

UNIVERSITÀ  
DEGLI STUDI  
DI PADOVA



UNIVERSITY OF PADUA  
DEPARTMENT OF INFORMATION ENGINEERING  
MASTER'S DEGREE IN COMPUTER ENGINEERING

# Dynamic $k$ -Center Clustering with Lifetimes

**Supervisor**

Prof. Pietracaprina Andrea

**Candidate**

Moretti Simone

**Co-supervisor**

Prof. Pucci Geppino

**Student ID**

2139223

ANNO ACCADEMICO 2025-2026

Data di laurea 07/07/2026



# Abstract

The  $k$ -center problem is a fundamental clustering variant with applications in learning systems and data summarization. In several real-world scenarios, the dataset to be clustered is not static, but evolves over time, as new data points arrive and old ones become stale. To account for dynamicity, the  $k$ -center problem has been mainly studied under the sliding window setting, where only the  $N$  most recent points are considered non-stale, or the fully dynamic setting, where arbitrary sequences of point arrivals and deletions without prior notice may occur. In this thesis, we introduce the dynamic setting with lifetimes, which bridges the two aforementioned classical settings by still allowing arbitrary arrivals and deletions, but making the deletion time of each point known upon its arrival. Under this new setting, we devise a deterministic  $(2+\varepsilon)$ -approximation algorithm with  $O(k/\varepsilon)$  amortized update time and memory usage linear in the number of currently active points. Moreover, we develop a deterministic  $(6+\varepsilon)$ -approximation algorithm that, under "tame" update sequences, has  $O(k/\varepsilon)$  worst-case update time and heavily sublinear working memory.



# Contents

<b>Introduction</b>	<b>1</b>
<b>1 Preliminaries</b>	<b>5</b>
1.1 Clustering . . . . .	5
1.1.1 $k$ -Center Clustering . . . . .	6
1.2 Dynamic Algorithms . . . . .	8
<b>2 Dynamic <math>k</math>-Center: Minimizing Approximation</b>	<b>11</b>
2.1 A $O(k^2)$ Fully Dynamic Algorithm . . . . .	12
2.2 The State-of-the-Art of Fully Dynamic $k$ -Center . . . . .	14
2.3 With lifetimes . . . . .	19
2.3.1 Approximation . . . . .	24
2.3.2 Running time and memory complexity . . . . .	26
2.4 Bonus: A new fully dynamic algorithm for $k$ -center . . . . .	31
2.4.1 Comparison with previous work . . . . .	32
<b>3 Dynamic <math>k</math>-Center with Sublinear Memory</b>	<b>33</b>
3.1 Sliding Window . . . . .	33
3.2 With Lifetimes . . . . .	37
3.2.1 Approximation . . . . .	39
3.2.2 Running time and memory complexity . . . . .	40
3.2.3 Comparison with related work . . . . .	42
<b>4 Experiments</b>	<b>45</b>
4.1 Approximation . . . . .	47
4.2 Running Time . . . . .	49
4.3 Memory Consumption . . . . .	51
4.4 Conclusions . . . . .	52



# Introduction

Clustering is one of the most basic and widely used tools for data analysis. Given a collection of objects with a notion of distance measuring how dissimilar they are, a clustering algorithm groups the objects so that similar ones end up together and dissimilar ones get grouped apart. This shows structure and produces useful information of an otherwise useless dataset. Among the many formalizations of this intuition, the *k-center* problem occupies a prominent place: given a set of points in a metric space and a number  $k$ , it asks to find  $k$  points, called *centers*, that minimize the maximum distance from any point to its closest center. It is a natural primitive in facility location, data summarization, coresets construction, feature selection, outlier detection and it is frequently used as a building block inside larger learning pipelines.

The problem is *NP*-hard, and, unless  $P = NP$ , no polynomial-time algorithm can approximate its optimal radius within a factor smaller than 2 [1]. This lower bound is matched by a remarkably simple greedy algorithm due to Gonzalez [1], which repeatedly selects as the next center the point farthest from those chosen so far and attains a 2-approximation in  $O(nk)$  time.

Furthermore, in a growing number of applications, the dataset to be clustered is not fixed, but *evolves over time*. A series of *insertions* and *deletions* of points are provided in input, along with *queries* that ask for a clustering of the current *active* set of points at time  $t$ :  $X_t$ . Recomputing a clustering from scratch for each query is too expensive, so we seek *dynamic* algorithms that maintain a good solution incrementally, updating a small amount of internal state per operation and spending time and memory that scale with the active set rather than with the entire history of the data.

Insertions and deletions can be very structured or arbitrary, giving rise to two opposed models, well analyzed in the literature. In the *sliding-window* model, a point is considered active only as long as it is among the  $N$  most recent arrivals: expiration is completely determined by arrival order, so points leave in exactly the order in which they entered. Cohen-Addad et al. [2] obtained a  $(6 + \varepsilon)$ -approximation for *k-center* in sliding windows

saving only  $O((k/\varepsilon) \log \Delta)$  points, obtaining a working memory *independent* of the size of the active set, where  $\Delta$  denotes the aspect ratio of the data. At the other extreme lies the *fully dynamic* model, where insertions and deletions are arbitrary and not known in advance. This is the most general and demanding setting. Chan et al. [3] gave the first fully dynamic algorithm for  $k$ -center, and Bateni et al. [4] later refined the approach into the current state of the art, maintaining a  $(2 + \varepsilon)$ -approximation with total running time  $O(n \cdot k \frac{\log \Delta}{\varepsilon} \log n)$  over a sequence of  $n$  operations.

In many real scenarios the assumption of unknown deletion times is far too pessimistic. Cached entries and DNS records carry an explicit time-to-live; reservations, licenses, and sessions come with an expiry date; a check-in in a location-based social network remains the user’s current position exactly until their next check-in; and buffered items in packet routers or manufacturing pipelines are consumed within a bounded delay after being produced. The sliding-window model does assume predictable expiration, but only in the most rigid possible form, forcing every point to survive for exactly the same span and to expire in strict arrival order.

**The dynamic setting with lifetimes.** Motivated by this gap, this thesis studies  $k$ -center in the *dynamic setting with lifetimes*, which we position as a bridge between the two classical models. As in the fully dynamic setting, arrivals and deletions are arbitrary; but, as in the sliding-window setting, expirations are known in advance. We investigate this highly unexplored region of the literature to find algorithms that beat the state-of-the-art in update time or approximation efficiency.

## Contributions

The contributions of this thesis are the following:

- We introduce and formalize the dynamic setting with lifetimes for the metric  $k$ -center problem, and we show how it relates to the sliding-window and fully dynamic settings.
- We devise a deterministic  $(2 + \varepsilon)$ -approximation algorithm for this setting. We build on the structure of the fully dynamic algorithm of Chan et al. [3], but exploiting the known deletion times, achieving an amortized update time of  $O((\log \Delta / \varepsilon) k)$ , thereby removing the multiplicative  $\log n$  factor incurred by the state-of-the-art fully dynamic algorithm of Bateni et al. [4], while using memory linear in the size

of the active set. The resulting update time is optimal up to a  $O(\log \Delta/\varepsilon)$  factor, matching the  $\Omega(|X|k)$  lower bound of [4] which holds even on offline algorithms.

- As a by-product, we show that dropping the knowledge of deletion times from this algorithm yields a new, purely fully dynamic  $(2 + \varepsilon)$ -approximation. We conjecture that it retains the same expected update time and thus would improve on the fully dynamic state of the art, and, although we leave the analysis open, we include the variant in our experimental evaluation.
- We devise a deterministic  $(6 + \varepsilon)$ -approximation algorithm with *sublinear* working memory. Extending the sliding-window construction of Cohen-Addad et al. [2] to arbitrary, known deletion times, the algorithm attains  $O((\log \Delta/\varepsilon)k)$  *worst-case* update time and working memory independent of the active-set size, whenever the update sequence is “tame”. Tameness is captured by a novel parameter,  $H$ -orderedness, that quantifies how far deletions may depart from arrival order. To the best of our knowledge, this is the first  $(6 + \varepsilon)$ -approximation for  $k$ -center that operates outside the sliding-window model.
- We complement the theory with an experimental study. We implement our algorithms alongside those of Chan et al. [3] and Bateni et al. [4], and evaluate them on synthetic data and on the real-world Gowalla check-in dataset, comparing approximation quality, update and query times, and memory usage.

## Structure of the thesis

The remainder of the thesis is organized as follows. Chapter 1 introduces metric spaces and the  $k$ -center problem, recalls the hardness of the problem and Gonzalez’ greedy algorithm, and formalizes the dynamic setting together with the fully dynamic and lifetimes models. Chapter 2 is devoted to minimizing the approximation ratio: it reviews the fully dynamic algorithms of Chan et al. [3] and Bateni et al. [4], presents our  $(2 + \varepsilon)$ -approximation with lifetimes and its analysis, and discusses the fully dynamic variant obtained by discarding deletion times. Chapter 3 turns to memory efficiency: it recalls the sliding-window algorithm of Cohen-Addad et al. [2], introduces the notion of  $H$ -orderedness, and develops and analyzes our sublinear-memory  $(6 + \varepsilon)$ -approximation. Chapter 4 reports the experimental evaluation and draws conclusions.



# Chapter 1

## Preliminaries

In this chapter we provide the basic definitions needed to understand the subsequent work. We first introduce metric spaces, which give the notion of distance our clustering relies on, and then formalize the  $k$ -center problem in terms of this distance. Finally, we describe the dynamic setting in which our algorithms operate.

### 1.1 Clustering

**Definition 1.1** (Metric Space). *Let  $\mathcal{X}$  be a set and  $d : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$  a function with the following properties:*

1.  $d(x,y) = 0 \iff x = y \quad \forall x,y \in \mathcal{X}$ ;
2.  $d(x,y) = d(y,x) \quad \forall x,y \in \mathcal{X}$  (Symmetry);
3.  $d(x,y) \leq d(x,z) + d(z,y) \quad \forall x,y,z \in \mathcal{X}$  (Triangle Inequality).

*Then the pair  $(\mathcal{X}, d)$  is called a **metric space**, and  $d$  is called the **distance**. An element  $x \in \mathcal{X}$  is called a **point**.*

Throughout, we let  $d_{\min}$  and  $d_{\max}$  denote the minimum and maximum distances between any two distinct points of  $X$ , and define the **aspect ratio** of  $X$  as  $\Delta = d_{\max}/d_{\min}$ . As is standard in most previous works, our algorithms assume that  $d_{\min}$  and  $d_{\max}$  are known in advance.

The distance gives an indication of the similarity between points: the smaller the distance between two points, the more similar they are. Finding a meaningful notion of clustering in a fully general setting is notoriously elusive, so we restrict our attention to

*center-based* clustering, in which each cluster is represented by one of its points, called a *center*, and every point is assigned to the center it is closest to. To make this precise, we first extend the distance from points to sets.

**Definition 1.2** (Point-to-set distance and radius). *Let  $X \subseteq \mathcal{X}$  be a finite set of points. For a point  $p \in X$  and a subset  $C \subseteq X$ , the distance of  $p$  from  $C$  is*

$$d(p, C) := \min_{q \in C} d(p, q).$$

*For any  $S \subseteq X$ , the **radius** of  $C$  with respect to  $S$  is*

$$r_C(S) := \max_{p \in S} d(p, C).$$

A set of centers  $C \subseteq X$  induces a partition of  $X$  into  $|C|$  clusters, obtained by assigning each point to its closest center (with ties broken arbitrarily). With this assignment,  $r_C(X)$  is the largest distance between any point and the center of its cluster, and thus measures how tight the induced clustering is: the smaller the radius, the more compact the clusters. The  $k$ -center problem asks for the choice of at most  $k$  centers that makes this radius as small as possible.

### 1.1.1 $k$ -Center Clustering

**Definition 1.3** ( $k$ -Center Clustering). *Given a finite set  $X \subseteq \mathcal{X}$  and an integer  $k \in \mathbb{N}_{>0}$ , the  **$k$ -center** problem requires to find a subset  $C \subseteq X$  with  $|C| \leq k$  that minimizes the radius  $r_C(X)$ , i.e.,*

$$C^* = \arg \min_{C \subseteq X, |C| \leq k} r_C(X).$$

*The points of  $C$  are called **centers**, and we denote the radius of the optimal solution by  $r_k^*(X) = r_{C^*}(X)$ .*

**Observation 1.** *To specify a solution it suffices to provide the set of centers  $C$ : the clusters are recovered by assigning each point to its closest center, i.e., cluster  $i$  is the set  $\{x \in X \mid i = \arg \min_j \{d(x, c_j)\}\}$ .*

We now present a simple brute-force algorithm for solving  $k$ -center clustering, which enumerates all possible sets of  $k$  centers and keeps the one of minimum radius.

---

**Algorithm 1:** K-CENTERBRUTEFORCE( $X, k$ )

---

```
1  $opt \leftarrow \infty$ 
2  $c^* \leftarrow (\mathbf{0}, \dots, \mathbf{0})$ 
3  $C^* \leftarrow (\emptyset, \dots, \emptyset)$ 
4 for  $(c_1, \dots, c_k) \in X^k$  do
5    $t \leftarrow r_{\{c_1, \dots, c_k\}}(X)$ 
6   if  $opt > t$  then
7      $opt \leftarrow t$ 
8      $c^* \leftarrow (c_1, \dots, c_k)$ 
9     for  $1 \leq i \leq k$  do
10     $C_i^* \leftarrow \{x \in X \mid i = \arg \min_j \{d(x, c_j)\}\}$ 
11 return  $(c^*, C^*)$ 
```

---

**Proposition 1.1.** *The  $k$ -center clustering problem is NP-hard. In other words, it is impossible to find a polynomial time algorithm that computes an optimal solution for an arbitrary  $X$ , unless  $P = NP$ . Moreover, it is impossible to achieve an approximation factor  $2 - \varepsilon$ , for any fixed  $\varepsilon > 0$ , unless  $P = NP$ .*

Given the NP-hardness of the problem, approximate solutions are sought. One such approximation algorithm was given by Gonzalez [1]:

---

**Algorithm 2:** GONZALEZ( $X, k$ )

---

```
1  $c \leftarrow (\mathbf{0}, \dots, \mathbf{0})$ 
2  $c_1 \leftarrow$  arbitrary  $x \in X$ 
3 for  $1 < i \leq k$  do
4    $c_i \leftarrow \arg \max_{x \in X} d(x, \{c_1, \dots, c_{i-1}\})$ 
5  $C \leftarrow (\emptyset, \dots, \emptyset)$ 
6 for  $1 \leq i \leq k$  do
7    $C_i \leftarrow \{x \in X \mid i = \arg \min_j \{d(x, c_j)\}\}$ 
8 return  $(c, C)$ 
```

---

The algorithm picks the centers iteratively, choosing each new center to be the point that is farthest from the set of centers selected so far.

**Proposition 1.2.** *Gonzalez' algorithm produces a 2-approximation of  $k$ -center in time  $O(nk)$ .*

## 1.2 Dynamic Algorithms

In some applications the pointset  $X$  is not fixed, but changes over time. We define the *dynamic setting* to model such cases. In the dynamic setting:

- $X$  is the set of all points that ever appear during the whole process;
- $X_t$  is the set of *active* points at time  $t$ ;
- $\text{INSERT}(p, t)$  denotes the insertion of a point in the active set at time  $t$ ;
- $\text{DELETE}(p, t)$  denotes the deletion of a point from the active set at time  $t$ ;
- for a given point  $p \in X$  we denote with  $t_{\text{arr}}(p)$  its insertion (*arrival*) time and with  $t_{\text{del}}(p)$  its deletion time;
- all insertion operations happen at distinct times (i.e.,  $p_1 = p_2 \iff t_{\text{arr}}(p_1) = t_{\text{arr}}(p_2) \forall p_1, p_2 \in X$ ), whereas deletions may share the same time;
- a point  $p$  is active during the time interval  $[t_{\text{arr}}(p), t_{\text{del}}(p))$  (i.e.,  $p \in X_t \iff t_{\text{arr}}(p) \leq t < t_{\text{del}}(p) \forall p \in X$ ).

The goal of a dynamic algorithm is to maintain a solution for the problem on the active set  $X_t$  without recomputing it from scratch at each step, so that the cost of an update depends on the size of the active set  $X_t$  rather than on the size of the whole pointset  $X$ . For the  $k$ -center problem in particular, the algorithm should keep, at any time  $t$ , a set of at most  $k$  centers  $C_t \subseteq X_t$  whose radius  $r_{C_t}(X_t)$  is close to the optimal one  $r_k^*(X_t)$ , while keeping the time and memory it uses small.

The way in which deletions are handled distinguishes the two main dynamic settings that we consider.

**Definition 1.4** (Fully dynamic setting). *In the **fully dynamic setting**, insertions and deletions are arbitrary and issued by an external source: at any time, a new point may be inserted into the active set, or an active point may be deleted from it. The deletion time  $t_{\text{del}}(p)$  of a point  $p$  is not known in advance, and is revealed only when  $p$  is actually deleted.*

**Definition 1.5** (Dynamic setting with lifetimes). *In the **dynamic setting with lifetimes**, whenever a point  $p$  is inserted at time  $t_{\text{arr}}(p)$ , its deletion time  $t_{\text{del}}(p) > t_{\text{arr}}(p)$  is specified as well, thereby fixing the lifetime  $[t_{\text{arr}}(p), t_{\text{del}}(p))$  of  $p$  at the moment of its arrival. We say that  $p$  expires at time  $t_{\text{del}}(p)$ , when it leaves the active set.*

The algorithms presented in this thesis operate in the setting with lifetimes, and crucially exploit the knowledge of the deletion times to obtain better guarantees than what is possible in the fully dynamic setting.



# Chapter 2

## Dynamic $k$ -Center: Minimizing Approximation

In this chapter we will look at algorithms for maintaining a  $(2 + \varepsilon)$ -approximation of the dynamic  $k$ -center. We will first look at 2 previous algorithms for the fully-dynamic case, then we will present our improvement in the lifetimes setting. The first two sections will describe respectively the work of Chan et al. [3] and of Bateni et al. [4]. The work of Bateni is the current state of the art for  $k$ -center approximation in the fully dynamic setting, achieving a  $(2 + \varepsilon)$ -approximation in time  $O(n \cdot k^{\frac{\log \Delta}{\varepsilon}} \log n)$ . We will show how the setting with lifetimes allows to remove the  $\log n$  multiplicative factor from the complexity, achieving a nearly optimal running time. The work of Chan, whilst requiring a worse runtime of  $\Omega(n \cdot k^{\frac{2 \log \Delta}{\varepsilon}})$ , is still presented as our algorithm is based on it.

All the 3 presented algorithms make a set of guesses  $\Gamma$  of the optimal radius and keep for each of them a suitable data structure. To better understand the power of guessing the radius we will state the following proposition:

**Proposition 2.1.** *Let  $\gamma \in \Gamma$  be a guess of the optimal radius. If at time  $t \exists x_1, \dots, x_{k+1} \in X_t$  s.t.  $d(x_i, x_j) > 2\gamma \forall i \neq j$ , then  $r_k^*(X_t) > \gamma$ .*

*Proof.* Suppose that  $r_k^*(X_t) \leq \gamma$ . Then  $\exists c_1, \dots, c_k \in X_t$  centers with respective clusters  $C_1, \dots, C_k$  s.t.  $d(x, c_i) \leq \gamma \forall 1 \leq i \leq k, x \in C_i$ . By the pigeonhole principle, at least 2 of the  $k + 1$   $x_i$  will be in the same cluster  $C_j$ . Assume wlog that  $x_1, x_2 \in C_1$ . Then we will have:

$$2\gamma < d(x_1, x_2) \leq d(x_1, c_1) + d(x_2, c_1) \leq \gamma + \gamma = 2\gamma$$

Where we used the triangle inequality.

We obtained that  $2\gamma < 2\gamma$ , which is absurd, meaning that  $r_k^*(X_t) \leq \gamma$  must be false.  $\square$

In all 3 algorithms, for each guess  $\gamma \in \Gamma$ , we will at each time find either  $k$  centers with respective clusters at distance  $\leq 2\gamma$  or  $k + 1$  points at pairwise distance  $> 2\gamma$ . We set:

$$\Gamma = \{(1 + \beta)^i : \lfloor \log_{1+\beta} d_{\min} \rfloor \leq i \leq \lceil \log_{1+\beta} d_{\max} \rceil\}. \quad (2.1)$$

**Proposition 2.2.** *Let  $\varepsilon > 0$ . Suppose that  $\Gamma$  is as defined before with  $\beta = \varepsilon/2$  and that for each  $\gamma$  we either have a clustering with radius  $2\gamma$  or have  $k + 1$  points at distance  $> 2\gamma$ . Let  $\gamma' \in \Gamma$  be the smallest  $\gamma \in \Gamma$  s.t. we have a clustering of radius  $2\gamma$ .  $2\gamma'$  is a  $(2 + \varepsilon)$ -approximation of the optimal radius.*

*Proof.* Given that for  $\gamma'$  we know  $k$  centers at distance  $\leq 2\gamma'$  from all the other points we have that  $r_k^*(X_t) \leq 2\gamma'$ . It remains to prove that  $r_k^*(X_t) \geq 2\gamma'/(2 + \varepsilon)$ .

Let  $i$  be s.t.  $(1 + \beta)^i = \gamma'$ . If  $i = \lfloor \log_{1+\beta} d_{\min} \rfloor$  then given that  $d_{\min}$  is the minimal pairwise distance we have  $r_k^*(X_t) \geq d_{\min}$  (unless  $k = |X_t|$ ), therefore:  $2\gamma'/(2 + \varepsilon) < \gamma' \leq d_{\min} \leq r_k^*(X_t)$ .

If  $i > \lfloor \log_{1+\beta} d_{\min} \rfloor$  by the fact that  $\gamma'$  is minimal we know that a solution with radius  $(1 + \beta)^{i-1}$  is impossible, therefore:

$$r_k^*(X_t) > (1 + \beta)^{i-1} = \frac{(1 + \beta)^i}{1 + \beta} = \frac{\gamma'}{1 + \varepsilon/2} = \frac{2\gamma'}{2 + \varepsilon}.$$

$\square$

## 2.1 A $O(k^2)$ Fully Dynamic Algorithm

In this section we present the work of Chan et al. [3], who provided the first fully dynamic algorithm for the  $k$ -center problem. For ease of readability we will omit the subscript  $t$  indicating the time at which each object exists every time that it is not strictly necessary. Just remember that, being this a dynamic algorithm, all the presented data structure will change over time. Moreover, both proofs of the correctness and running time are omitted. The correctness is later proved in the analysis with lifetimes, the proofs for the running time can be found in the original paper.

As discussed before, the algorithm creates the set of guesses  $\Gamma$  and maintains a data structure  $\mathcal{L}_\gamma$  for each of them.  $\mathcal{L}^\gamma$  is composed by:

- A list  $C^\gamma = (c_1^\gamma, \dots, c_{\ell^\gamma}^\gamma)$  of  $\ell^\gamma < k$  centers s.t. for all  $x, y \in C^\gamma$ ,  $d(x, y) > 2\gamma$ .
- A list  $D^\gamma = (D_1^\gamma, \dots, D_{\ell^\gamma}^\gamma)$  of disjoint clusters s.t.  $\forall 1 \leq i \leq \ell^\gamma, x \in D_i^\gamma, d(x, c_i^\gamma) \leq 2\gamma$ .
- A set  $U^\gamma = X \setminus \cup_{1 \leq i \leq \ell^\gamma} D_i^\gamma$  of unclustered points s.t.  $\forall 1 \leq i \leq \ell^\gamma, x \in U^\gamma, d(c_i^\gamma, x) \geq 2\gamma$ . Moreover, we require that  $U^\gamma \neq \emptyset \implies \ell^\gamma = k$ .

Having defined the data structure this way, it is quite simple to see that for each  $\gamma$  if  $U^\gamma \neq \emptyset$  then there are at least  $k + 1$  points at distance  $> 2\gamma$ , otherwise a clustering with radius  $2\gamma$  is given by  $(C^\gamma, D^\gamma)$ .

The insertion procedure is quite simple:

---

**Algorithm 3:** INSERT( $p, \gamma$ )

---

```

1 for  $i = 1, \dots, \ell^\gamma$  do
2   if  $d(p, c_i^\gamma) \leq 2\gamma$  then
3     Insert  $p$  in  $D_i^\gamma$ 
4     return
5 if  $\ell^\gamma < k$  then
6   Insert  $p$  in  $C^\gamma$  and in  $D_{\ell^\gamma+1}^\gamma$   $\ell^\gamma \leftarrow \ell^\gamma + 1$ 
7 else Insert  $p$  in  $U^\gamma$ ;

```

---

It inserts  $p$  in the cluster of  $D^\gamma$  of lowest index whose center is at distance at most  $2\gamma$  from  $p$ . If no such cluster exists, if  $D^\gamma$  comprises  $\ell^\gamma < k$  clusters, then the procedure creates a new cluster by making  $p$  its center, otherwise, it inserts  $p$  in  $U^\gamma$ .

Deleting a point  $p$  is a bit trickier: if  $x \notin C^\gamma$  it can be simply removed from its cluster without changing anything else. However if  $x \in C^\gamma$  we need to do additional work to maintain the properties of our data structure. Let's look at the deletion procedure:

---

**Algorithm 4:** DELETE( $p, \gamma$ )

---

```
1 if  $p \notin C^\gamma$  then
2   | Remove  $p$  from its respective cluster in  $D^\gamma$  or from  $U^\gamma$ 
3   | return
4 Let  $1 \leq i \leq \ell^\gamma$  s.t.  $p = c_i^\gamma$ 
5  $\hat{X} \leftarrow (\cup_{i \leq j \leq \ell^\gamma} D_j^\gamma) \cup U^\gamma$ 
6  $(\hat{C}, \hat{D}, \hat{U}) \leftarrow \text{RANDRECLUST}(\hat{X}, \gamma, k - i + 1)$ 
7  $C^\gamma \leftarrow (c_1^\gamma, \dots, c_{i-1}^\gamma) \oplus \hat{C}$ 
8  $D^\gamma \leftarrow (D_1^\gamma, \dots, D_{i-1}^\gamma) \oplus \hat{D}$ 
9  $\ell^\gamma \leftarrow i + |\hat{C}| - 1$ 
10  $U^\gamma \leftarrow \hat{U}$ 
```

---

Where the  $\oplus$  operator denotes the concatenation of 2 lists. As we can see if a point is not a center it gets simply deleted, otherwise all of the points that are in clusters with index  $\geq$  than the index of the deleted center get reclustered. The reclustering procedure (Algorithm 5) takes this set in input, alongside a value  $\kappa$ , and returns up to  $\kappa$  clusters with corresponding centers picked at random from the set of points. The points left unclustered by this procedure, if any, are then put in  $U$ .

---

**Algorithm 5:** RANDRECLUST( $X, \gamma, \kappa$ )

---

```
1  $U \leftarrow X$ 
2  $\ell \leftarrow 0$ 
3 while  $U \neq \emptyset$  and  $\ell < \kappa$  do
4   |  $\ell \leftarrow \ell + 1$ 
5   | Pick  $c_\ell$  from  $U$  uniformly at random
6   |  $D_\ell \leftarrow \{x \in U \mid d(x, c_\ell) \leq 2\gamma\}$ 
7   |  $U \leftarrow U \setminus D_\ell$ 
8 return  $((c_1, \dots, c_\ell), (D_1, \dots, D_\ell), U)$ 
```

---

**Theorem 2.1.** *The algorithm described above obtains a  $(2 + \varepsilon)$ -approximation of the  $k$ -center Problem and runs in  $O(n \cdot k^2 \frac{\log \Delta}{\varepsilon})$  amortized time in expectation.*

## 2.2 The State-of-the-Art of Fully Dynamic $k$ -Center

In this section we present the work of Bateni et al. [4], which is the current state of the art for the fully dynamic  $k$ -center problem, lowering the amortized update time from the  $\Omega(n \cdot k^2 \frac{\log \Delta}{\varepsilon})$  of [3] to  $O(n \cdot k^2 \frac{\log \Delta}{\varepsilon} \log n)$  while retaining the same  $(2 + \varepsilon)$ -approximation.

As in the previous section, for ease of readability we will omit the subscript  $t$  indicating the time at which each object exists whenever it is not strictly necessary, keeping in mind that, this being a dynamic algorithm, all the presented data structures change over time. We also omit both correctness and running time proofs: the correctness follows from Propositions 2.1 and 2.2, while the running time proofs can be found in the original paper.

The key idea is to abandon the explicit clustering maintained by [3] in favour of the classical reduction from  $k$ -center to Maximal Independent Set (MIS) due to Hochbaum and Shmoys [5]. For a guess  $\gamma \in \Gamma$  we define the *threshold graph*  $G^\gamma = (X, E^\gamma)$  on the active points, where

$$(x, y) \in E^\gamma \iff d(x, y) \leq 2\gamma. \quad (2.2)$$

An independent set of  $G^\gamma$  is then exactly a set of points at pairwise distance  $> 2\gamma$ , while a *maximal* independent set  $I^\gamma$  enjoys the additional property that every point of  $X$  is within distance  $2\gamma$  of some vertex of  $I^\gamma$ . Therefore, if  $|I^\gamma| \leq k$ , taking  $I^\gamma$  as centers yields a clustering of radius  $2\gamma$ ; if instead  $|I^\gamma| \geq k + 1$ , then  $I^\gamma$  contains  $k + 1$  points at pairwise distance  $> 2\gamma$  and, by Proposition 2.1,  $r_k^*(X) > \gamma$ . This is precisely the dichotomy required by Proposition 2.2, so maintaining a maximal independent set of  $G^\gamma$  for each  $\gamma \in \Gamma$  and returning the centers of the smallest feasible guess gives a  $(2 + \varepsilon)$ -approximation.

Notice that for the purpose of the dichotomy above it is never necessary to compute the whole MIS: as soon as an independent set of size  $k + 1$  is found we can already certify  $r_k^*(X) > \gamma$ . This motivates the *k-Bounded MIS* problem: maintain either a maximal independent set of size at most  $k$ , or an arbitrary independent set of size  $k + 1$ . The advantage is that, while the insertion of a point can create up to  $\Omega(n)$  edges in  $G^\gamma$ , only the  $O(k)$  potential edges between the new point and the current independent set need to be inspected to decide whether the latter is still maximal.

**Maintaining the LFMIS** The structure that the algorithm keeps is the *Lexicographically First Maximal Independent Set* (LFMIS) of  $G^\gamma$  with respect to a random ranking  $\pi : X \rightarrow [0,1]$ . The LFMIS is obtained greedily by repeatedly adding the vertex of smallest rank, removing it together with all its neighbours, and iterating until no vertex is left. Each point draws its rank  $\pi(x)$  once, uniformly and independently, upon insertion. Maintaining the LFMIS under a random ranking has two convenient properties: it is *history independent*, in that once  $\pi$  is fixed the current LFMIS depends only on the current

graph and not on the order of the updates that produced it; and the insertion of a vertex  $v$  is unlikely to cause many changes, since this requires  $\pi(v)$  to be small.

In line with the  $k$ -Bounded MIS problem, the algorithm only maintains the *top- $k$  LFMIS*  $L^\gamma$ , defined as the first  $\min\{k+1, |\text{LFMIS}^\gamma|\}$  vertices of the LFMIS in rank order. The difficulty is to keep track of the “excess” vertices that belong to the LFMIS but not to  $L^\gamma$ , so that they can quickly re-enter  $L^\gamma$  once a vertex of smaller rank leaves it. To this end, for each vertex  $u$  not in the LFMIS we define its *eliminator*  $\text{elim}^\gamma(u)$  as the neighbour of smallest rank in  $N^\gamma(u) \cap \text{LFMIS}^\gamma$ , where  $N^\gamma(u) = \{x \in X : x \neq u, d(u, x) \leq 2\gamma\}$  denotes the neighbourhood of  $u$  in  $G^\gamma$ . The eliminator certifies that  $u$  cannot enter the independent set. The algorithm maintains a *leader* mapping  $\text{lead}^\gamma(u) = \text{elim}^\gamma(u)$ , so that whenever a vertex  $v$  leaves the LFMIS only the vertices in its *follower set*  $F_v^\gamma = \{u : \text{lead}^\gamma(u) = v\}$  need to be examined.

Concretely, the data structure kept for each  $\gamma \in \Gamma$  is composed by:

- $L^\gamma$ , an ordered list of at most  $k+1$  vertices, sorted by rank  $\pi$ , holding the top- $k$  LFMIS;
- $Q^\gamma$ , a priority queue, ordered by rank  $\pi$ , of the unclustered vertices, i.e. the excess vertices waiting to be (re)inserted;
- a follower set  $F_v^\gamma$  for each leader  $v$ , together with the leader mapping  $\text{lead}^\gamma$  pointing each follower to its eliminator.

Accordingly, at any time each vertex is either an *active leader* (in  $L^\gamma$ ), an *inactive leader* (a vertex that still owns a non-empty follower set but has been pushed into  $Q^\gamma$ ), a *follower* (a vertex  $u$  with  $\text{lead}^\gamma(u) \in L^\gamma$ ), or *unclustered* (in  $Q^\gamma$  with no leader).

**Update procedure** Every insertion or deletion is handled, for each  $\gamma \in \Gamma$ , by the procedure  $\text{PROCESSUPDATE}(v, \sigma)$  (Algorithm 6), where  $\sigma \in \{+, -\}$  marks the update as an insertion or a deletion. After dispatching the update to  $\text{INSERT}$  or  $\text{DELETE}$ , a final loop drains the queue  $Q^\gamma$ , reinserting its vertices in increasing rank order until either  $Q^\gamma$  is empty or no queued vertex has rank small enough to enter  $L^\gamma$ ; this restores the invariant  $L^\gamma = \text{LFMIS}_{k+1}^\gamma$ .

---

**Algorithm 6:** PROCESSUPDATE( $v, \sigma, \gamma$ )

---

```
1 if  $\sigma = +$  then
2   | Draw  $\pi(v)$  uniformly at random in  $[0,1]$ 
3   |  $\text{lead}^\gamma(v) \leftarrow \perp$ 
4   | INSERT( $v, \gamma$ )
5 else DELETE( $v, \gamma$ );
6 while  $Q^\gamma \neq \emptyset$  and  $(|L^\gamma| \leq k$  or  $\min_{w \in Q^\gamma} \pi(w) < \max_{w \in L^\gamma} \pi(w))$  do
7   |  $u \leftarrow \arg \min_{w \in Q^\gamma} \pi(w)$ 
8   | Remove  $u$  from  $Q^\gamma$ 
9   | INSERT( $u, \gamma$ )
```

---

The insertion of a vertex  $v$  (Algorithm 7) first discards the trivial case in which  $L^\gamma$  is already full and  $v$  has the largest rank, so that  $v$  cannot belong to the top- $k$  LFMIS and is simply queued. Otherwise it computes the set  $S = L^\gamma \cap N^\gamma(v)$  of the current leaders adjacent to  $v$ , which costs  $O(k)$  distance queries since  $|L^\gamma| \leq k + 1$ . If  $S = \emptyset$  then  $v$  has no neighbour in  $L^\gamma$  and can enter it, possibly evicting the leader of largest rank into  $Q^\gamma$  should  $L^\gamma$  exceed  $k + 1$  vertices. If instead  $v$  has a neighbour of smaller rank, the smallest-rank such neighbour  $u^*$  becomes the leader of  $v$ . Finally, if  $v$  has smaller rank than all of its neighbours in  $S$ , it enters the LFMIS and displaces them: the displaced leaders become followers of  $v$ , while their own followers, having lost their eliminator, are released into  $Q^\gamma$ .

---

**Algorithm 7:** INSERT( $v, \gamma$ )

---

```
1 if  $|L^\gamma| = k + 1$  and  $\pi(v) > \max_{u \in L^\gamma} \pi(u)$  then
2   |   Insert  $v$  in  $Q^\gamma$ 
3   |   return
4  $S \leftarrow L^\gamma \cap N^\gamma(v)$ 
5 if  $S = \emptyset$  then
6   |   Insert  $v$  in  $L^\gamma$ 
7   |   if  $|L^\gamma| = k + 2$  then
8     |   |    $u \leftarrow \arg \max_{u' \in L^\gamma} \pi(u')$ 
9     |   |   Remove  $u$  from  $L^\gamma$  and insert it in  $Q^\gamma$ 
10 else
11   |    $u^* \leftarrow \arg \min_{u' \in S} \pi(u')$ 
12   |   if  $\pi(u^*) < \pi(v)$  then
13     |   |   if  $v$  is a leader then
14       |   |   |   for  $w \in F_v^\gamma$  do
15         |   |   |   |   Insert  $w$  in  $Q^\gamma$ ;  $\text{lead}^\gamma(w) \leftarrow \perp$ 
16         |   |   |   |   Delete  $F_v^\gamma$ 
17         |   |   |   Add  $v$  to  $F_{u^*}^\gamma$ ;  $\text{lead}^\gamma(v) \leftarrow u^*$ 
18     |   |   else
19       |   |   |   for  $w \in \cup_{u \in S} F_u^\gamma$  do
20         |   |   |   |   Insert  $w$  in  $Q^\gamma$ ;  $\text{lead}^\gamma(w) \leftarrow \perp$ 
21         |   |   |   for  $u \in S$  do
22         |   |   |   |    $\text{lead}^\gamma(u) \leftarrow v$ ; remove  $u$  from  $L^\gamma$ ; delete  $F_u^\gamma$ 
23         |   |   |   Insert  $v$  in  $L^\gamma$  with  $F_v^\gamma \leftarrow S$ 
```

---

The deletion of a vertex  $v$  (Algorithm 8) is cheap unless  $v$  is a leader. If  $v$  is a follower it is simply removed from its leader's follower set. If  $v$  sits in  $Q^\gamma$  as an inactive leader, its followers are released into  $Q^\gamma$  before  $v$  is removed. Finally, if  $v$  is an active leader, all of its followers are released into  $Q^\gamma$  and  $v$  is removed from  $L^\gamma$ ; the queue draining at the end of PROCESSUPDATE then restores the invariant.

---

**Algorithm 8:** DELETE( $v, \gamma$ )

---

```
1 if  $v$  is a follower then
2   | Remove  $v$  from  $F_{\text{lead}^\gamma(v)}^\gamma$  and from  $X$ 
3   | return
4 else if  $v \in Q^\gamma$  then
5   | if  $v$  is an inactive leader then
6   |   | for  $w \in F_v^\gamma$  do
7   |   |   | Insert  $w$  in  $Q^\gamma$ ;  $\text{lead}^\gamma(w) \leftarrow \perp$ 
8   |   |   | Delete  $F_v^\gamma$ 
9   |   | Remove  $v$  from  $Q^\gamma$  and from  $X$ 
10  | return
11 else
12  | for  $w \in F_v^\gamma$  do
13  |   | Insert  $w$  in  $Q^\gamma$ ;  $\text{lead}^\gamma(w) \leftarrow \perp$ 
14  |   | Delete  $F_v^\gamma$ 
15  |   | Remove  $v$  from  $L^\gamma$  and from  $X$ 
```

---

A query is answered exactly as prescribed by Proposition 2.2: we return the centers of the smallest guess  $\gamma^* \in \Gamma$  for which  $|L^{\gamma^*}| \leq k$ , each non-center point  $x$  being assigned to the cluster of its leader  $\text{lead}^{\gamma^*}(x)$ .

The improvement over [3] stems from the random ranking: the cost of an update is, up to an  $O(k)$  factor, the number of vertices whose leader changes during that update, and the history independence of the LFMS guarantees that this number is small in expectation.

**Theorem 2.2.** *The algorithm described above obtains a  $(2 + \varepsilon)$ -approximation of the  $k$ -center Problem and runs in  $O(n \cdot k \frac{\log \Delta}{\varepsilon} \log n)$  amortized time in expectation.*

## 2.3 With lifetimes

In this section we present our  $(2 + \varepsilon)$ -approximation algorithm for dynamic  $k$ -center when viewed in the lifetimes setting, which attains a  $(2 + \varepsilon)$  approximation ratio in  $O((\log \Delta/\varepsilon) \cdot k + \max_t \log |X_t|)$  amortized update time. We recall that in the dynamic with lifetimes setting the algorithm knows  $t_{\text{del}}(x) \forall x \in X$  since the point is inserted.

The algorithm is based on the one by Chan et al.[3], presented in section 2.1. In their work, they used the randomness of the reclustering procedure to prove that the points to

be deleted are centers with low enough probability. However, knowing the deletion time of each point may allow us to choose better new centers at reclustering: we can indeed choose the ones with higher deletion times. An effortless reworking of the algorithm, in which the only change is that at line 5 of RANDRECLUST we pick the point with highest deletion time in  $U$ , does not achieve  $\tilde{O}(n \cdot k)$  amortized update time. We will show this with a counterexample.

**Proposition 2.3.** *There exists a sequence of updates on which the modified algorithm runs in  $\Theta(n \cdot k^2)$  amortized time.*

*Proof.* We will give a counterexample that doesn't even take into account the fact that deletion times are known and works with any policy of center selection during reclustering. Let  $\gamma \in \mathbb{R}^+$  and  $k \in \mathbb{N}^*$ , we take  $n = 2k$  points  $x_1, \dots, x_n$  s.t.:

- $d(x_{2i-1}, x_{2i}) \leq 2\gamma \forall 1 \leq i \leq k.$
- $d(x_i, x_j) > 2\gamma$  for all other pairs of different indexes.
- $t_{\text{arr}}(x_i) = i \forall 1 \leq i \leq n.$
- $t_{\text{del}}(x_{2i-1}) = n + k - i + 1 \forall 1 \leq i \leq k.$
- $t_{\text{del}}(x_{2i}) = n + 2k - i + 1 \forall 1 \leq i \leq k.$

Meaning that the  $n$  points are grouped into  $k$  clusters of 2 points each. By how the insertion algorithm 3 works, points with index  $2i - 1$  will become the  $i$ -th center and points with index  $2i$  will enter the  $i$ -th cluster. Given that the deletions are made in the opposite order of insertion, we will first delete the  $k$ -th center, then the  $k - 1$ -th and so on. The reclustering algorithm 5 only affects clusters with index  $\geq$  than the one of the center being deleted, so we are sure that all original centers will remain centers until their deletion. When deleting the  $i$ -th center, the only points remaining in clusters with index  $j \geq i$  are  $x_{2j}$  and they all need to be put on different clusters. So the reclustering of points when deleting the  $i$ -th center takes  $\Theta((k - i)^2)$  time. The total reclustering cost is then:

$$\sum_{i=1}^k \Theta((k - i)^2) = \Theta(k^3) = \Theta(n \cdot k^2)$$

Given that  $n = 2k$ . □

To fully exploit the deletion time knowledge we will need to modify the deletion algorithm, as well as being more conservative of the times that we call the reclustering procedure.

The data structure maintained for each  $\gamma \in \Gamma$  is identical of the one used in [3], the real change is in how it is modified by deletions. To recall we save the sets:

- $C^\gamma = (c_1^\gamma, \dots, c_{\ell^\gamma}^\gamma)$ , a list of  $\ell^\gamma \leq k$  centers, such that for each  $1 \leq i \neq j \leq \ell^\gamma$ , we have  $d(c_i^\gamma, c_j^\gamma) > 2\gamma$ ;
- $D^\gamma = (D_1^\gamma, \dots, D_{\ell^\gamma}^\gamma)$ , a list of  $\ell^\gamma \leq k$  disjoint sets of points, such that for each  $1 \leq i \leq \ell^\gamma$  and  $x \in D_i^\gamma$ , we have  $d(x, c_i^\gamma) \leq 2\gamma$ ;
- $U^\gamma$ , a set of unclustered points such that, for each  $x \in U^\gamma$  and  $1 \leq i \leq \ell^\gamma$ , we have  $d(x, c_i^\gamma) > 2\gamma$ . Moreover  $U^\gamma \neq \emptyset$  implies  $\ell^\gamma = k$ .

Given that the deletion times are given during insertion the streaming model is a bit different than the previously described ones. Indeed, whilst in the fully dynamic model both insertion and deletions are issued by the model itself, in the setting with lifetimes the deletions need to be handled automatically by the algorithm. To account for that, we save the inserted points in a priority queue, ordered by smallest deletion time, which is the reason we get an additive  $O(\max_t \log |X_t|)$  in the final complexity.

**Update procedure** Procedure  $\text{UPDATE}_{2+\varepsilon}(p, t)$  (see Algorithm 9 for the pseudocode) is called when a new point  $p$  arrives (in this case  $t = t_{\text{arr}}(p)$ ), or at any time  $t$  when expired points, if any, need to be deleted from the data structure (in this case,  $p = \text{NULL}$ ). First, for each  $q \in Q$  with  $t_{\text{del}}(q) \leq t$ ,  $q$  is removed from  $Q$  and procedure  $\text{DELETE}(q, \gamma)$  is invoked for each  $\gamma \in \Gamma$ . Then, if  $p \neq \text{NULL}$ , procedure  $\text{INSERT}(p, \gamma)$  is invoked for each  $\gamma \in \Gamma$ , and  $p$  is added to the priority queue  $Q$ .

$\text{INSERT}(p, \gamma)$  (see Algorithm 10 for the pseudocode) inserts  $p$  in the cluster of  $D^\gamma$  of lowest index whose center is at distance at most  $2\gamma$  from  $p$ . If no such cluster exists, then if  $D^\gamma$  comprises  $\ell^\gamma < k$  clusters, then the procedure creates a new cluster by making  $p$  its center, otherwise, it inserts  $p$  in  $U^\gamma$ . At the end of  $\text{INSERT}(p, \gamma)$  a reclustering procedure, described later, is called.

$\text{DELETE}(p, \gamma)$ , where  $p$  is an expired point, works as follows (see Algorithm 11). If  $p$  is not a center, it is simply removed from its cluster in  $D^\gamma$ . If  $p$  is a center  $c_i^\gamma$ , the procedure deletes  $c_i^\gamma$  and tries to recluster all points of  $D_i^\gamma$  in clusters with index  $> i$ . Intuitively, each of these points is inserted in the cluster where it would have been inserted if  $c_i^\gamma$  had

---

**Algorithm 9:** UPDATE<sub>2+ε</sub>( $p, t$ )

---

```
1 while  $\min_{q \in Q} t_{\text{del}}(q) \leq t$  do
2   | Remove  $q$  from  $Q$ 
3   | for  $\gamma \in \Gamma$  do
4   |   | DELETE( $q, \gamma$ )
5 if  $p \neq \text{NULL}$  then
6   | for  $\gamma \in \Gamma$  do
7   |   | INSERT( $p, \gamma$ )
8   |   Insert  $p$  in  $Q$ 
```

---

---

**Algorithm 10:** INSERT( $p, \gamma$ )

---

```
1 for  $i = 1, \dots, \ell^\gamma$  do
2   | if  $d(p, c_i^\gamma) \leq 2\gamma$  then
3   |   | Insert  $p$  in  $D_i^\gamma$ , and goto
4   |   | line 8
5 if  $\ell^\gamma < k$  then
6   |   Insert  $p$  in  $C^\gamma$  and in  $D_{\ell^\gamma+1}^\gamma$ 
7   |    $\ell^\gamma \leftarrow \ell^\gamma + 1$ 
7 else Insert  $p$  in  $U^\gamma$ ;
8 RECLUSTERING( $\gamma$ )
```

---

not existed. Then, all cluster indexes greater than  $i$  are decremented by 1, making them contiguous again, and, if  $U^\gamma \neq \emptyset$ , the point of  $U^\gamma$  with the latest deletion time becomes the center of a new cluster. Finally, the reclustering procedure is called.

Note that the expensive part of the deletion procedure, namely, re-assigning points and creating a new cluster, only happens if a cluster center is deleted. In order to limit these costly operations, we aim at selecting centers with largest deletion times. In particular, this is enforced when a new center from  $U^\gamma$  must be selected (Lines 17÷23 of DELETE( $p, \gamma$ )), and when the reclustering is performed at the end of INSERT and DELETE, which is described next.

**Definition 2.1.** For a given  $\gamma \in \Gamma$ , we define as **persistent** those clustered points which expire after their respective clusters' centers, and **vanishing** those which either expire earlier than their clusters' centers or are unclustered.

Note that only persistent points can change cluster because of the deletion of their center, as vanishing points expire before the deletion of their center. We partition each cluster  $D_i^\gamma$  into persistent and vanishing points as  $P_i^\gamma := \{x \in D_i^\gamma : t_{\text{del}}(c_i^\gamma) < t_{\text{del}}(x)\}$  and  $V_i^\gamma := \{x \in D_i^\gamma : t_{\text{del}}(c_i^\gamma) \geq t_{\text{del}}(x)\}$ . In practice, we maintain two counters per cluster for  $|P_i^\gamma|$  and  $|V_i^\gamma|$ , which can be updated in  $O(1)$  time for each insertion and deletion.

Procedure RECLUSTERING( $\gamma$ ) (see Algorithm 12), first searches for the smallest index  $i$  such that

$$\sum_{j=i}^{\ell^\gamma} |P_j^\gamma| > |U^\gamma| + \sum_{j=i}^{\ell^\gamma} |V_j^\gamma|. \quad (2.3)$$

If such  $i$  exists, then it reclusters the points of  $U^\gamma \cup (\cup_{j=i}^{\ell^\gamma} D_j^\gamma)$  into  $\ell^\gamma - i + 1$  clusters, selecting, for each such new cluster, a center with latest deletion time. Note that, after reclustering, some points might remain unclustered, but the newly created clusters comprise only vanishing points.

---

**Algorithm 11: DELETE( $p, \gamma$ )**

---

```
1 if  $p \notin C^\gamma$  then
2   | Remove  $p$  from  $U^\gamma$  or  $D^\gamma$ 
3   | goto line 24
4 Let  $i$  be such that  $p = c_i^\gamma$ 
5 for  $x \in D_i^\gamma$  with  $x \neq p$  do
6   | for  $j = i + 1, \dots, \ell^\gamma$  do
7     | | if  $d(x, c_j^\gamma) \leq 2\gamma$  then
8       | | | Insert  $x$  in  $D_j^\gamma$ , and
9       | | | break
10    | | if  $x$  was not inserted then
11      | | | if  $\ell^\gamma < k$  then
12        | | | | Insert  $x$  in  $C^\gamma$  and in
13        | | | |  $D_{\ell^\gamma+1}^\gamma$ ;
14        | | | |  $\ell^\gamma \leftarrow \ell^\gamma + 1$ 
15      | | | else Insert  $x$  in  $U^\gamma$ ;
16 Remove  $c_i^\gamma$  and  $D_i^\gamma$ ;
17  $\ell^\gamma \leftarrow \ell^\gamma - 1$ ;
18 Rename indexes of  $C^\gamma, D^\gamma$  to be
19   contiguous
20 if  $U^\gamma \neq \emptyset$  then
21   |  $\ell^\gamma \leftarrow \ell^\gamma + 1$ ;
22   | Pick  $u \in U^\gamma$  s.t.  $t_{\text{del}}(u)$  is
23   |   maximal;
24   | Remove  $u$  from  $U^\gamma$ , insert in
25   |    $C^\gamma, D^\gamma$  as  $c_{\ell^\gamma}^\gamma$ 
26   | for  $x \in U^\gamma$  do
27     | | if  $d(x, u) \leq 2\gamma$  then
28       | | | Remove  $x$  from  $U^\gamma$ ,
29       | | |   insert in  $D_{\ell^\gamma}^\gamma$ 
30 RECLUSTERING( $\gamma$ )
```

---

---

**Algorithm 12: RECLUSTERING( $\gamma$ )**

---

```
1  $i \leftarrow \min$  index  $h$  s.t.
2    $\sum_{j=h}^{\ell^\gamma} |P_j^\gamma| > |U^\gamma| + \sum_{j=h}^{\ell^\gamma} |V_j^\gamma|$ 
3 if no such  $i$  exists then
4   | return
5  $U' \leftarrow U^\gamma \cup (\cup_{j=i}^{\ell^\gamma} D_j^\gamma)$ 
6 Remove centers  $c_j^\gamma$  and clusters  $D_j^\gamma$ 
7    $\forall j \geq i$ ;
8  $\ell^\gamma \leftarrow i - 1$ ;
9 for  $j = i, \dots, k$  do
10  | if  $U' = \emptyset$  then
11    | | break
12    | |  $\ell^\gamma \leftarrow j$ ;
13    | | Pick  $u \in U'$  s.t.  $t_{\text{del}}(u)$  is maximal;
14    | | Remove  $u$  from  $U'$ , insert in  $C^\gamma, D^\gamma$ 
15    | |   as  $c_j^\gamma$ ;
16    | | for  $x \in U'$  do
17      | | | if  $d(x, u) \leq 2\gamma$  then
18        | | | | Remove  $x$  from  $U'$ , insert in
19        | | | |  $D_j^\gamma$ 
20  $U^\gamma \leftarrow U'$ 
```

---

---

**Algorithm 13: QUERY $_{2+\varepsilon}(t)$** 

---

```
1  $\gamma^* \leftarrow \min\{\gamma \in \Gamma : U^\gamma = \emptyset\}$ ;
2 return  $C^{\gamma^*}$ 
```

---

The following remark shows that the reclustering is crucial to obtain sublinear amortized update time.

**Remark.** *There exists a sequence of  $O(n)$  calls to UPDATE $_{2+\varepsilon}$  which, without the use of RECLUSTERING, would perform a total of  $\Theta(n^2)$  operations.*

*Proof.* We describe an adversarial point set for a given guess  $\gamma \in \Gamma$ . At time  $t = 1$ , we insert a point  $p_1$  with lifetime such that  $t_{\text{del}}(p_1) = t_{\text{arr}}(p_1) + n + 1$ .

At time  $t = 2, \dots, n$ , we insert  $n - 1$  points  $p_2, \dots, p_n$ , each with lifetime such that  $t_{\text{del}}(p) = t_{\text{arr}}(p) + 2n - 1$ , and such that each  $p_i$  with  $i > 1$  is at distance  $3/2\gamma$  from  $p_1$  (and, e.g., at distances  $\gamma/10$  between each other). Clearly, all these points are assigned to cluster  $D_1$ , with the first point  $p_1$  being the center. Then, at time  $t = n + 1, \dots, 2n$ , we

insert  $n$  additional points  $p_{n+1}, \dots, p_{2n}$ , each with lifetime such that  $t_{\text{del}}(p_i) = t_{\text{arr}}(p_i) + 2$  and such that each such  $p_i$  is at distance  $5/2\gamma$  from  $p_1$ , at distance  $5/2\gamma$  from the other points  $p_j$  with  $n < j \leq 2n$ , and at distance  $3/2\gamma$  from points  $p_2, \dots, p_n$ .

The first  $n$  calls of  $\text{UPDATE}_{2+\varepsilon}$  involve no deletions, so they take  $O(k)$  time each. The call at time  $n + 1$  inserts  $p_{n+1}$  as a new center, since it is at distance  $> 2\gamma$  from  $p_1$ , in  $O(k)$  time. However, the call at time  $n + 2$  first involves the deletion of  $p_1$  and moving  $p_2, \dots, p_n$  to the cluster with center  $p_{n+1}$ , and then inserts  $p_{n+2}$  as a new center. This takes  $\Theta(n)$  time. Each call at time  $t = n + 3, \dots, 2n$  involves the deletion of point  $p_{t-2}$ , moving  $p_2, \dots, p_n$  to the cluster with center  $p_{t-1}$  and the insertion of  $p_t$  as a new center, which takes  $\Theta(n)$ . Therefore, the algorithm performs a total of  $\Theta(n^2)$  operations.  $\square$

**Query procedure** At any time  $t$ , an approximate solution to  $k$ -center for the active set  $X_t$  can be computed calling  $\text{QUERY}_{2+\varepsilon}(t)$  (see Algorithm 13 for the pseudocode). We assume that when  $\text{QUERY}_{2+\varepsilon}(t)$  is invoked, the data structure contains no expired points. This can be ensured by checking that the minimum deletion time of a point in  $Q$  is  $> t$ , and, if this is not the case, invoking  $\text{UPDATE}_{2+\varepsilon}(\text{NULL}, t)$  before  $\text{QUERY}_{2+\varepsilon}(t)$ . The procedure simply returns the set  $C^{\gamma^*}$ , with  $\gamma^* := \min_{\gamma \in \Gamma} \{\gamma : U^\gamma = \emptyset\}$ , as a set of centers.

The following subsections analyze the approximation ratio, and the time and memory requirements of the above algorithm. For convenience, we define a *time step* as one call of INSERT or DELETE. Since arrival times are distinct and  $Q$  breaks ties for deletion times based on arrival times, we can assume, without loss of generality, that each call corresponds to a unique time step. For any time step  $t \geq 0$ , we use  $X_t$  to denote the active set, and  $C^{\gamma,t}, D^{\gamma,t}, U^{\gamma,t}, Q^t$  to denote the various components of the data structure, *after* the  $t$ <sub>th</sub> call of INSERT or DELETE. Time step  $t = 0$  refers to the initial phase of the algorithm, when nothing is yet inserted. We may drop  $t$  when clear from the context.

### 2.3.1 Approximation

First, we prove that the algorithm maintains the following invariant. This can be then used to prove the approximation guarantee for  $\text{QUERY}_{2+\varepsilon}$ .

**Lemma 2.1.** *For any time step  $t$ , the following properties hold for each  $\gamma \in \Gamma$ , and  $x \in X_t$ :*

1. *if  $d(x, C^{\gamma,t}) > 2\gamma$ , then  $x \in U^{\gamma,t}$ ;*
2. *if  $d(x, C^{\gamma,t}) \leq 2\gamma$ , then  $x \in D_i^{\gamma,t}$  with  $i = \arg \min_{j=1, \dots, \ell^{\gamma,t}} \{c_j \in C^{\gamma,t} : d(x, c_j) \leq 2\gamma\}$ .*

Moreover, if  $U^{\gamma,t} \neq \emptyset$  then  $|C^{\gamma,t}| = k$ .

*Proof.* At step  $t = 0$ , the property holds by vacuity. Suppose inductively that at step  $t \geq 0$ , the invariant holds for all points in  $X_t$ . We recall that the superscript (e.g.  $U^{\gamma,t}$ ) denotes the state of a set *after* the  $t$ -th insert or delete operation. When the superscript is missing, it denotes a state during the execution of the operations. At time step  $t + 1$ :

- If  $\text{INSERT}(p, \gamma)$  is called, then all points in  $X_t$  are left untouched. Moreover, the invariant for  $p$  holds by construction.
- If  $\text{DELETE}(p, \gamma)$  is called and  $p$  is not a center, then all points in  $X_t \setminus \{p\}$  are left untouched. Let instead  $p$  be a center  $c_i^{\gamma,t}$ . Then, for each  $x \notin D_i^{\gamma,t} \cup U^{\gamma,t}$ , note that the center and cluster index renaming does not affect the invariant. For each  $x \in D_i^{\gamma,t}$ , either  $x$  is inserted in  $D_h$  with  $h = \arg \min_{j=i+1, \dots, \ell} \{c_j \in C: d(x, c_j) \leq 2\gamma\}$  (possibly by creating new clusters), or it is inserted in  $U$  if no center at distance  $\leq 2\gamma$  is found. In the first case, note that  $\min_{j=1, \dots, i-1} d(x, c_j) > 2\gamma$ , otherwise by the inductive hypothesis  $x$  would not belong to  $D_i$ . Therefore, we have that  $h = \arg \min_{j=1, \dots, \ell} \{c_j \in C: d(x, c_j) \leq 2\gamma\}$ , thus fulfilling the invariant.

Finally, let  $U \neq \emptyset$  and consider a  $x \in U$ , either because it belonged in  $U^{\gamma,t}$  or because it was inserted in  $U$  on line 13. In this case, one point  $u$  will be selected as a new center, and thus  $|C| = k$ . We have that if  $d(x, C^{\gamma,t+1}) > 2\gamma$ , then we have  $d(x, u) > 2\gamma$  and therefore  $x$  remains in  $U$ , satisfying the first case of the invariant. Otherwise, since  $x \in U$  implies that  $d(x, c_i) > 2\gamma \forall i = 1, \dots, \ell - 1$ , we have that  $d(x, u) \leq 2\gamma$ , and thus  $x$  is inserted in  $D_\ell$ , thus satisfying the second case of the invariant.

Finally, after a call to  $\text{INSERT}$  or  $\text{DELETE}$ , the  $\text{RECLUSTERING}(\gamma)$  procedure might be called. In that case, we note that each  $x \in \bigcup_{j=1}^{i-1} D_j^{\gamma,t}$  is left untouched, and the invariant is still valid. On the other hand, if  $x \in U'$  after line 4, then it could not be inserted in any of the clusters  $D_j$  for  $j = 1, \dots, i-1$ , by the inductive hypothesis. Moreover, by construction it will be inserted in the cluster  $D_h$  with  $h = \arg \min_{j=i, \dots, \ell} \{c_j \in C: d(x, c_j) \leq 2\gamma\}$ , as a point can be inserted in a cluster only if it is not inserted in a cluster with a lower index. At the end of the loop (i.e.,  $|C| = k$ ), any point  $x$  such that  $d(x, C) > 2\gamma$  is left in  $U$ .

We therefore have the claim by induction.  $\square$

**Theorem 2.3.** *For any time step  $t$ , the solution returned by  $\text{QUERY}_{2+\varepsilon}(t)$  is a  $(2 + \varepsilon)$ -approximation for the  $k$ -center problem on the active set  $X_t$ .*

*Proof.* Recall that the algorithm returns  $C^{\gamma^*,t}$ , with  $\gamma^* = \min_{\gamma \in \Gamma} \{\gamma : U^{\gamma,t} = \emptyset\}$ . If  $\gamma^* = d_{\min}$ , the theorem trivially holds. Suppose  $\gamma^* > d_{\min}$ , and let  $\gamma' = \gamma^*/(1 + \beta) \in \Gamma$ . Clearly,  $U^{\gamma',t} \neq \emptyset$ , hence  $|C^{\gamma',t}| = k$ . Let  $W = C^{\gamma',t} \cup \{x\}$ , for an arbitrary  $x \in U^{\gamma',t}$ . By construction and by Lemma 2.1 we have that  $\forall p \neq q \in W$ ,  $d(p, q) > 2\gamma'$ . Moreover, since  $|W| = k + 1$ , we have that  $r_k^*(X_t) \geq \min_{p \neq q \in W} d(p, q)/2 > \gamma' = \gamma^*/(1 + \beta)$ , as at least two distinct points of  $W$  are in the same optimal cluster. By Lemma 2.1,  $\max_{x \in X_t} d(x, C^{\gamma^*,t}) \leq 2\gamma^* < 2(1 + \beta)r_k^*(X_t) \leq (2 + \varepsilon)r_k^*(X_t)$ .  $\square$

### 2.3.2 Running time and memory complexity

Consider a sequence of  $n$  consecutive calls to the  $\text{UPDATE}_{2+\varepsilon}$  procedure, which must involve  $\Theta(n)$  insertions and  $O(n)$  deletions. In what follows, we analyze the performance of our algorithm relatively to these calls. Let us focus on an arbitrary guess  $\gamma \in \Gamma$ . We first analyze the overall cost of the reclustering operations.

**Lemma 2.2.** *After a call of  $\text{RECLUSTERING}(\gamma)$  every point that was reclustered, i.e., every point belonging to the set  $U'$  in  $\bigcup_{j=i}^{\ell^\gamma} D_j^\gamma \cup U^\gamma$  before the reclustering, becomes vanishing for  $\gamma$ .*

*Proof.* Since at each step  $j$  the new center  $c_j$  gets chosen from all remaining points as the one with maximal deletion time, any point  $x$  that gets inserted into  $D_j$  must have deletion time  $t_{\text{del}}(x) < t_{\text{del}}(c_j)$ , and is therefore vanishing. If a point is left in  $U$ , then it is vanishing by definition.  $\square$

**Lemma 2.3.** *If a point  $x \in X$  is vanishing for  $\gamma$  at time  $t_0$ , then  $x$  remains vanishing for all  $t$  such that  $t_0 \leq t < t_{\text{del}}(x)$ .*

*Proof.* We will show that the vanishing property (for a fixed guess  $\gamma$ ) is invariant under the call of  $\text{INSERT}$ ,  $\text{DELETE}$  or  $\text{RECLUSTERING}$  algorithms, unless  $t = t_{\text{del}}(x)$  and the point is thus removed. The property holds at  $t = 0$  by vacuity and the proof follows by induction.

Fix some  $x \in X$  which is vanishing at time  $t$ .

- **Insertion:** If  $x$  is vanishing then it is already inserted, so any insertion operation will regard a new point and cannot modify either  $x$  or its center. Then,  $x$  will remain vanishing after any insertion operation.
- **Deletion:** Suppose that we call  $\text{DELETE}(p, \gamma)$ . If  $p = x$  the current time is  $t_{\text{del}}(x)$  and the proposition holds. Furthermore, if  $p \notin C^\gamma$  (i.e.  $p$  is not a center) its deletion

will only alter  $p$  itself and no other point. Therefore, we can assume  $p \neq x$  and  $\exists i$  s.t.  $p = c_i$ . We then have 3 cases:

1.  $x \in D_i$ . This would mean that the center of  $x$  has a deletion time smaller than that of  $x$ , which is in contradiction to the fact that  $x$  is vanishing.
2.  $x \in D_j$ , for some  $j \neq i$ . In the case that  $p = c_i$  the deletion algorithm affects clusters with index different from  $i$  by possibly adding more points to them (Line 8) and by relabeling their indexes (Line 16). Therefore, no changes are made to  $x$  or to its center  $c_j$ , and thus  $x$  remains vanishing.
3.  $x \in U$ . In this case,  $x$  remains in  $U$  or is inserted in  $D_l$  (Line 23). Either way, it remains vanishing, since we picked  $c_l = u$  s.t.  $t_{\text{del}}(u)$  is maximal among all points in  $U$ .

- **Reclustering:** Suppose that we call  $\text{RECLUSTERING}(i, \gamma)$ . If  $\exists j < i$  such that  $x \in D_j$ , then the reclustering operation does not affect  $D_j$  in any way, meaning that  $x$  remains vanishing. The other case is covered by Lemma 2.2, and thus  $x$  remains vanishing.

□

**Proposition 2.4.** *The total number of operations made by all calls of  $\text{RECLUSTERING}(\gamma)$  is  $O(n \cdot k)$ .*

*Proof.* The main for loop of Algorithm 12 is repeated  $O(k)$  times. In each iteration we need to scan  $U'$  linearly to first find the element with maximal  $t_{\text{del}}$  and then to check whether some other point in  $U'$  is close to it. So the main loop is executed in  $O(k \cdot |U'|)$ . The other parts of the algorithm can be executed in  $O(|U'|)$ , for a total running time of  $O(k \cdot |U'|)$ . Then:

$$k|U'| = k \left| U \cup \bigcup_{j=i}^k D_j \right| \leq k \left( |U| + \sum_{j=i}^k (|P_j| + |V_j|) \right) \leq 2k \sum_{j=i}^k |P_j|,$$

where the last inequality is guaranteed by the fact that Property (2.3) holds for  $i$  when calling the reclustering.

This means that the number of operations is  $O\left(k \cdot \sum_{j=i}^l |P_j|\right)$  and thus bounded by the number of persistent points involved in the reclustering. Given that by Lemma 2.2 all persistent points in  $P_{j \geq i}$  will become vanishing and that by Lemma 2.3 every vanishing point remains vanishing, the number of persistent points in all reclusterings is bounded

by  $n$  and thus the total number of operations made by all calls of  $\text{RECLUSTERING}(i, \gamma)$  is  $O(n \cdot k)$   $\square$

We now estimate the overall costs of insertions and deletions. For insertions, the following proposition holds because of Proposition 2.4 and because, besides the call to  $\text{RECLUSTERING}(\gamma)$ ,  $\text{INSERT}(p, \gamma)$  takes  $O(k)$  time.

**Proposition 2.5.** *The total number of operations made by all calls of  $\text{INSERT}(p, \gamma)$  is  $O(n \cdot k)$ .*

*Proof.* The proof is trivial since the insertion algorithm simply scans the up to  $k$  centers present in the data structure, which is  $O(k)$ . It can also insert the point inside a cluster or create a new cluster, both of which are  $O(1)$  operations. Given that the total number of insertions is  $n$ , the total number of operations made by all calls of  $\text{INSERT}(p, \gamma)$  must be  $O(n \cdot k)$ .  $\square$

Consider now the deletions. In what follows, for time  $t$  and  $x \in X_t$ , and for each  $\gamma \in \Gamma$ , let  $c_{ID}^{\gamma, t}(x)$  be the index of the cluster in  $D^\gamma$  that  $x$  is in at time  $t$ , and let  $c_{ID}^{\gamma, t}(x) = k + 1$  if  $x \in U^\gamma$ .

**Lemma 2.4.** *Let  $x \in X$  and  $\gamma \in \Gamma$ . Call  $\mathcal{Q}_x^\gamma$  the set of centers  $c$  such that at time  $t = t_{\text{del}}(c) - 1$  it holds that  $x$  is vanishing for  $\gamma$  and  $c_{ID}^{\gamma, t}(c) \leq c_{ID}^{\gamma, t}(x)$ . Then,  $|\mathcal{Q}_x^\gamma| \leq k + 1$ .*

*Proof.* We will show that for  $c \neq x$  to satisfy the two properties implies the fact that when  $c$  was last picked as a center  $x$  was either not inserted or persistent. Suppose by contradiction that, when  $c$  was last picked as a center,  $x$  was vanishing. We have 3 possibilities:

1.  $c$  became a center for the last time at time  $t'$  during an insertion operation (Line 5) or during a re-assignment of persistent points after a deletion (Line 12). In both of these cases  $c_{ID}^{\gamma, t'}(c) > c_{ID}^{\gamma, t'}(x)$  as  $c$  becomes the center of the cluster with the highest index. Given that we are assuming that  $t'$  is the last time that  $c$  was chosen as a center, there cannot be any reclustering involving  $c$  after time  $t'$ . Furthermore, if no reclustering is made, clusters stay in their relative order and point  $x$  will remain inside its cluster until its deletion. So the fact that  $c_{ID}^{\gamma, t'}(c) \geq c_{ID}^{\gamma, t'}(x)$  is in contradiction with  $c_{ID}^{\gamma, t}(c) \leq c_{ID}^{\gamma, t}(x)$ .
2.  $c$  became a center for the last time after a deletion. If  $x \notin U$  the analysis is akin to the one in the first case. If  $x \in U$  we know that  $t_{\text{del}}(x) < t_{\text{del}}(c)$ , thus at time  $t$   $x$  will already be deleted and therefore cannot be vanishing.

3.  $c$  became a center for the last time during the execution of  $\text{RECLUSTERING}(i, \gamma)$ . If  $x \in D_j$  s.t.  $j < i$  the analysis is akin to the one in the first case. Let instead  $x \in U'$  (same  $U'$  as in Algorithm 12). If  $x$  ends up in a cluster with smaller index than the one of  $c$ , then we break the assumption that  $c_{ID}^{\gamma, t}(c) \leq c_{ID}^{\gamma, t}(x)$ . If instead it ends up in a cluster with index greater or equal to the one of  $c$ , then we must have  $t_{\text{del}}(x) < t_{\text{del}}(c)$ , and thus  $x$  cannot be vanishing at  $t = t_{\text{del}}(c) - 1$  because it has been already deleted.

Given that by Lemma 2.3 a vanishing point remains vanishing, calling  $t'$  the time at which  $x$  becomes vanishing, every center  $c$  which satisfies the two properties of the Lemma must be picked at time  $< t'$  and carried on up to a time  $\geq t'$ . Given that at most  $k$  centers can coexist simultaneously, the maximum number of centers  $c \neq x$  satisfying the properties is  $k$ . Accounting for the possibility that  $c = x$ , we obtain at most  $k + 1$  centers.  $\square$

**Proposition 2.6.** *The total number of operations made by all calls of  $\text{DELETE}(p, \gamma)$  is  $O(n \cdot k)$ .*

*Proof.* First, we bound the number of operations done by the re-assignment of cluster points in the loop at Line 5. We begin by noticing that when  $c_i$  is deleted, only persistent points contribute to the loop at Line 5. For a single such point  $x \in D_i, x \neq c_i$ , Lines 6-13 perform  $O(c_{ID}^{\gamma, t}(x) - c_{ID}^{\gamma, t-1}(x))$  operations, i.e., proportional to the difference between the current index  $i$  and the index of the cluster that it will end up in. If  $c_{ID}^{\gamma, t}(x)$  was non-decreasing over time, then the maximum number of operations for  $x$  across all timesteps would be  $O(k)$ . However this is not the case because of re-indexing due to center deletions.

Given that  $x$  must be a persistent point, we are only interested in events in which  $c_{ID}^{\gamma, t}(x)$  decreases and  $x$  remains a persistent point. The only such case occurs during the relabeling process in a deletion that involves a center (Line 16) with index  $< c_{ID}^{\gamma, t-1}(x)$ . In this case  $c_{ID}^{\gamma, t}(x) = c_{ID}^{\gamma, t-1}(x) - 1$ , so each center deletion adds 1 to the maximum number of re-assignment operations for each persistent point with index greater or equal to its own.

Let  $\mathcal{D} := \{c : c \in C^{\gamma, t_{\text{del}}(c)-1}\}$  be the set of all deleted centers at any time. We remark that, for  $x \in X$ , superscripts  $t_{\text{del}}(x) - 1$  denote states right before the deletion of  $x$ . We get that the total number of re-assignment operations is:

$$O\left(n \cdot k + \sum_{c \in \mathcal{D}} \sum_{i \geq c_{ID}^{\gamma, t_{\text{del}}(c)-1}(c)} |P_i^{\gamma, t_{\text{del}}(c)-1}|\right) = O\left(nk + \sum_{c \in \mathcal{D}} \left(|U^{\gamma, t_{\text{del}}(c)-1}| + \sum_{i \geq c_{ID}^{\gamma, t_{\text{del}}(c)-1}(c)} |V_i^{\gamma, t_{\text{del}}(c)-1}|\right)\right),$$

where we used the Property (2.3). If we fix  $c \in \mathcal{D}$  in the last summation and call  $j := c_{ID}^{\gamma, t_{\text{del}}(c)-1}(c)$  we get that, for each  $x \in U^{\gamma, t_{\text{del}}(c)-1} \cup \bigcup_{i \geq j} V_i^{\gamma, t_{\text{del}}(c)-1}$ ,  $x$  is by definition vanishing and  $c_{ID}^{\gamma, t_{\text{del}}(c)-1}(c) \leq c_{ID}^{\gamma, t_{\text{del}}(c)-1}(x)$ . Therefore we have:

$$\begin{aligned} \sum_{c \in \mathcal{D}} \left( |U^{\gamma, t_{\text{del}}(c)-1}| + \sum_{i \geq c_{ID}^{\gamma, t_{\text{del}}(c)-1}(c)} |V_i^{\gamma, t_{\text{del}}(c)-1}| \right) &\leq \sum_{x \in X} \left| \left\{ c \in \mathcal{D} : \begin{array}{l} x \text{ is vanishing at time } t_{\text{del}}(c)-1 \\ \text{and } c_{ID}^{\gamma, t_{\text{del}}(c)-1}(c) \leq c_{ID}^{\gamma, t_{\text{del}}(c)-1}(x) \end{array} \right\} \right| \\ &= \sum_{x \in X} |\mathcal{Q}_x^\gamma| \leq n \cdot (k+1), \end{aligned}$$

where we used a double counting argument by rearranging the summation order, and then applied Lemma 2.4. Hence the total number of operations made by the re-assignment of cluster points in the deleting function over all calls of  $\text{DELETE}(p, \gamma)$  is  $O(n \cdot k)$ . Second, we bound the number of operations made by the creation of a new cluster in the deleting function (Line 17). When creating a new cluster we scan each element in  $U$  for finding the one with greater  $t_{\text{del}}$  and then we do the same for finding which points are within distance  $2\gamma$ . Therefore, the total number of operations is  $O(\sum_{c \in \mathcal{D}} |U^{\gamma, t_{\text{del}}(c)-1}|)$ , which, as above, is  $O(n \cdot k)$ .

Finally, all other operations in the deletion procedure take  $O(k)$  time, which concludes the proof.  $\square$

Based on Propositions 2.5 and 2.6, we can prove the bound for update times. We also state the bound for query times, as well as the memory requirements of our algorithm.

**Theorem 2.4.** *The total number of operations made by a sequence of  $n$  consecutive calls to  $\text{UPDATE}_{2+\varepsilon}$  is  $O(n(\max_t \log |X_t| + (\log \Delta/\varepsilon)k))$ .*

*Proof.* By implementing the priority queue  $Q$  as a heap, each of the  $\Theta(n)$  insertions and  $O(n)$  remove-min operations in  $Q$  costs  $O(\log |X_t|)$ , so their overall cost is  $O(n \cdot \max_t \log |X_t|)$ . By Propositions 2.5 and 2.6, for each  $\gamma \in \Gamma$ , we have that the overall costs of all calls  $\text{INSERT}(p, \gamma)$  and  $\text{DELETE}(p, \gamma)$  is  $O(nk)$ . Since  $|\Gamma| \in O(\log_{1+\varepsilon} \Delta) = O(\log \Delta/\varepsilon)$ , the theorem follows.  $\square$

**Theorem 2.5.** *The number of operations made by each  $\text{QUERY}_{2+\varepsilon}$  call is  $O((\log \Delta/\varepsilon) + k)$ .*

*Proof.* Note that we iterate through the guesses  $\gamma \in \Gamma$ , and for each guess we check in constant time the size of  $C^\gamma$ . This takes  $O(|\Gamma|) = O(\log \Delta/\varepsilon)$  time. Finally, returning the solution takes  $O(|C^{\gamma^*}|) = O(k)$  time.  $\square$

**Theorem 2.6.** *The maximum number of points maintained in memory when processing the pointset  $X$  is  $O((\log \Delta/\varepsilon) \max_t |X_t|)$ .*

*Proof.* Note that, by construction, after the call of  $\text{UPDATE}(p, t)$ , only points of  $X_t$  are maintained in memory. Moreover, for a fixed guess  $\gamma$ , any such point  $p \in X_t$  can belong, by construction, to only one of the sets  $D_i^\gamma, i = 1, \dots, \ell^\gamma$  or  $U^\gamma$ . We therefore have that  $\sum_{i=1}^{\ell^\gamma} |D_i^\gamma| + |U^\gamma| \leq |X_t|$ . Moreover, clearly  $|C^\gamma| \leq |X_t|$  as the centers are all distinct. Since we have  $|\Gamma| \in O(\log \Delta/\varepsilon)$ , the data structure maintains at most  $O((\log \Delta/\varepsilon)|X_t|)$  points. Moreover, the priority queue  $Q$  contains each point of the active set exactly once, for a total of  $|X_t|$  points. Finally, until the following call of  $\text{UPDATE}$ , the number of points stored in the data structures can only decrease (i.e., if  $\text{QUERY}$  calls  $\text{DELETE}$ ), thus concluding the proof.  $\square$

Interestingly, our update times are tight, up to a factor  $O(\log \Delta/\varepsilon)$ , due to the lower bound of  $\Omega(|X|k)$  distance queries given by [4, Theorem 54], which holds even in the offline setting.

## 2.4 Bonus: A new fully dynamic algorithm for $k$ -center

The knowledge of the deletion times in the previously described algorithm is used in 2 cases:

1. When choosing a new center from  $U$  we need to find the one with higher deletion time (Line 19 of Algorithm 11 and Line 11 of Algorithm 12).
2. In the definition of persistent and vanishing points (Definition 2.1).

By modifying these 2 parts of our algorithm we can obtain a fully dynamic algorithm that does not need to know deletion times in advance:

1. Points from  $U$  can be chosen randomly, just like in the algorithm from Chan et al. (Section 2.1).
2. Vanishing points are defined as the points that entered a cluster after choosing a new center from  $U$  (Line 23 of Algorithm 11 and Line 15 of Algorithm 12), persistent points are all the remaining ones.

The rest of the algorithm remains the same.

We conjecture that this variant of our algorithm achieves the same update running time of  $O(n(\max_t \log |X_t| + (\log \Delta/\varepsilon)k))$  in expectation, lowering the SOTA running time for fully dynamic algorithms. Sadly, we didn't manage to prove or disprove this conjecture, but we will still include this variant in the experiments to see if it has some practical applications.

### 2.4.1 Comparison with previous work

Our algorithm mimics the fully-dynamic algorithm from [3], yet by exploiting the known deletion times it lowers the amortized update times from  $O(\frac{\log \Delta}{\varepsilon}k^2)$  to  $O(\frac{\log \Delta}{\varepsilon}k)$ , while having the same approximation ratio. In fact, this allows us to have faster updates even compared to the more complex  $(2 + \varepsilon)$ -approximation algorithm by [4], which has amortized  $O((\log \Delta/\varepsilon) \max_t \log |X_t|(k + \max_t \log |X_t|))$  update time. Importantly, since our algorithm is deterministic, it works even against metric-adaptive adversaries, while the aforementioned two can only deal with oblivious ones. Therefore, in settings where lifetimes are available, it would be wasteful to use a fully-dynamic algorithm, as they are more computationally expensive and with less guarantees, rather than our algorithm with lifetimes. In fact, making use of lifetimes, our algorithm can *sidestep* the metric-adaptive lower bound by [4, Thm. 3], for which no fully-dynamic algorithm with  $O(\text{polylog}|X_t|)$  update times can have  $O(1)$  approximation for large enough  $k$ .

The recent preprint by [6] addresses a dynamic setting akin to the one with lifetimes, while focusing primarily on the memory requirements. Indeed, their algorithm requires  $O((\log \Delta/\varepsilon) \cdot k^3)$  memory, as opposed to our requirements, which are linear in the active set size. However, their algorithm features a high  $O((\log \Delta/\varepsilon) \cdot k^4)$  update time, and, more importantly, features only a  $(6k + 2 + \varepsilon)$  approximation ratio, which is much worse than our almost-tight  $(2 + \varepsilon)$ -approximation.

# Chapter 3

## Dynamic $k$ -Center with Sublinear Memory

This chapter is dedicated to analyzing memory efficient approaches to the  $k$ -Center problem. The algorithms presented in the previous one worked in total memory linear in the number of active points  $|X_t|$ . However, when dealing with large datasets, a working memory linear in the size of the active set might not fit in the fast access memory. In order to develop algorithm that utilize sublinear working memory a compromise on the approximation accuracy must be done. We will now see two  $(6 + \varepsilon)$ - approximation algorithms that build on each other: the first one by Cohen-Addad et al.[2] is limited to the sliding window setting (a special case of the lifetimes setting), the second one is our extension in the lifetimes setting.

### 3.1 Sliding Window

In this section we present the work of Cohen-Addad et al.[2], who gave the first algorithm for the  $k$ -center problem in the *sliding window* model, achieving a  $(6 + \varepsilon)$ -approximation while storing only  $O((k/\varepsilon) \log \Delta)$  points, i.e., a number of points independent of the size of the active set. As in the previous chapter, for ease of readability we will omit the subscript  $t$  indicating the time at which each object exists whenever it is not strictly necessary, keeping in mind that, this being a dynamic algorithm, all the presented data structures change over time. We also omit both correctness and running time proofs as they both follow as a special case of the analysis with lifetimes or can be found on the original paper [2]. In the sliding window model, the stream is a (potentially infinite)

sequence of points and, for a fixed integer  $N > 0$  called the *window size*, the active set  $X_t$  consists of the  $N$  most recent points of the stream. Equivalently, every point  $p$  satisfies  $t_{\text{del}}(p) = t_{\text{arr}}(p) + N$ , so that points expire in the very same order in which they arrive (FIFO). In particular, at each time step exactly one point arrives and at most one point expires, and the deletion order is completely determined by the arrival order.

As in the algorithms of the previous chapter, for each guess  $\gamma \in \Gamma$ , defined by Equation 2.1 and here instantiated with  $\beta = \varepsilon/6$ , the algorithm maintains a data structure that either exhibits  $k + 1$  points at pairwise distance  $> 2\gamma$ , which by Proposition 2.1 certifies  $r_k^*(X_t) > \gamma$ , or produces a clustering of radius  $6\gamma$ . Unlike the  $(2 + \varepsilon)$ -approximations seen so far, the second branch does not yield a clustering of radius  $2\gamma$ , which is precisely the reason for the looser approximation factor.

The data structure kept for each  $\gamma \in \Gamma$  is composed by:

- $A^\gamma = (a_1^\gamma, \dots, a_{\ell^\gamma}^\gamma) \subseteq X_t$ , a list of  $\ell^\gamma \leq k + 1$  *attractors* such that  $d(a_i^\gamma, a_j^\gamma) > 2\gamma$  for all  $i \neq j$ . For  $x \in X_t$  and  $a \in A^\gamma$ , we say that  $x$  is *attracted* by  $a$  if  $t_{\text{arr}}(x) \geq t_{\text{arr}}(a)$  and  $d(x, a) \leq 2\gamma$ .
- $R_{\text{act}}^\gamma = (r_1^\gamma, \dots, r_{\ell^\gamma}^\gamma) \subseteq X_t$ , a set of *active representatives*, where, for  $1 \leq i \leq \ell^\gamma$ ,  $r_i^\gamma = \text{rep}^\gamma(a_i^\gamma)$  is the point with latest deletion time (equivalently, the most recent point) among those attracted by  $a_i^\gamma$ .
- $R_{\text{orph}}^\gamma \subseteq X_t$ , a set of *orphaned representatives*, i.e., representatives whose attractor has expired but which are themselves still in the window. We define  $R^\gamma = R_{\text{act}}^\gamma \cup R_{\text{orph}}^\gamma$ .

As before, if  $|A^\gamma| \geq k + 1$  then the attractors are  $k + 1$  points at pairwise distance  $> 2\gamma$  and form the desired certificate, so we only ever need to produce a clustering when  $|A^\gamma| \leq k$ . In this case, as we will show in the next section (Lemma 3.1), every active point lies within distance  $4\gamma$  of  $R^\gamma$ .

**Update procedure** Since in the sliding window model points arrive one at a time and expire FIFO, the update procedure  $\text{UPDATE}_{6+\varepsilon}^{\text{sw}}(p)$  (see Algorithm 14 for the pseudocode) is invoked once per arriving point  $p$ . For each guess  $\gamma \in \Gamma$ , it first removes from  $R_{\text{orph}}^\gamma$  all orphaned representatives that have expired, and, if the oldest attractor has expired, it calls  $\text{DELETEATTRACTION}$  to discard it. The latter procedure (Algorithm 15) turns the representative of the expiring attractor into an orphan, so that it remains available

to cover future points, and only then removes the attractor itself. The new point  $p$  is finally handed to INSERT (Algorithm 16).

INSERT( $p, \gamma$ ) computes the set  $D$  of attractors within distance  $2\gamma$  of  $p$ . If  $D = \emptyset$ , then  $p$  is far from every attractor and becomes a new attractor, representing itself. If this makes  $|A^\gamma| = k + 2$ , the attractor with earliest deletion time is removed via DELETEATTRACTION. Moreover, whenever  $|A^\gamma| = k + 1$ , all orphaned representatives that expire before the earliest-expiring attractor are discarded: the  $k + 1$  attractors already certify  $r_k^*(X_{t'}) > \gamma$  for every  $t'$  until that attractor expires, so  $\gamma$  cannot be a useful guess in this interval and there is no need to keep points covering it. If instead  $D \neq \emptyset$ , then  $p$ , being the most recent point, becomes the new representative of every attractor in  $D$ .

---

**Algorithm 14:** UPDATE $_{6+\varepsilon}^{\text{sw}}(p)$

---

```

1 foreach  $\gamma \in \Gamma$  do
2   foreach  $q \in R_{\text{orph}}^\gamma$  with  $t_{\text{del}}(q) \leq t_{\text{arr}}(p)$  do
3      $R_{\text{orph}}^\gamma \leftarrow R_{\text{orph}}^\gamma \setminus \{q\}$ 
4   if  $\exists a \in A^\gamma$  with  $t_{\text{del}}(a) \leq t_{\text{arr}}(p)$  then
5     DELETEATTRACTION( $a, \gamma$ )
6   INSERT( $p, \gamma$ )

```

---



---

**Algorithm 15:** DELETEATTRACTION( $a, \gamma$ )

---

```

1  $R_{\text{orph}}^\gamma \leftarrow R_{\text{orph}}^\gamma \cup \{\text{rep}^\gamma(a)\};$ 
2  $R_{\text{act}}^\gamma \leftarrow R_{\text{act}}^\gamma \setminus \{\text{rep}^\gamma(a)\};$ 
3  $A^\gamma \leftarrow A^\gamma \setminus \{a\};$ 

```

---

---

**Algorithm 16:** INSERT( $p, \gamma$ )

---

```
1  $D \leftarrow \{a \in A^\gamma : d(p, a) \leq 2\gamma\};$ 
2 if  $D = \emptyset$  then
3    $A^\gamma \leftarrow A^\gamma \cup \{p\};$ 
4    $rep^\gamma(p) \leftarrow p; \quad R_{act}^\gamma \leftarrow R_{act}^\gamma \cup \{p\};$ 
5   if  $|A^\gamma| = k + 2$  then
6      $a_{old} \leftarrow \arg \min_{a \in A^\gamma} t_{del}(a);$ 
7     DELETEATTRACTION( $a_{old}, \gamma$ );
8   if  $|A^\gamma| = k + 1$  then
9      $t_{min} \leftarrow \min_{a \in A^\gamma} t_{del}(a);$ 
10    foreach  $q \in R_{orph}^\gamma$  with  $t_{del}(q) < t_{min}$  do
11       $R_{orph}^\gamma \leftarrow R_{orph}^\gamma \setminus \{q\}$ 
12 else
13   foreach  $a \in D$  do
14      $R_{act}^\gamma \leftarrow (R_{act}^\gamma \setminus \{rep^\gamma(a)\}) \cup \{p\};$ 
15      $rep^\gamma(a) \leftarrow p;$ 
```

---

**Query procedure** To answer a query, the algorithm iterates through the guesses  $\gamma \in \Gamma$  in increasing order and looks for the smallest  $\gamma$  such that  $|A^\gamma| \leq k$  and a  $2\gamma$ -cover of  $R^\gamma$  of size at most  $k$  can be computed by the greedy strategy of Hochbaum and Shmoys [5]; the centers of this cover are returned as the solution. Since every active point is within  $4\gamma$  of  $R^\gamma$  and every representative is within  $2\gamma$  of a returned center, the triangle inequality yields a clustering of radius  $6\gamma$ , while a guess for which more than  $k$  centers are required provides, by Proposition 2.1, a certificate that  $r_k^*(X_t) > \gamma$ .

**Theorem 3.1** ([2]). *Given a stream with aspect ratio  $\Delta$  and a window of size  $N$ , the algorithm described above maintains a  $(6+\varepsilon)$ -approximate solution for the metric  $k$ -center problem, storing at most  $O((k/\varepsilon) \log \Delta)$  points, with an update time of  $O(k^2 \varepsilon^{-1} \log \Delta)$  per point.*

The space bound rests on the observation that each of  $A^\gamma$ ,  $R_{act}^\gamma$  and  $R_{orph}^\gamma$  contains at most  $k + 1$  points, hence at most  $3(k + 1)$  points are kept for each of the  $|\Gamma| = O(\varepsilon^{-1} \log \Delta)$  guesses. Crucially, this bound exploits the FIFO expiration order to argue that orphaned representatives cannot accumulate. In the next section we show how to lift this construction to the setting with lifetimes, where deletion times are arbitrary, recovering the same bound when deletions follow the arrival order.

## 3.2 With Lifetimes

The algorithm developed in the previous section achieves an almost-tight approximation ratio, but requires memory which is linear in the size of the active set. In this section, we devise an algorithm which requires much smaller memory, constant with respect to the size of the active set, at the cost of a slightly worse approximation guarantee.

The algorithm is based on the sliding-window algorithm of Cohen-Addad et al.[2] presented in Section 3.1. In the sliding window model, deletions occur in the same order as arrivals, which lets the algorithm discard expired attractors and orphaned representatives simply by scanning for the oldest ones. In the setting with lifetimes, instead, deletion times are arbitrary, and we generalize the construction so that the decision of which attractors and representatives to retain is driven by the (known) deletion times rather than by the arrival order. As we will see, the resulting data structure coincides with the one of Section 3.1 when points expire in FIFO order.

Let  $\varepsilon > 0$  be a user-defined approximation tolerance, and let  $\Gamma$  be the set of guesses defined by Equation 2.1, instantiated now with  $\beta = \varepsilon/6$ . At any time  $t$ , the algorithm maintains, for each  $\gamma \in \Gamma$ , a data structure, generalizing the one of the sliding-window algorithm of Section 3.1, which is able to provide a solution of radius  $6\gamma$  or a certificate that no solution of radius  $\gamma$  exists. The data structure consists of the sets:

- $A^\gamma = (a_1^\gamma, \dots, a_{\ell^\gamma}^\gamma) \subseteq X_t$  a set of  $\ell^\gamma \leq k+1$  *attractors* such that  $d(a_i^\gamma, a_j^\gamma) > 2\gamma, \forall i \neq j$ . For  $x \in X_t, a \in A^\gamma$ , we say that  $x$  is attracted by  $a$  if  $t_{\text{arr}}(x) \geq t_{\text{arr}}(a)$  and  $d(x, a) \leq 2\gamma$ .
- $R_{\text{act}}^\gamma = (r_1^\gamma, \dots, r_{\ell^\gamma}^\gamma) \subseteq X_t$  a set of *representatives*, where, for  $1 \leq i \leq \ell^\gamma$ ,  $r_i$  is the point with latest deletion time among the ones attracted by  $a_i \in A^\gamma$ , and it is denoted with  $\text{rep}^\gamma(a_i) = r_i$ . We use  $R_{\text{orph}}^\gamma \subseteq X_t$  to denote the set of representatives whose attractors are no longer in  $A^\gamma$ , and define  $R^\gamma = R_{\text{act}}^\gamma \cup R_{\text{orph}}^\gamma$ .

Observe that if  $|A^\gamma| \geq k+1$ , there are  $k+1$  points at distance  $> 2\gamma$ , which implies that  $\gamma < r_k^*(X_t)$ . Conversely, if  $|A^\gamma| \leq k$ , we will show that any point of  $X_t$  is at distance at most  $6\gamma$  from  $R^\gamma$ .

**Update procedure** When a new point  $p$  arrives, procedure  $\text{UPDATE}_{6+\varepsilon}(p)$  is called (see Algorithm 17 for the pseudocode) to handle the insertion of  $p$ , as well as possible deletions of expired points. For every  $\gamma \in \Gamma$ , the data structures are updated as follows. First, all points expired before time  $t_{\text{arr}}(p)$  are removed from  $A^\gamma$  and  $R^\gamma$ . Then, if  $p$

---

**Algorithm 17:** UPDATE<sub>6+ $\varepsilon$</sub> ( $p$ )

---

```
1 foreach  $\gamma \in \Gamma$  do
2   Remove points  $x$  of  $A^\gamma$  and  $R^\gamma$  with
    $t_{\text{del}}(x) \leq t_{\text{arr}}(p)$ 
3    $E \leftarrow \{a \in A^\gamma : d(a,p) \leq 2\gamma\}$ 
4   if  $E = \emptyset$  then
5      $A^\gamma \leftarrow A^\gamma \cup \{p\}$ 
6      $\text{rep}^\gamma(p) \leftarrow p$ ;  $R^\gamma \leftarrow R^\gamma \cup \{p\}$ 
7     CLEANUP( $p, \gamma$ )
8   else
9      $E' \leftarrow$ 
    $\{a \in E : t_{\text{del}}(\text{rep}^\gamma(a)) < t_{\text{del}}(p)\}$ 
10    if  $E' \neq \emptyset$  then
11       $a \leftarrow$  arbitrary attractor in  $E'$ 
12       $R^\gamma \leftarrow R^\gamma \setminus \{\text{rep}^\gamma(a)\}$ 
13       $\text{rep}^\gamma(a) \leftarrow p$ ;  $R^\gamma \leftarrow R^\gamma \cup \{p\}$ 
```

---

---

**Algorithm 18:** CLEANUP( $p, \gamma$ )

---

```
1 if  $|A^\gamma| = k + 2$  then
2    $a_{\text{min}} \leftarrow$ 
    $\arg \min_{a \in A^\gamma} t_{\text{del}}(a)$ ;
3   Remove  $a_{\text{min}}$  from  $A^\gamma$ ;
4 if  $|A^\gamma| = k + 1$  then
5    $t_{\text{min}} \leftarrow \min_{a \in A^\gamma} t_{\text{del}}(a)$ ;
6   Remove all points  $q$ 
   with  $t_{\text{del}}(q) < t_{\text{min}}$ 
   from  $R^\gamma$ 
```

---

is at distance  $> 2\gamma$  from every attractor  $a \in A^\gamma$ , then  $p$  is added to  $A^\gamma$ , and to  $R^\gamma$  as representative of itself (i.e.,  $p = \text{rep}^\gamma(p)$ ), and procedure CLEANUP( $p, \gamma$ ) is invoked to possibly evict unnecessary points from  $A^\gamma$  and  $R^\gamma$ . Otherwise, an attractor  $a \in A^\gamma$  with  $d(a,p) \leq 2\gamma$  and  $t_{\text{del}}(\text{rep}^\gamma(a)) < t_{\text{del}}(p)$ , if any, is arbitrarily chosen and  $p$  becomes the new representative  $\text{rep}^\gamma(a)$ . If no such attractor is found,  $p$  is discarded.

Procedure CLEANUP( $p, \gamma$ ) (see Algorithm 18 for the pseudocode), works as follows. If  $|A^\gamma| = k + 2$ , then the attractor with minimum deletion time is removed from the set. After this removal, or if the size of  $A^\gamma$  is  $k + 1$  at the start of the procedure, all representatives  $q$  such that  $t_{\text{del}}(q) < t_{\text{min}}$ , where  $t_{\text{min}} = \min_{a \in A^\gamma} t_{\text{del}}(a)$  is the minimum deletion time of an attractor, are removed from  $R^\gamma$ . These removals are justified by the following argument. If  $|A^\gamma| = k + 1$ , then these  $k + 1$  points, which are at distance greater than  $2\gamma$  from one another, provide a certificate that no solution of radius  $\leq \gamma$  exists for any  $X_{t'}$  with  $t \leq t' < t_{\text{min}}$ , hence  $\gamma$  will not be a useful guess until  $t_{\text{min}}$ , and there is no need to keep points that expire in this time interval.

### Query procedure

At any time  $t$ , to obtain a solution to  $k$ -center for the active set  $X_t$ , procedure QUERY<sub>6+ $\varepsilon$</sub> ( $t$ ) is invoked (see Algorithm 19). First, for each  $\gamma \in \Gamma$ , all points  $p$  with  $t_{\text{del}}(p) \leq t$  are removed from  $A^\gamma$  and  $R^\gamma$ . Then, the procedure searches the minimum guess  $\gamma \in \Gamma$  such that  $|A^\gamma| \leq k$  and a  $2\gamma$ -covering  $C^\gamma$  of  $R^\gamma$  of size at most  $k$  can be found using the simple greedy strategy by [5], and this set  $C^\gamma$  is returned as the solution.

---

**Algorithm 19:** QUERY<sub>6+ε</sub>( $t$ )

---

```
1 foreach  $\gamma \in \Gamma$  do
2   | Remove points  $x$  of  $A^\gamma$  and  $R^\gamma$  with  $t_{\text{del}}(x) \leq t$ 
3 for increasing values of  $\gamma \in \Gamma$  such that  $|A^\gamma| \leq k$  do
4   |  $C \leftarrow \emptyset$ ;
5   | foreach  $q \in R^\gamma$  do
6     |   if  $C = \emptyset$  or  $d(q, C) > 2\gamma$  then  $C \leftarrow C \cup \{q\}$ ;
7     |   if  $|C| > k$  then go to next  $\gamma$ ;
8   | return  $C$ 
```

---

### 3.2.1 Approximation

In what follows, we denote with superscript  $t$  the states of data structures either after the call of UPDATE<sub>6+ε</sub>( $p$ ), with  $t_{\text{arr}}(p) = t$ , or after the execution of the first for loop in QUERY<sub>6+ε</sub>( $t$ ). To prove the approximation ratio, we first prove that the algorithm maintains the following invariant.

**Lemma 3.1.** *After any execution of UPDATE<sub>6+ε</sub>( $p$ ), the following properties hold for each  $\gamma \in \Gamma$ , and  $t = t_{\text{arr}}(p)$ :*

1. *If  $|A^{\gamma,t}| \leq k$ , then  $\forall x \in X_t$  we have  $d(x, R^{\gamma,t}) \leq 4\gamma$ .*
2. *If  $|A^{\gamma,t}| = k + 1$ , then  $\forall x \in X_t$  such that  $t_{\text{del}}(x) \geq \min_{a \in A^{\gamma,t}} t_{\text{del}}(a)$ , we have  $d(x, R^{\gamma,t}) \leq 4\gamma$ .*

*Proof.* Let us focus on an arbitrary guess  $\gamma$ . We say that a point  $q \in X_t$  is *relevant* for  $\gamma$  and  $t$  if  $|A^{\gamma,t}| \leq k$  or  $|A^{\gamma,t}| > k$  and  $t_{\text{del}}(q) \geq \min_{a \in A^{\gamma,t}} t_{\text{del}}(a)$ . Clearly, the proof concerns only relevant points. Also, for every point  $q \in X_t$ , we let  $\psi^\gamma(q)$  denote the attractor for which  $q$  was selected as representative by UPDATE<sub>6+ε</sub>( $q$ ) (possibly,  $q$  itself) or any attractor belonging to the set  $E$  during the execution of UPDATE<sub>6+ε</sub>( $q$ ). Clearly,  $d(q, \psi^\gamma(q)) \leq 2\gamma$ . We will argue that for every point  $q \in X_t$  which is relevant for  $\gamma$  and  $t$ , there exists a point  $r \in R^{\gamma,t}$  such that  $\psi^\gamma(r) = \psi^\gamma(q)$ .

If  $q$  is the point arrived at time  $t$  (i.e.,  $q = p$  and  $t_{\text{arr}}(q) = t$ ) then we set  $r = \text{rep}^{\gamma,t}(\psi^\gamma(q))$ , which is possibly  $q$  itself. If  $q$  arrived at an earlier time, it is immediate to argue that since  $q$  is relevant at time  $t$ , it has also been relevant at all times  $t'$ , with  $t_{\text{arr}}(q) \leq t' < t$ , when a point arrived or a query was invoked. Set  $r_0 = \text{rep}^{\gamma,t_{\text{arr}}(q)}(\psi^\gamma(q))$ . We have that  $\psi^\gamma(r_0) = \psi^\gamma(q)$ . Also, it holds that  $t_{\text{del}}(r_0) \geq t_{\text{del}}(q)$ , therefore  $r_0$  is also

relevant at all times  $t'$ , with  $t_{\text{arr}}(q) \leq t' < t$ . Being relevant,  $r_0$  has never been removed from  $R^{\gamma, t'}$  during some invocation of CLEANUP, which only removes non-relevant points. Consequently, since  $r_0$  is in  $R^{\gamma, t_{\text{arr}}(q)}$ , either  $r_0$  is still in  $R^{\gamma, t}$  (and we set  $r = r_0$ ), or it has been expunged from some set  $R^{\gamma, t_{\text{arr}}(r_1)}$ , due to the arrival of some longer-lived representative  $r_1$  at time  $t_{\text{arr}}(r_1) > t_{\text{arr}}(q)$ . Observe that  $\psi^\gamma(r_1) = \psi^\gamma(r_0)$ , and that  $r_1$  is also relevant at time  $t$ . If  $r_1 \in R^{\gamma, t}$ , we set  $r = r_1$ , otherwise, we apply the same reasoning to  $r_1$ . Iterating the argument, we determine a sequence of relevant points  $r_0, r_1, \dots, r_h$ , with  $t_{\text{arr}}(q) \leq t_{\text{arr}}(r_0) < t_{\text{arr}}(r_1) \dots < t_{\text{arr}}(r_h)$  and  $\psi^\gamma(q) = \psi^\gamma(r_0) = \dots = \psi^\gamma(r_h)$ , with  $r_h \in R^{\gamma, t}$ . We then set  $r = r_h$ . The lemma follows immediately since  $\psi^\gamma(r) = \psi^\gamma(q)$  implies  $d(q, r) \leq 4\gamma$ .  $\square$

Based on the invariants, we can prove the approximation ratio of the solutions returned by  $\text{QUERY}_{6+\varepsilon}$ .

**Theorem 3.2.** *For any time  $t$ , the solution returned by  $\text{QUERY}_{6+\varepsilon}(t)$  is a  $(6 + \varepsilon)$ -approximation for the  $k$ -center problem on the active set  $X_t$ .*

*Proof.* Let  $\gamma = (1 + \beta)^i \in \Gamma$  be the guess such that the returned solution  $C$  is computed from  $R^{\gamma, t}$ , for some  $\lfloor \log_{1+\beta} d_{\min} \rfloor \leq i \leq \lceil \log_{1+\beta} d_{\max} \rceil$ . By construction,  $|A^{\gamma, t}| \leq k$ , and, for each  $p \in R^{\gamma, t}$ ,  $d(p, C) \leq 2\gamma$ . Using Lemma 3.1 and the triangle inequality, we have that for each  $p \in X_t$ ,  $d(p, C) \leq 6\gamma$ . Now, if  $i = \lfloor \log_{1+\beta} d_{\min} \rfloor$  then  $\gamma \leq r_k^*(X_t)$  and the theorem holds. Instead, if  $i > \lfloor \log_{1+\beta} d_{\min} \rfloor$ , we have, by construction, that for  $\gamma' = (1 + \beta)^{i-1} = \gamma / (1 + \beta)$  either  $|A^{\gamma', t}| = k + 1$  or  $k + 1$  points of  $R^{\gamma', t}$  have been found at distance  $> 2\gamma'$  from one another. From what we observed earlier, in either case we have  $\gamma' < r_k^*(X_t)$ , hence  $\gamma < (1 + \beta)r_k^*(X_t)$ . Thus, for each  $p \in X_t$ ,  $d(p, C) \leq 6\gamma < 6(1 + \beta)r_k^*(X_t) \leq (6 + \varepsilon)r_k^*(X_t)$ , and the theorem follows.  $\square$

### 3.2.2 Running time and memory complexity

We now analyze the number of points maintained by the data structure employed by our algorithm, which directly influences the time and memory requirements. The analysis is made parametric with respect to a specific notion of ordering of the input dataset  $X$ , which is defined below and it is a crucial novel ingredient introduced by this thesis.

**Definition 3.1.** *For an integer parameter  $H \geq 0$ , we say that the set  $X$  is  $H$ -ordered if for every pair of points  $p, q \in X$  such that  $t_{\text{arr}}(p) < t_{\text{arr}}(q)$  and  $|\{x \in X : t_{\text{arr}}(p) < t_{\text{arr}}(x) < t_{\text{arr}}(q)\}| \geq H$ , i.e., at least  $H$  points arrive between  $p$  and  $q$ , we have that  $t_{\text{del}}(p) < t_{\text{del}}(q)$ .*

Intuitively, the above notion provides a quantitative characterization of how arbitrary the points' arrival and departure times are, and a low value of  $H$  means that points are deleted roughly in the order in which they arrive. For example, in the sliding-window setting, or in general when points are deleted in the order they arrive, the pointsets are 0-ordered. The assumption of  $H$ -ordered sets is realistic: consider for example a stream with batch processing (e.g., buffers in TCP packet routers, or intermediate storage for industrial manufacturing). Items arrive in sequence and are grouped into batches of size  $H$ , with each batch being processed as a unit. While the deletion order of items within a batch may change, batches themselves are served in arrival order, resulting in an  $H$ -ordered set.

Importantly, the algorithm does not require the knowledge of  $H$ , which is only used by the analysis. Indeed, under  $H$ -ordered sets, we have the following result.

**Lemma 3.2.** *Suppose that  $X$  is  $H$ -ordered, for some  $H \geq 0$ . Then, for any guess  $\gamma \in \Gamma$  and time  $t$ ,*

$$|A^{\gamma,t}| + |R^{\gamma,t}| = O(\min\{(k+H), |X_t|\}).$$

*Proof.* Let us fix arbitrary  $\gamma$  and  $t$ , and observe that, straightforwardly,  $|A^{\gamma,t}| + |R^{\gamma,t}| \leq 2|X_t|$ . We will now show that  $|A^{\gamma,t}| + |R^{\gamma,t}| = O(k+H)$ . By construction, we know  $|A^{\gamma,t}| \leq k+1$ , so we only have to show that  $|R^{\gamma,t}| = O(k+H)$ . Let  $A_{\text{all}}^\gamma(t)$  denote the set of all attractors ever arrived up to time  $t$ , and let  $q$  be the one with largest arrival time among those not in  $A^{\gamma,t}$ , namely,  $q = \arg \max_{a \in A_{\text{all}}^\gamma(t) \setminus A^{\gamma,t}} t_{\text{arr}}(a)$ . Hence,  $q$  arrived at some time  $t_{\text{arr}}(q) < t$  and was removed from  $A^\gamma$  at some time  $t'$ , with  $t_{\text{arr}}(q) \leq t' < t$ , either by Line 2 of the UPDATE procedure, because it expired, or by Line 3 of the CLEANUP procedure. In the time interval  $(t_{\text{arr}}(q), t]$  no more than  $k+1$  new attractors arrived since, otherwise, one of the attractors arrived in  $(t_{\text{arr}}(q), t]$  would not belong to  $A^{\gamma,t}$ , because  $A^\gamma$  can never contain more than  $k+1$  points, and this would contradict the choice of  $q$ . For a point  $p \in R^{\gamma,t}$ , let  $\psi^\gamma(p)$  denote the attractor for which  $p$  is representative. We partition  $R^{\gamma,t}$  into 3 subsets:

- $R^{\gamma,t}(1) = \{p \in R^{\gamma,t} : t_{\text{arr}}(\psi^\gamma(p)) \in (t_{\text{arr}}(q), t] \}$
- $R^{\gamma,t}(2) = \{p \in R^{\gamma,t} : \psi^\gamma(p) \in A^{\gamma, t_{\text{arr}}(q)} \}$
- $R^{\gamma,t}(3) = \{p \in R^{\gamma,t} : \psi^\gamma(p) \notin A^{\gamma, t_{\text{arr}}(q)} \text{ and } t_{\text{arr}}(\psi^\gamma(p)) < t_{\text{arr}}(q) \}$

It is immediate to see that  $R^{\gamma,t} = R^{\gamma,t}(1) \cup R^{\gamma,t}(2) \cup R^{\gamma,t}(3)$ . Moreover,  $|R^{\gamma,t}(1)| \leq k+1$  from what we observed before, and  $|R^{\gamma,t}(2)| \leq k+1$  since  $|A^{\gamma, t_{\text{arr}}(q)}| \leq k+1$  by construction. We now argue that  $|R^{\gamma,t}(3)| = O(H)$ . First note that all points of  $R^{\gamma,t}(3)$

must have arrived prior to  $t_{\text{arr}}(q)$ , since, by definition of  $R^{\gamma,t}(3)$ , the attractors for which they are representatives, entered and left  $A^\gamma$  before  $t_{\text{arr}}(q)$ . Let  $p$  be the point of  $R^{\gamma,t}(3)$  with earliest arrival time. So, all other points of  $R^{\gamma,t}(3)$  must have arrived in the open interval  $(t_{\text{arr}}(p), t_{\text{arr}}(q))$ . Thus, to show that  $|R^{\gamma,t}(3)| = O(H)$  it is sufficient to argue that less than  $H$  points arrived in  $(t_{\text{arr}}(p), t_{\text{arr}}(q))$ . By contradiction, suppose that this is not the case. Then, the assumption that  $X$  is  $H$ -ordered implies that  $t_{\text{del}}(p) < t_{\text{del}}(q)$ . Since  $q$  was removed from  $A^\gamma$  at time  $t' < t$ , we have two cases:

- $q$  expired at  $t'$  and was removed at Line 2 of the UPDATE procedure. Hence,  $t_{\text{del}}(p) < t_{\text{del}}(q) = t' < t$ , and  $p$  cannot be in  $R^{\gamma,t}$ .
- $q$  was removed from  $A^\gamma$  at time  $t'$  at Line 3 of the CLEANUP procedure, since  $q$  had minimum deletion time among  $k + 2$  attractors. Then, right after the removal of  $q$  from  $A^\gamma$ ,  $t_{\text{min}}$  was set to a value  $\geq t_{\text{del}}(q)$  and all representatives  $r$  with  $t_{\text{del}}(r) < t_{\text{del}}(q) \leq t_{\text{min}}$  were removed. Therefore, since  $t_{\text{del}}(p) < t_{\text{del}}(q)$ , then  $p$  would have been removed as well.

In either case,  $p$  could not be in  $R^{\gamma,t}$ , which contradicts the choice of  $p$ .  $\square$

The next theorems, based on Lemma 3.2, state upper bounds for the update and query running times, and for the memory complexity. Let  $\text{pr}(t)$  be the time of the last update/query operation before  $t$ .

**Theorem 3.3.**  $\text{UPDATE}_{6+\varepsilon}(p)$  requires  $O((\log \Delta/\varepsilon) \min\{(k+H), |X_{\text{pr}(t_{\text{arr}}(p))}|\})$  operations.

**Theorem 3.4.**  $\text{QUERY}_{6+\varepsilon}(t)$  requires  $O((\log \Delta/\varepsilon)k \cdot \min\{(k+H), |X_{\text{pr}(t)}|\})$  operations.

**Theorem 3.5.** The maximum number of points maintained in memory when processing the pointset  $X$  is  $O((\log \Delta/\varepsilon) \min\{(k+H), \max_t |X_t|\})$ .

Note that, for  $H \in O(k)$ , update times and the memory complexity are bounded by  $O((\log \Delta/\varepsilon)k)$ .

### 3.2.3 Comparison with related work

The memory and time complexities of the algorithm presented in this chapter depend on the tameness of the update sequences. Importantly, the memory usage never exceeds the one of the  $(2+\varepsilon)$ -approximation algorithm presented in Section 2.3, and, for small  $H$  (e.g.,  $H \in O(k)$ ), it is significantly smaller memory, while attaining comparable update time.

Moreover, unlike any existing fully-dynamic algorithm in general metrics, the bound on the operations is *worst-case* rather than amortized. These gains, however, come at the cost of a (slightly) worse approximation ratio.

Our setting can be viewed as a generalization of the sliding window one. Indeed, our algorithm extends the sliding-window one of [2], described in Section 3.1, by allowing arbitrary deletions. Despite the increased generality, we match the same memory and time complexities when  $H = 0$ . As discussed in Section 2.4.1, [6] present an algorithm for  $k$ -center in a setting akin the one with lifetimes, which stores  $O((\log \Delta/\varepsilon) k^3)$  points regardless of  $H$ -ordering and is more memory-efficient than ours for large  $H$  (i.e.,  $H \in \Omega(k^3)$ ). However, their algorithm has expensive update times, and, more importantly, a worse non-constant approximation ratio, growing linearly with  $k$ .



# Chapter 4

## Experiments

In this chapter we show the empirical experiments we did on the previously presented algorithms. All the experiments were ran on a laptop with 8 core AMD Ryzen 7 5800U processor and 16GB of RAM. You can find all the source code used on github [7]. Of all the shown algorithms we made our own implementation. Only exceptions are Chan's algorithm, of which we also picked the original implementation of the authors [8], and the Cohen-Addad's one, which we did not include as it is just a special case of our low memory algorithm and no public implementation was made available by the authors.

We did make use of generative AI to help with the process of writing code, although we tried our best to not blindly trust the models on everything. Specifically:

- All algorithms were first implemented from scratch by hand and then fed to Claude for debugging.
- The python scripts for coordinating the experiments, generating the datasets and plotting the results were generated by Claude and reviewed afterwards.

In the experiments we mix fully dynamic algorithms with algorithms with lifetimes and, for simplicity, we feed all algorithms the same datasets. This means that fully dynamic algorithms will read in input the deletion times, but they discard it immediately. Furthermore, algorithms with lifetimes will have access to deletions in the stream the moment they happen, but they will ignore them as in the model with lifetimes only insertions should be streamed. We divided the measured execution time in **update** and **query time**, we also left the input streaming out of the measured time, in order to more accurately measuring the difference between implementations. We also kept track of the **maximal number of saved points**, for evaluating our  $(6 + \varepsilon)$ -approximation.

We will refer to the algorithms by the following names:

- **gonzalez**: Gonzalez’ 2-approximation, described in Chapter 1. This algorithm is included as a sanity check for empirically estimating the approximation of our implementations and it is not modified for cleverly handling insertions or deletions.
- **chan-original**: Chan’s fully dynamic algorithm described in section 2.1.
- **bateni**: Our implementation of Bateni’s fully dynamic algorithm described in section 2.2 (no implementation was provided by the authors).
- **ours2eps**: Our  $(2 + \varepsilon)$ -approximation algorithm with lifetimes, described in section 2.3.
- **ours2epsrandom**: The fully dynamic random variant of our  $(2 + \varepsilon)$ -approximation algorithm, as described in 2.4.
- **ours6eps**: Our  $(6 + \varepsilon)$ -approximation algorithm with low memory described in section 3.2.

We compare our algorithms on 2 types of datasets:

- **random**: In this dataset points are picked from a uniform distribution in  $[0,1]^d$ . Insertion and deletion times are picked randomly with constraints as to satisfy the maximal number of active points and possible  $H$ -orderdness of the dataset.
- **gowalla**: This is a real life dataset taken from SNAP [9]. It represents a social network in which people can share with others information on their real life locations by check-ins. We are not interested in the network of user friendships, rather than on the check-ins of each user. Each check-in has a clear timestamp, along with latitudinal and longitudinal coordinates of the location the user was when it took the check-in. The deletion time of each check-in is represented by the following time the user made another check-in (the last check-in is discarded). You can look at the following table to see an example of the raw data:

User	Check-in Time	Latitude	Longitude
196514	2010-07-24T13:41:10Z	53.364905	-2.270824
196514	2010-07-24T13:43:18Z	53.3679640626	-2.2792943689
196514	2010-07-24T13:44:08Z	53.3675663377	-2.278631763
196514	2010-07-24T13:44:26Z	53.3674087524	-2.2783813477
196514	2010-07-24T13:44:38Z	53.3663709833	-2.2700764333
196514	2010-07-24T13:44:46Z	53.3653895945	-2.2754087046
196514	2010-07-24T13:44:58Z	53.360511233	-2.276369017
196514	2010-07-24T13:45:06Z	53.3648119	-2.2723465833

Table 4.1: Example of user check-ins in Gowalla dataset.

Given that coordinates are on a spherical surface we use the great-circle distance as our metric for this dataset, whilst the other two are Euclidean.

We will investigate the impact on the performance of many different parameters, the following table is a recap of what each parameter is, along with its default value if not otherwise specified:

Parameter	Description	Default Value
$N$	Total number of insertions	10000
<code>max_n_active</code>	Maximal number of active points	1000
$d$	Dimensionality of the points	3
$k$	Number of centers	10
$\varepsilon$	Approximation accuracy parameter	0.5
$H$	$H$ -ordering of the dataset, as defined in Section 3.2	$\infty$

Table 4.2: Parameter Descriptions

## 4.1 Approximation

We will now show the comparison of the algorithms against Gonzalez’ 2-approximation in order to check whether we made a major implementation mistake. We are forced to use Gonzalez’ approximation instead of finding the optimal radius as it will be unfeasible to compute for a sensible number of active points and  $k$  (a solution which checks every possible center selection runs in  $\Omega(|X_t|^k)$ ). Even with the Gonzalez’ algorithm we are limited to a low number of active points, because each query requires  $\Omega(n \cdot k)$  to run.

The comparison is made by running all algorithms in the exact same datasets with the same queries. Afterwards, for each query, the radius induced by the centers found by Gonzalez is taken as a baseline for the optimal solution and is compared with the other algorithms. We test the algorithms across various values of  $k$  on the random and Gowalla datasets. We fix  $\varepsilon = 0.5$ . For each run (i.e. each different pair of  $k$  and dataset), we randomly ask a query after each insertion with 5% probability. The plots show the average ratio of the radius found by each algorithm and the radius found by Gonzalez across all queries made in one run.




Figure 4.1: random

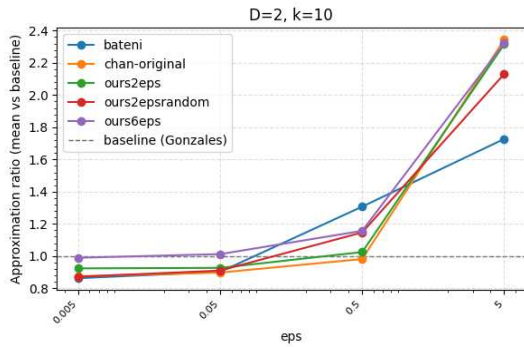


Figure 4.2: gowalla

We can see that all  $(2 + \varepsilon)$ -approximation algorithms exhibit a clear degradation of the approximation factor as  $\varepsilon$  increases as expected. The found solution is even better than Gonzalez' for  $\varepsilon = 0.05$ . For higher values of  $\varepsilon$ , if we assume that Gonzalez found a 2-approximation of the optimal solution and therefore we multiply the average ratio by 2, we do not encounter an approximation higher than  $2 + \varepsilon$  for all  $(2 + \varepsilon)$ -approximation algorithms, with just one exception made by Bateni in the Gowalla dataset with  $\varepsilon = 0.5$ . The exception is almost surely due to the fact that Gonzalez can still find a solution more precise than a worse possible 2-approximation and it isn't strong enough to conclude that Bateni's algorithms might produce an approximation worse than  $(2 + \varepsilon)$ .

In the random dataset, the  $(6 + \varepsilon)$ -approximation, while being clearly worse than the counterparts, doesn't come near the  $6 + \varepsilon$  mark. For example, with  $\varepsilon = 0.2$ , we obtain a 1.4 ratio with respect to Gonzalez, which is less than a  $2.8 < 6.2$  approximation of the optimal radius. Moreover, in the Gowalla dataset the  $(6 + \varepsilon)$ -approximation is nearly indistinguishable from the other algorithms. This means that  $6 + \varepsilon$  is a purely theoretical upper bound which can be much lower in practice.

## 4.2 Running Time

Here we evaluate the timing performance of the algorithms starting with the update time.

We start by looking at the impact of parameters  $\varepsilon$  and  $d$ . Both plots are scaled logarithmically on both axes. We consider the random dataset.

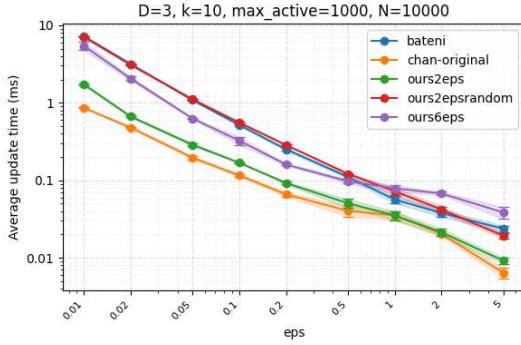


Figure 4.3: Plot over  $\varepsilon$

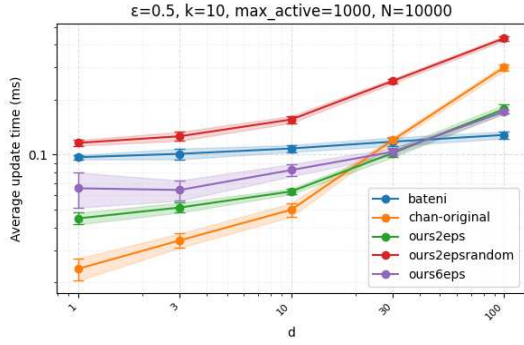


Figure 4.4: Plot over  $d$

The plot of  $\varepsilon$  shows exactly what we expected: the running time decreases with  $\varepsilon$ . The one of  $d$ , whilst still having an upwards trend across all the algorithms as expected, allows us to guess the proportion of the time each algorithm dedicates to computing distances. This is due to the fact that the only part in which  $d$  influences the running time is inside the distance calculation, which is the same function across all algorithms. Bateni shows an almost flat graph, indicating that distance computations do not occupy the majority of the time, probably due to the priority queue updates at each guess. The algorithm from Chan is the most efficient in terms of non-distance computations, as we can see by its low running time when  $d = 3$ , but pays the  $O(k^2)$  distance computed per-update when  $d$  gets large enough.

We now show the difference in runtime across multiple values of  $k$ . As before, plots are in loglog scale. We will look at both the random and Gowalla dataset to also get a sense of how well the algorithm run on real world data.

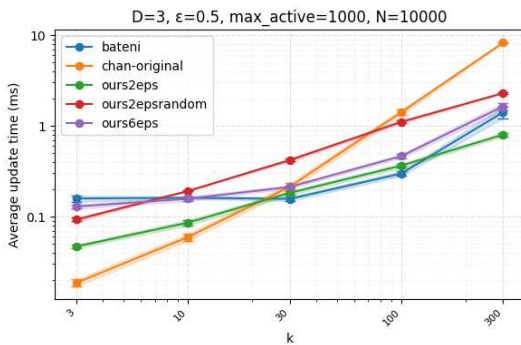


Figure 4.5: Random dataset

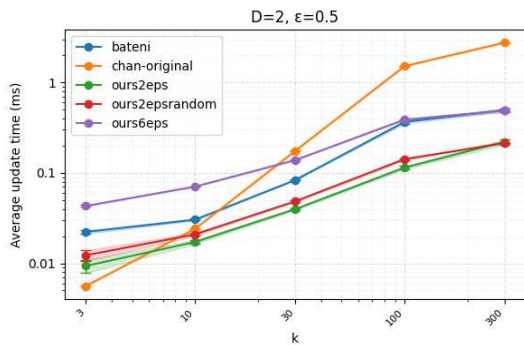


Figure 4.6: Gowalla dataset

In both figures we can visually see the quadratic dependency on  $k$  of Chan’s algorithm, which in loglog plot is indicated by a steeper line. Bateni’s algorithm has a weird behaviour on the fully random dataset, as it seems to first lower the update time up to  $k = 30$  before finally starting to grow. This behaviour is probably due to the geometrical disposition of the random point cloud and might need a complicated and ultimately useless analysis to justify. What’s interesting is that both our lifetimes algorithm and its fully dynamic counterpart perform really well on real world data, despite showing no practical improvement on random data against the state of the art.

The last parameter we want to control is the number of active points, given that the theoretical improvement of our algorithm against Bateni’s stands from the removal of a multiplicative logarithm in this number. For this experiment we will only compare our  $(2 + \varepsilon)$ -algorithms against Bateni and we will do so in the random dataset, which lets us control the maximal number of active points as oppose to the Gowalla dataset. Given that we might need to have an high number of active points for this experiment we set  $n = 300000$ . The plots are only logarithmic in the x-axis.

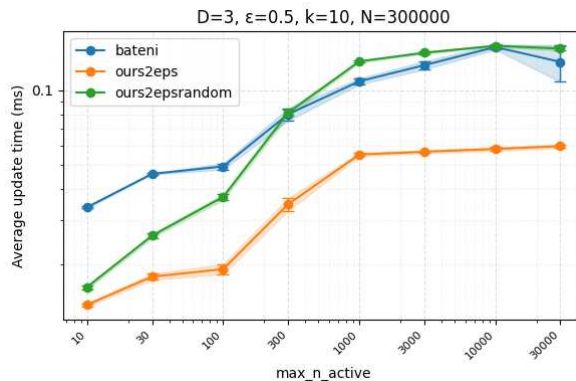


Figure 4.7: Plot over `max_n_active`

While we see a better performance of our algorithm with lifetimes against Bateni, the picture is not clear enough. We see from all algorithm an increase of runtime from 100 to 1000 active points, for it to remain almost constant in the remaining part of the plot. The same behaviour was observed with other values of  $N$  and  $k$ , and it is probably due to the random dataset not being able to produce an adversarial sequence of operations that make Bateni run in the worst possible running time. We could try to find an ad-hoc adversarial dataset, but that would be beyond the scope of the thesis.

To conclude this section, we analyze the query time. We mainly do this analysis as our  $(6 + \varepsilon)$ -approximation is influenced by parameters  $k$  and  $H$  in the query cost. Therefore, we look at how query time behaves across both  $k$  and  $H$  on the random dataset.

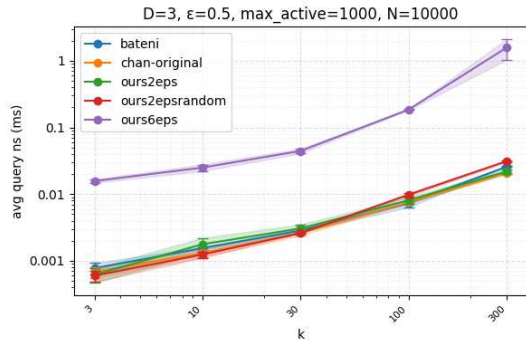


Figure 4.8: Query time over  $k$

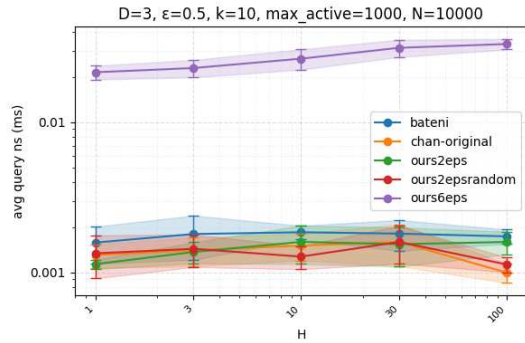


Figure 4.9: Query time over  $H$

The  $(6 + \varepsilon)$ -approximation exhibits a clear difference in query time with respect to the others, this is due to the fact that the query algorithm needs to find a clustering across all representative and does not maintain  $k$  centers per guess like the other algorithms. We can see an increase of query cost with respect to  $k$  across all algorithms, as the  $k$  centers still have to be printed in output. The  $H$  value does not affect the  $(2 + \varepsilon)$ -approximations at all as expected, the only affected algorithm is the  $(6 + \varepsilon)$  one, with a slight increase in query time.

### 4.3 Memory Consumption

In the analysis of our  $(6 + \varepsilon)$ -approximation the notion of  $H$ -ordering comes up. This brief section aims to find if the dependency on the  $H$ -ordering of the dataset on the number of saved points is purely theoretical or has some real impact in the number of saved points. We ran our experiment on the random dataset across different values of  $H$  and  $k$ , here we report the maximal total number of points saved by the algorithm.

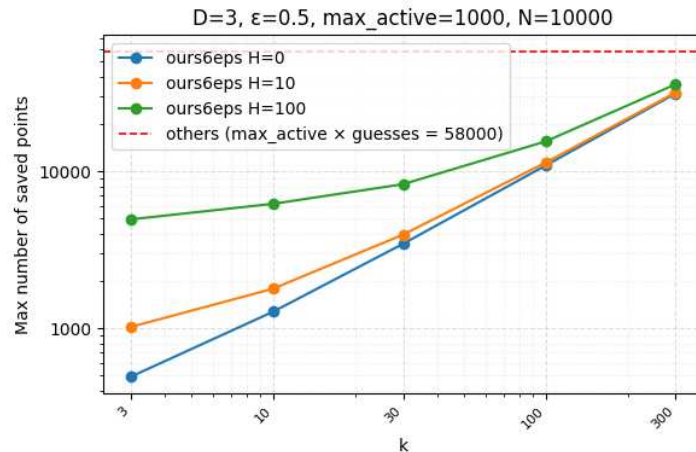


Figure 4.10: Memory consumption over  $k$

We see that  $H$  has a tangible additive impact on the number of saved points, confirming theoretical analysis. We also see a linear increase in points saved with respect to  $k$ , but this was obvious by the fact that we need to save up to  $k + 1$  attractors to certify the unfeasibility of a guess. It remains unclear whether this dependency on  $H$  is due only to the algorithm in question or if it exists an information theoretical lower bound that proves this for any algorithm, further research in this direction is needed.

## 4.4 Conclusions

The theoretical advantage of utilizing our algorithms when the knowledge of lifetimes is available is undisputable: we obtain a  $O(\log(n))$  reduction on the running time of the  $(2 + \varepsilon)$ -approximation and set a new baseline for the low memory case, as no algorithms achieved a  $6 + \varepsilon$  approximation outside the sliding window model.

For the practicality of our approach the picture is not so decisive, but it is still in our favour, at least on some cases.

As of these implementations, the algorithm from Chan et al. seems to perform better than the counterparts for small values of  $k$ . This can be inherent to the additional complexity of needing to handle persistent and vanishing points, but we suspect that the original implementation of the authors might simply be better than what we came up with. Moreover, for not so high values of  $k$  (e.g. 30) the algorithm becomes worse than the others.

Across all the experiments, the algorithm from Bateni et al. managed to outperform us only on the random dataset with a really high value of  $d$  and matched our performances in the random dataset with values of  $k$  between 30 to 100. Moreover, in the real world case the performance of Bateni stays consistently above both our algorithm with lifetimes and our fully dynamic variant.

The  $(6 + \varepsilon)$ -approximation also seems to perform well on real world data, maintaining an approximation on par with every other algorithm, making it competitive even against theoretically better algorithms.

This means that the algorithms we developed have not only good theoretical guarantees, but promise better performances on practical applications and could be used in real world scenarios to do fast data analysis.

# Bibliography

- [1] T. F. Gonzalez, “Clustering to minimize the maximum intercluster distance,” *Theoretical Computer Science*, vol. 38, pp. 293–306, 1985.
- [2] V. Cohen-Addad, C. Schwiegelshohn, and C. Sohler, “Diameter and k-center in sliding windows,” in *International Colloquium on Automata, Languages, and Programming (ICALP)*, 2016, pp. 19–1.
- [3] T. H. Chan, A. Guerqin, and M. Sozio, “Fully dynamic k-center clustering,” in *World Wide Web Conference (WWW)*, 2018, pp. 579–587.
- [4] M. Bateni, H. Esfandiari, H. Fichtenberger, *et al.*, “Optimal fully dynamic k-center clustering for adaptive and oblivious adversaries,” in *Symposium on Discrete Algorithms (SODA)*, 2023.
- [5] D. S. Hochbaum and D. B. Shmoys, “A unified approach to approximation algorithms for bottleneck problems,” *Journal of the ACM (JACM)*, vol. 33, no. 3, pp. 533–550, 1986.
- [6] L. Blank, S. Cabello, M. Hajiaghayi, *et al.*, “The general expiration streaming model: Diameter,  $k$ -center, counting, sampling, and friends,” *arXiv preprint arXiv:2509.07587*, 2025.
- [7] S. Moretti, *Experiments included in this thesis*, <https://github.com/simpatine/dynamic-k-center-lifetimes>, Accessed: 2026-07-01.
- [8] T. H. Chan, A. Guerqin, and M. Sozio, *Fully dynamic k-center clustering experiments’ github page*, <https://github.com/fe6Bc5R4JvLkFkSeExHM/k-center>, Accessed: 2026-07-01.
- [9] E. Cho, E. Myers, and J. Leskovec, *Snap: Gowalla dataset*, <https://snap.stanford.edu/data/loc-Gowalla.html>, Accessed: 2026-06-27.