(c_j, S_{j-1})$ at the τ-th adversarial point insertion, the minimum pairwise distance between centers at time τ could also be achieved with c_i. Since $i > j$, the incremental algorithm should have replaced c_i before replacing c_j, which leads to a contradiction. ◀

Recall that the optimal k-center cost of Definition 3 is denoted by R^*, and the optimal improper k-center cost of Definition 4 is denoted by $\hat{R}^*$.

▶ **Observation 23.** *Consider a k-center instance (P, k) subject to adversarial point insertions. Then $\hat{R}^*$ is non-decreasing at all times.*

Proof. Consider an optimal improper k-center solution $\hat{S}_1^*$ before the adversarial point insertion, where $\mathrm{cost}^{(\mathrm{old})}(\hat{S}_1^*) = (\hat{R}^*)^{(\mathrm{old})}$. Suppose to the contrary that after an adversarial point insertion, the optimal improper k-center cost decreases (i.e., $\hat{R}^* < (\hat{R}^*)^{(\mathrm{old})}$). Under adversarial point insertions, the cost of $\hat{S}_1^*$ is non-decreasing (i.e., $\mathrm{cost}^{(\mathrm{old})}(\hat{S}_1^*) \leq \mathrm{cost}(\hat{S}_1^*)$). Thus, there must be another optimal k-center solution $\hat{S}_2^*$ after the adversarial point insertion, where $\mathrm{cost}(\hat{S}_2^*) = \hat{R}^*$.

Since $\hat{S}_2^*$ is a subset of the metric space and $|\hat{S}_2^*| \leq k$, the set $\hat{S}_2^*$ is also a feasible improper k-center solution before the adversarial point insertion. The cost of $\hat{S}_2^*$ is non-decreasing under adversarial point insertions, and so $\mathrm{cost}^{(\mathrm{old})}(\hat{S}_2^*) \leq \mathrm{cost}(\hat{S}_2^*)$. Combining all these together, it follows that $\mathrm{cost}^{(\mathrm{old})}(\hat{S}_2^*) \leq \mathrm{cost}(\hat{S}_2^*) = \hat{R}^* < (\hat{R}^*)^{(\mathrm{old})}$, contradicting the fact that $\hat{S}_1^*$ was an optimal improper k-center solution before the adversarial point insertion. Therefore, the optimal improper k-center cost cannot decrease when a new point is inserted by the adversary. ◀

▶ **Lemma 24.** *Let $\alpha := \frac{3+\sqrt{5}}{2}$ and $\beta := \frac{1+\sqrt{5}}{2}$. After an adversarial point insertion, for every point $p \in P$ it holds that:*

$$\mathrm{dist}(p, S) \ \leq\ (3+\sqrt{5})\hat{R}^* \ \leq\ (3+\sqrt{5})R^*.$$

Proof. Based on Lemma 18, every point $p \in P$ is either within distance βr from the set S_{init} or within distance αr from a leader center in S_{ldr}. In the latter case, by Lemma 17 we have $r \leq 2\hat{R}^* \leq 2R^*$, and as $S_{\mathrm{ldr}} \subseteq S$ it follows that $\mathrm{dist}(p, S) \leq \alpha r \leq 2\alpha\hat{R}^* = (3+\sqrt{5})\hat{R}^* \leq (3+\sqrt{5})R^*$, as needed. Since $\beta < \alpha$, the same conclusion holds for any point $p \in P$ satisfying $\mathrm{dist}(p, c) \leq \beta r$ for some center $c \in S_{\mathrm{init}} \cap S_{\mathrm{reg}}$. Notice that when a new phase begins then $S_{\mathrm{init}} = S_{\mathrm{reg}} \cup S_{\mathrm{ldr}}$, implying the correctness of the statement by Lemma 18 together with the fact that $r \leq 2\hat{R}^* \leq 2R^*$. Hence, it remains to analyze points $p \in P$ such that $\mathrm{dist}(p, S_{\mathrm{ldr}}) > \alpha r$ and $\mathrm{dist}(p, S_{\mathrm{init}}) = \mathrm{dist}(p, q) \leq \beta r$ for some $q \in S_{\mathrm{init}} \setminus S$ during a fixed phase. Namely, the point q has been replaced as a center during this phase.

The proof is by induction on the adversarial point insertions within a single phase. The base case is at the beginning of the phase, where $S_{\mathrm{init}} = S$ and the statement holds as explained earlier. For the induction step, consider an adversarial point insertion p^+, and let P be the updated point set including p^+. In the scenario where $\mathrm{dist}(p^+, S_{\mathrm{reg}}^{(\mathrm{old})}) \leq \beta r^{(\mathrm{old})}$ or $\mathrm{dist}(p^+, S_{\mathrm{ldr}}^{(\mathrm{old})}) \leq \alpha r^{(\mathrm{old})}$, neither the set of centers nor the radius is modified (i.e., $S = S^{(\mathrm{old})}$ and $r = r^{(\mathrm{old})}$). Thus the statement holds by the induction hypothesis combined with Observation 23 and Observation 16. Otherwise it holds that $\mathrm{dist}(p^+, S^{(\mathrm{old})}) > \beta r^{(\mathrm{old})}$, and the incremental algorithm continues either with Case 1 or Case 2. According to Lemma 9, during Case 2 of the incremental Algorithm 1, a new phase begins. Therefore, we only need to consider adversarial point insertions for which the incremental algorithm proceeds with Case 1 in the current phase.

For clarity regarding the reset of c_i in Line 27 of Algorithm 1, let $c_i^{(\mathrm{old})}$ denote the old center (referring to the old location) determined in Line 24. Since the radius does not increase in Case 1, we have $r = r^{(\mathrm{old})}$. Observe that the newly inserted point p^+ becomes a leader in Line 26 of Algorithm 1, and by construction we have $S_{\mathrm{reg}} = S_{\mathrm{reg}}^{(\mathrm{old})} \setminus \{c_i^{(\mathrm{old})}\}$ and $S_{\mathrm{ldr}} = S_{\mathrm{ldr}}^{(\mathrm{old})} \cup \{p^+\}$ (i.e., $S = (S^{(\mathrm{old})} \setminus \{c_i^{(\mathrm{old})}\}) \cup \{p^+\}$). This implies that for all points whose nearest center in $S^{(\mathrm{old})}$ is not $c_i^{(\mathrm{old})}$, the claim follows by the induction hypothesis combined with Observation 23 and Observation 16. Hence, it suffices to consider points $p \in P^{(\mathrm{old})}$ whose nearest center in $S^{(\mathrm{old})}$ is the point $c_i^{(\mathrm{old})}$. For a fixed such point $p \in P^{(\mathrm{old})}$, let $q \in S_{\text{init}}$ be its nearest center at the beginning of the current phase. Recall that for a point $p \in P$, its closest point q in S_{init} may not belong to the set of centers before or after the adversarial update (i.e., S_{init} can differ from $S^{(\text{old})}$ and S). There are two cases to analyze depending on which subphase of the incremental Algorithm 1 is active:

The first subphase is active. Let i be the index determined in Line 24 of Algorithm 1. Since the first subphase is active (see Definition 19), it holds that $\text{dist}(c_i^{(\text{old})}, S^{(\text{old})} \setminus \{c_i^{(\text{old})}\}) = \textit{minDistS}^{(\text{old})} \leq \beta r$, which implies that $\text{dist}(c_i^{(\text{old})}, S^{(\text{old})} \setminus \{c_i^{(\text{old})}\}) < \text{dist}(p^+, S^{(\text{old})})$. Since the distance between $c_i^{(\text{old})}$ and $S^{(\text{old})} \setminus \{c_i^{(\text{old})}\}$ is the minimum over all pairwise distances of points in $S^{(\text{old})}$, the pairwise distances of points in $S^{(\text{old})} \cup \{p^+\}$ are at least $\text{dist}(c_i^{(\text{old})}, S^{(\text{old})} \setminus \{c_i^{(\text{old})}\})$. Based on Lemma 14 we have $|S^{(\text{old})} \cup \{p^+\}| = k+1$, and observe that $S^{(\text{old})} \cup \{p^+\} \subseteq P$. As a consequence, by Lemma 15 it holds that $\text{dist}(c_i^{(\text{old})}, S^{(\text{old})} \setminus \{c_i^{(\text{old})}\}) \leq 2\hat{R}^*$.

Observe that if $q = c_i^{(\text{old})}$ then by the triangle inequality and by Observation 16, we have $\text{dist}(p, S) \leq \beta r + \text{dist}(q, S) \leq \beta r + 2\hat{R}^* \leq (1+\beta)2\hat{R}^* = (3+\sqrt{5})\hat{R}^* \leq (3+\sqrt{5})R^*$, as needed for p. Otherwise $q \notin S^{(\text{old})}$ and based on Lemma 22, we can infer that $q \notin S_{i-1}$. Due to the order induced by the Gonzalez Operation, we have $\text{dist}(q, S_{i-1}) \leq \text{dist}(c_i^{(\text{old})}, S_{i-1})$. Again by Lemma 22, it follows that $\text{dist}(q, S) \leq \text{dist}(q, S_{i-1})$, and so $\text{dist}(q, S) \leq \text{dist}(c_i^{(\text{old})}, S_{i-1})$. Based on Lemma 21 (note that i is the largest such index in Line 24), it holds that $\text{dist}(c_i^{(\text{old})}, S_{i-1}) \leq \text{dist}(c_i^{(\text{old})}, S^{(\text{old})} \setminus \{c_i^{(\text{old})}\})$. Combining these inequalities, we conclude that $\text{dist}(q, S) \leq \text{dist}(c_i^{(\text{old})}, S^{(\text{old})} \setminus \{c_i^{(\text{old})}\})$. Since we previously showed that $\text{dist}(c_i^{(\text{old})}, S^{(\text{old})} \setminus \{c_i^{(\text{old})}\}) \leq 2\hat{R}^*$, we deduce that $\text{dist}(q, S) \leq 2\hat{R}^*$. As a result, by the triangle inequality it follows that $\text{dist}(p, S) \leq 2\hat{R}^* + \beta r$. By Lemma 17 we have $r \leq 2\hat{R}^* \leq 2R^*$, and thus $\text{dist}(p, S) \leq (3+\sqrt{5})\hat{R}^* \leq (3+\sqrt{5})R^*$, as required for p.

The second subphase is active. At any moment during the second subphase, the minimum pairwise distance between centers is greater than $\beta r^{(\text{old})}$ (i.e., $\textit{minDistS}^{(\text{old})} > \beta r^{(\text{old})}$). As $r = r^{(\text{old})}$ we have $\text{dist}(p^+, S^{(\text{old})}) > \beta r$, which means that the pairwise distances of points in $S^{(\text{old})} \cup \{p^+\}$ are at least βr. According to Lemma 14 we have $|S^{(\text{old})} \cup \{p^+\}| = k+1$, and thus by Lemma 15 it holds that $\beta r \leq 2\hat{R}^*$. Note that the point q belongs to S_{init} but is no longer a center (i.e., q was already replaced by another point in the set of centers during this phase). By the construction of the incremental Algorithm 1, only regular centers are removed from the set of centers. In turn, according to Lemma 8 (Invariant 3) and since the radius remains unchanged during a fixed phase, there is a leader center within distance αr from the point q. Therefore, by the triangle inequality and as $r \leq \frac{2\hat{R}^*}{\beta}$ and $\hat{R}^* \leq R^*$, it follows that $\text{dist}(p, S) \leq \alpha r + \beta r \leq \frac{\alpha}{\beta} 2\hat{R}^* + 2\hat{R}^* = (\frac{\alpha}{\beta} + 1)2\hat{R}^* = (3+\sqrt{5})\hat{R}^* \leq (3+\sqrt{5})R^*$, as required for p. ◀

Proof of Theorem 1. The proof follows by Lemmas 14 and 24 and Corollary 13. ◀

5 Lower Bound on Worst-Case Recourse

In this section, we show a lower bound on the worst-case recourse. In the following, we give the proof for the incremental setting. For the decremental setting, this proof can be "reversed", starting with p^+ as part of the original point set. Then, the adversary deletes that point and the original setting of the following proof needs to be restored, requiring k centers to be moved.

▶ **Theorem 2.** *Consider a point set P in a metric space subject to point insertions, an integer $k \geq 1$, and an approximation factor $\alpha \in [1, \sqrt{2})$. Then any partially dynamic α-approximate k-center algorithm has a worst-case recourse of k.*

Proof. Consider the following k-center instance containing $3k-2$ points $p_1, p_2, \ldots, p_{3k-2}$ such that:

- $\mathrm{dist}(p_1, p_2) = \mathrm{dist}(p_2, p_3) = d'$,
- $\mathrm{dist}(p_{i-1}, p_i) = 2d'$, for all $i \geq 4$ with $i \equiv 1 \pmod 3$,
- $\mathrm{dist}(p_{i-1}, p_i) = d$, for all $i \geq 5$ with $i \not\equiv 1 \pmod 3$, and
- $\mathrm{dist}(p_1, p_i) = \mathrm{dist}(p_1, p_{i-1}) + \mathrm{dist}(p_{i-1}, p_i)$ for all i,

where d and d' satisfy $\max(\frac{2}{3}, \frac{\alpha}{2})d < d' < \frac{d}{\alpha}$.

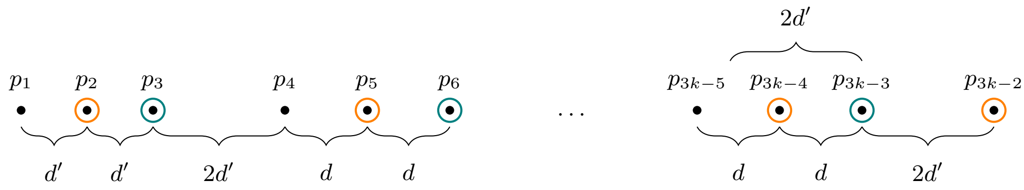

Figure 1 The k-center instance given to an α-approximation algorithm before the adversarial insertion of the point p^+. The orange circles mark the α-approximate placement of the k centers. The teal circles mark the α-approximate placement of $k-1$ centers.

First, we show that there is a unique α-approximate k-center solution. This k-center solution consists of the $k-1$ points p_i with $i \equiv 2 \pmod 3$ and the last point at p_{3k-2} (see Figure 1); the k-center cost of this solution is d. Regarding the uniqueness, each triple p_i, p_{i+1}, p_{i+2} (for every $1 \leq i \leq 3k-5$ with $i \equiv 1 \pmod 3$) and the point p_{3k-2} must contain a center, as otherwise the cost would be at least $2d'$. Moreover, each such center must be placed at the middle point of each triple, as otherwise the cost would be at least $2d'$ (for the leftmost triple) or at least $\min(2d, 2d') = 2d'$ (for any other triple). A strict lower bound on $2d'$ is:

$$2d' \; > \; 2\frac{\alpha}{2}d \; > \; \alpha d,$$

and thus placing a center at the middle point of each triple and a center at p_{3k-2} is the only possible α-approximate k-center solution.

Assume now that the adversary inserts a point p^+ at a distance greater than $2\alpha d$ from any other point $p_1, p_2, \ldots, p_{3k-2}$. We show that there is a unique α-approximate k-center solution whose cost is $2d'$. To obtain such a solution, one center has to move to p^+, as otherwise the cost would be at least $2\alpha d > \alpha \cdot 2d'$. Hence, we have $k-1$ centers remaining to cover the original $3k-2$ points. To obtain this unique k-center solution, the center from p_{3k-2} is moved to p^+ and all the remaining centers are shifted to the right; namely the $k-1$ remaining centers are placed at p_i where $i \equiv 0 \pmod 3$ (see points marked in teal in Figure 1). The cost of this solution is $2d'$ because $2d < 3d'$ and thus $d < 2d'$; the leftmost point in each triple (except the leftmost one) is covered by the center in the left neighboring triple.

We now show that no other placement of $k-1$ centers among the old points can achieve a cost smaller or equal to $\alpha \cdot 2d'$. Leaving a triple without a center would lead to a cost of at least:

$$\min(4d', 2d' + d) \; > \; 2d \; > \; \alpha \cdot 2d'.$$

Hence, for any α-approximate solution each triple has to have exactly one center. Suppose to the contrary that in the triple p_i, p_{i+1}, p_{i+2} for some $i \equiv 1 \pmod 3$ the center is not at p_{i+2}. If it is the rightmost triple ($i = 3k-5$), the cost of the solution is at least $2d' + d > 3d' > \alpha \cdot 2d'$.

Otherwise, consider the largest i with $i \equiv 1 \pmod 3$ such that the center for the triple p_i, p_{i+1}, p_{i+2} is not at p_{i+2}. If $i \neq 1$, the cost would be at least $\min(d + 2d', 2d) = 2d > 2\alpha d'$. If $i = 1$, the cost would be at least $\min(3d', 2d) = 2d > 2\alpha d'$.

After the adversarial insertion of p^+, the unique α-approximate solution is given by the set of centers $\{p_i \mid 3 \leq i \leq 3k-3, i \equiv 0 \pmod 3\} \cup \{p^+\}$. As a consequence, the recourse of the solution before and after the adversarial point insertion is k, concluding the proof. ◀

References

1 Kiarash Banihashem, Jeff Giliberti, Samira Goudarzi, MohammadTaghi Hajiaghayi, Peyman Jabbarzade, and Morteza Monemizadeh. Dynamic diameter in high-dimensions against adaptive adversary and beyond, 2025. doi:10.48550/arXiv.2511.01065.

2 MohammadHossein Bateni, Hossein Esfandiari, Hendrik Fichtenberger, Monika Henzinger, Rajesh Jayaram, Vahab Mirrokni, and Andreas Wiese. Optimal fully dynamic k-center clustering for adaptive and oblivious adversaries. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 2677–2727. SIAM, 2023. doi:10.1137/1.9781611977554.CH101.

3 Aaron Bernstein, Sebastian Forster, and Monika Henzinger. A deamortization approach for dynamic spanner and dynamic maximal matching. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 1899–1918. SIAM, 2019. doi:10.1137/1.9781611975482.115.

4 Sayan Bhattacharya, Ruoxu Cen, and Debmalya Panigrahi. Fully dynamic set cover: Worst-case recourse and update time. *CoRR*, abs/2511.08485, 2025. To appear in STOC'26. doi:10.48550/arXiv.2511.08485.

5 Sayan Bhattacharya, Martín Costa, and Ermiya Farokhnejad. Fully dynamic k-median with near-optimal update time and recourse. In Michal Koucký and Nikhil Bansal, editors, *Proceedings of the 57th Annual ACM Symposium on Theory of Computing, STOC 2025, Prague, Czechia, June 23-27, 2025*, pages 1166–1177. ACM, 2025. doi:10.1145/3717823.3718176.

6 Sayan Bhattacharya, Martín Costa, Ermiya Farokhnejad, Silvio Lattanzi, and Nikos Parotsidis. Almost optimal fully dynamic k-center clustering with recourse. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu, editors, *Forty-second International Conference on Machine Learning, ICML 2025, Vancouver, BC, Canada, July 13-19, 2025*, volume 267 of *Proceedings of Machine Learning Research*. PMLR / OpenReview.net, 2025. URL: https://proceedings.mlr.press/v267/bhattacharya25a.html.

7 Moses Charikar, Chandra Chekuri, Tomás Feder, and Rajeev Motwani. Incremental clustering and dynamic information retrieval. In Frank Thomson Leighton and Peter W. Shor, editors, *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, El Paso, Texas, USA, May 4-6, 1997*, pages 626–635. ACM, 1997. doi:10.1145/258533.258657.

8 Moses Charikar and Shay Solomon. Fully dynamic almost-maximal matching: Breaking the polynomial worst-case time barrier. In Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella, editors, *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, Prague, Czech Republic, July 9-13, 2018*, LIPIcs, pages 33:1–33:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.ICALP.2018.33.

9 Vincent Cohen-Addad, Chris Schwiegelshohn, and Christian Sohler. Diameter and k-center in sliding windows. In Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi, editors, *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*, volume 55 of *LIPIcs*, pages 19:1–19:12. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPIcs.ICALP.2016.19.

10 Hendrik Fichtenberger, Silvio Lattanzi, Ashkan Norouzi-Fard, and Ola Svensson. Consistent k-clustering for general metrics. In Dániel Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 2660–2678. SIAM, 2021. doi:10.1137/1.9781611976465.158.

11 Sebastian Forster and Antonis Skarlatos. Dynamic consistent k-center clustering with optimal recourse. In Yossi Azar and Debmalya Panigrahi, editors, *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2025, New Orleans, LA, USA, January 12-15, 2025*, pages 212–254. SIAM, 2025. doi:10.1137/1.9781611978322.7.

12 Teofilo F. Gonzalez. Clustering to minimize the maximum intercluster distance. *Theor. Comput. Sci.*, 38:293–306, 1985. doi:10.1016/0304-3975(85)90224-5.

13 Fabrizio Grandoni, Chris Schwiegelshohn, Shay Solomon, and Amitai Uzrad. Maintaining an EDCS in general graphs: Simpler, density-sensitive and with worst-case time bounds. *ACM Trans. Algorithms*, 22(1):10:1–10:16, 2026. doi:10.1145/3765900.

14 Mara Grilnberger and Antonis Skarlatos. Adaptive fully dynamic k-center clustering with (near-)optimal worst-case guarantees, 2026. doi:10.48550/arXiv.2604.01726.

15 Sudipto Guha. Tight results for clustering and summarizing data streams. In Ronald Fagin, editor, *Database Theory - ICDT 2009, 12th International Conference, St. Petersburg, Russia, March 23-25, 2009, Proceedings*, ACM International Conference Proceeding Series, pages 268–275. ACM, 2009. doi:10.1145/1514894.1514926.

16 Dorit S. Hochbaum and David B. Shmoys. A unified approach to approximation algorithms for bottleneck problems. *J. ACM*, 33(3):533–550, 1986. doi:10.1145/5925.5933.

17 Wen-Lian Hsu and George L. Nemhauser. Easy and hard bottleneck location problems. *Discret. Appl. Math.*, 1(3):209–215, 1979. doi:10.1016/0166-218X(79)90044-1.

18 László Kozma. The price of incrementality in k-center clustering, 2026. doi:10.48550/arXiv.2606.08713.

19 Silvio Lattanzi and Sergei Vassilvitskii. Consistent k-clustering. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 1975–1984. PMLR, 2017. URL: http://proceedings.mlr.press/v70/lattanzi17a.html.

20 Shay Solomon and Amitai Uzrad. Dynamic set cover with worst-case recourse. *CoRR*, abs/2511.07354, 2025. To appear in ICALP'26. doi:10.48550/arXiv.2511.07354.