
A Constant Approximation Algorithm for Sequential Random-Order No-Substitution k -Median Clustering

Tom Hess

Department of Computer Science
Ben-Gurion University of the Negev
Beer-Sheva, Israel
tomhe@post.bgu.ac.il

Michal Moshkovitz

Department of Computer Science
Tel-Aviv University*
Tel Aviv, Israel
mmoshkovitz@eng.ucsd.edu

Sivan Sabato

Department of Computer Science
Ben-Gurion University of the Negev
Beer-Sheva, Israel
sabatos@cs.bgu.ac.il

Abstract

We study k -median clustering under the sequential no-substitution setting. In this setting, a data stream is sequentially observed, and some of the points are selected by the algorithm as cluster centers. However, a point can be selected as a center only immediately after it is observed, before observing the next point. In addition, a selected center cannot be substituted later. We give the first algorithm for this setting that obtains a constant approximation factor on the optimal cost under a random arrival order, an exponential improvement over previous work. This is also the first constant approximation guarantee that holds without any structural assumptions on the input data. Moreover, the number of selected centers is only quasi-linear in k . Our algorithm and analysis are based on a careful cost estimation that avoids outliers, a new concept of a linear bin division, and a multiscale approach to center selection.

1 Introduction

Clustering is a fundamental unsupervised learning task used for various applications, such as anomaly detection [Leung and Leckie, 2005], recommender systems [Shepitsen et al., 2008] and cancer diagnosis [Zheng et al., 2014]. In recent years, the problem of *sequential clustering* has been actively studied, motivated by applications in which data arrives sequentially, such as online recommender systems [Nasraoui et al., 2007] and online community detection [Aggarwal, 2003].

In this work, we study k -median clustering in the sequential *no-substitution* setting, a term first introduced in Hess and Sabato [2020]. In this setting, a stream of data points is sequentially observed, and some of these points are selected by the algorithm as cluster centers. However, a point can be selected as a center only immediately after it is observed, before observing the next point. In addition, a selected center cannot be substituted later. This setting is motivated by applications in which center selection is mapped to a real-world irreversible action. For instance, consider a stream of website users, where the goal is to instantaneously identify users that serve as social cluster centers and provide them with a promotional gift while they are still on the website. As another example, consider recruiting representative participants for a clinical trial out of a stream of incoming patients.

*This work was done while the author was at the University of California San Diego.

The goal in the no-substitution k -median setting is to obtain a near-optimal k -median cost value, while selecting a number of centers that is as close as possible to k . For an adversarially ordered stream, it has been shown [Moshkovitz, 2021] that an algorithm that selects a number of centers that is sublinear in the stream length cannot obtain a constant approximation guarantee, unless some structural assumptions are imposed on the input data.

In this work, we show that in contrast, when the stream order is random, a constant approximation guarantee can be obtained without imposing structural assumptions on the data. Previous works on the random-order setting provide either an algorithm with an almost-constant approximation assuming a bounded metric space [Hess and Sabato, 2020] or an algorithm with an approximation factor that is exponential in k [Moshkovitz, 2021]. Thus, our guarantees are stronger and more general than the previous guarantees. Moreover, the number of centers selected by our algorithm is only quasi-linear in k and does not depend on the stream length.

Main result. We propose a new algorithm, MUNSC (Multiscale No-Substitution Clustering), that obtains a constant approximation factor with probability $1 - \delta$ over the random order of the stream, while selecting only $O(k \log^2(k/\delta))$ centers. MUNSC operates by executing several instances of a new center-selection procedure, called `SELECTPROC`, where each instance is applied to a stream prefix of a different scale. `SELECTPROC` decides which centers to select using a novel technique, in which a small stream prefix is used to estimate a truncated version of the optimal cost of the entire stream. This truncated version ignores outliers, and so can be estimated reliably. The value of the estimate is used to determine which centers to select from the rest of the stream. The multiscale use of `SELECTPROC` by MUNSC allows MUNSC to select only a quasi-linear number of centers.

As part of our analysis, we propose a new concept of a *linear bin division*, which allows a high-probability association of the costs of stream subsets, while selecting a number of centers that is independent of the stream length. This improves a previous construction of Meyerson et al. [2004].

The guarantees for MUNSC are provided in our main theorem, Theorem 4.1, stated in Section 4. Table 1 below compares our guarantees to previous works; See Section 2 for additional details. We note that our algorithm and guarantees are easily extendable to k -means clustering, as well as to any other clustering objective that satisfies the weak triangle inequality. MUNSC uses as a black box an offline k -median approximation algorithm with a constant approximation factor. Whenever the offline algorithm is efficient [e.g., Charikar et al., 2002, Li and Svensson, 2016], then so is MUNSC.

Limitations. The proposed algorithm is the first constant-approximation no-substitution clustering algorithm that makes no structural assumptions on the input data set. The only limitation is the assumption that the input order is random. This assumption is useful in cases where there is no adversary and any order is as likely, as in the example applications mentioned above. Moreover, a random order is a standard assumption in many streaming settings, satisfied also whenever the input stream is drawn i.i.d. from a distribution. In some cases, the randomness assumption might hold only approximately. In these cases, we expect a graceful degradation of the guarantees. We leave for future work a comprehensive treatment of input orders that are neither random nor adversarial. We expect the techniques of the current work to serve as a cornerstone in such a treatment.

Paper structure. We discuss related work in Section 2. The formal setting and notations are defined in Section 3. The algorithm and its guarantees are given in Section 4. In Section 5, we outline the proof of the main result. Some parts of the proof are deferred to appendices, which can be found in the supplementary material. We conclude with a discussion in Section 6.

2 Related Work

Several works have studied settings related to the no-substitution k -median clustering setting. Table 1 summarizes the upper bounds mentioned below. First, some works studied related settings under an adversarial arrival order. Liberty et al. [2016] studied online k -means clustering, in which centers are sequentially selected, and each observed point is allocated to one of the previously selected centers. The proposed algorithm can be applied to the no-substitution setting, yielding an approximation factor of $O(\log(n))$, where n is the stream length. For this algorithm to select a sublinear number of centers, the aspect ratio of the input data must be bounded. Bhaskara and Rwanpathirana [2020] studied the same setting as Liberty et al. [2016], improving the approximation factor to a constant, under

Table 1: Comparing our guarantees to previous works. n is the stream length. Abbreviations: LSS16: Liberty et al. [2016], BR20: Bhaskara and Rwanpathirana [2020], BM20: Bhattacharjee and Moshkovitz [2021] Mo21: Moshkovitz [2021], HS20: Hess and Sabato [2020].

Reference	arrival order	assumptions	approximation factor	number of centers
LSS16	adversarial	bounded aspect ratio	$O(\log(n))$	$O(k \log^2(n))$
BR20	adversarial	bounded aspect ratio	constant	$O(k \log^2(n))$
BM20	adversarial	data properties	$O(k^3)$	$O(k \log(k) \log(n))$
Mo21	random	none	exponential in k	$O(k^5)$ or $O(k \log(n))$
HS20	random	bounded diameter	constant+additive	k
This work	random	none	constant	$O(k \log^2(k))$

the same assumptions. Bhattacharjee and Moshkovitz [2021] explicitly studied the no-substitution setting, and provided an approximation factor of $O(k^3)$, under a different assumption on the input data set.

The setting of a random arrival order has been studied in several recent works. Hess and Sabato [2020] proposed an algorithm that selects exactly k centers. They obtained a constant approximation factor, with an additional additive term that vanishes for large streams, under the assumption that the metric space has a bounded diameter. Their guarantee is with high probability. Moshkovitz [2021] obtained an approximation factor that is exponential in k , with a number of centers that is linear in k and logarithmic in the stream length n . Assuming a known stream length, the same work also obtained an approximation factor which is exponential in k while selecting $O(k^5)$ centers. Both guarantees hold only with a constant probability. As can be seen in Table 1, to date, the only known constant approximation algorithms for no-substitution k -median clustering impose structural assumptions on the input data.

A related sequential clustering setting is the streaming setting [e.g., Guha et al., 2000, Ailon et al., 2009, Chen, 2009, Ackermann et al., 2012, Braverman et al., 2016], in which the main restriction is the amount of memory available to the algorithm. This setting allows substituting selected centers, but algorithms in this setting can be used in the no-substitution setting, by collecting all the centers ever selected. However, we are not aware of any algorithm in this setting with a competitive bound on the total number of selected centers.

3 Setting and Notation

For an integer i , denote $[i] := \{1, \dots, i\}$. Let (X, ρ) be a finite metric space, where X is a set of size n and $\rho : X \times X \rightarrow \mathbb{R}_+$ is a metric. For a point $x \in X$ and a set $T \subseteq X$, we denote $\rho(x, T) := \min_{y \in T} \rho(x, y)$. For an integer $k \geq 2$, a k -clustering of X is a set of (at most) k points from X which represent cluster centers. Throughout this work, whenever an item is selected based on minimizing ρ , we assume that ties are broken based on a fixed arbitrary ordering. Given a set $S \subseteq X$, the k -median cost of T on S is $\text{cost}(S, T) := \sum_{x \in S} \rho(x, T)$. The k -median clustering problem aims to select a k -clustering T of X with a minimal overall cost $\text{cost}(X, T)$. We denote by OPT an optimal solution to this problem: $\text{OPT} \in \text{argmin}_{T \subseteq X, |T| \leq k} \text{cost}(X, T)$. In the *sequential no-substitution* k -median setting that we study, X is not known a priori. The points from X are presented to the algorithm sequentially, in a random order. We assume that the stream length, n , is provided as input to the algorithm; see Section 6 for a discussion on supporting an unknown stream length. The algorithm may select an observed point as a center only before observing the next point. Any selected point cannot later be removed or substituted. The goal of the algorithm is to select a small number of centers, such that with a high probability, the overall cost of the selected set on the entire X approximates the optimal k -median cost, $\text{cost}(X, \text{OPT})$. In other words, we seek an algorithm with a low competitive ratio, in the sense of competitive analysis [Borodin and El-Yaniv, 2005]. An *offline k -median algorithm* \mathcal{A} is an algorithm that takes as input a finite set

Algorithm 1 The sequential clustering algorithm: MUNSC

input $\delta \in (0, 1)$, $k \in \mathbb{N}$, $n \in \mathbb{N}$ (stream length), \mathcal{A} (an offline k -median algorithm), sequential access to the input stream.

- 1: $\alpha_1 \leftarrow \delta/(4k)$; For $i \in \mathbb{N}$, $\alpha_i \leftarrow \alpha_1 \cdot 2^{i-1}$.
- 2: $I \leftarrow \lfloor \log_2(1/(6\alpha_1)) \rfloor$, $\delta' \leftarrow \delta/(I+1)$. $\triangle I$ is set so that $\alpha_{I+1} \in (1/12, 1/6]$.
- 3: Prepare $I+1$ copies of `SelectProc`, indexed by $i \in [I+1]$, as follows:
 - In all copies, use inputs $k, n, \mathcal{A}, \tau \leftarrow \phi_{\alpha_{I+1}}$, and set the confidence parameter to δ' .
 - In copy $i \in [I+1]$, set $\alpha \leftarrow \alpha_i$.
 - In all but the last copy, set $M \leftarrow 1$ and $\gamma \leftarrow 2\alpha_i$.
 - In the last copy (index $I+1$), set $M \leftarrow \log(8k_+/\delta')$, and $\gamma \leftarrow 1 - 2\alpha_{I+1}$.
- 4: Read each point from the input stream and feed it to each of the copies of `SelectProc`. If any of the copies selected the point, then select it as a center.

of points S and the parameter k , and outputs a k -clustering of S , denoted $\mathcal{A}(S, k)$. For $\beta \geq 1$, we say that \mathcal{A} is a β -approximation offline k -median algorithm on (X, ρ) , if for all input sets $S \subseteq X$, $\text{cost}(S, \mathcal{A}(S, k)) \leq \beta \cdot \text{cost}(S, \text{OPT}_S)$, where OPT_S is an optimal solution with centers from S . Formally, $\text{OPT}_S \in \text{argmin}_{T \subseteq S, |T| \leq k} \text{cost}(S, T)$.

4 The algorithm

In this section, we describe the proposed algorithm, MUNSC. MUNSC receives as input the value of k , the confidence level $\delta \in (0, 1)$ and the total stream length n . We further assume access to some black-box offline k -median algorithm \mathcal{A} . The guarantees of MUNSC depend on the approximation factor guaranteed by \mathcal{A} . MUNSC has only sequential access to the input stream. Our main result is the following theorem, showing that MUNSC obtains a constant approximation factor with a quasi-linear number of centers. A proof sketch of the theorem is given in Section 5. The full proof is provided in the supplementary.

Theorem 4.1. *Let $k, n \in \mathbb{N}$. Let $\delta \in (0, 1)$. Let (X, ρ) be a metric space of size n . Let T_{out} be the set of centers selected by MUNSC for the input parameters δ, k, n . Suppose that the input stream is a random permutation of X , and that the input black-box algorithm \mathcal{A} is a β -approximation offline k -median algorithm. Then, with a probability at least $1 - \delta$,*

1. $|T_{\text{out}}| = O(k \log^2(k/\delta))$, and
2. $\text{cost}(X, T_{\text{out}}) \leq C\beta \cdot \text{cost}(X, \text{OPT})$, where $C > 0$ is a universal constant.

We first give, in Section 4.1, a short overview of the structure of MUNSC. A core element of the algorithm is the internal procedure `SelectProc`, which is used by MUNSC with several different sets of input parameters, creating a multiscale effect that we discuss below. We explain the details of `SelectProc` in Section 4.2. Then, in Section 4.3, we explain the design of MUNSC.

4.1 Algorithm structure overview

MUNSC is listed in Alg. 1. It works as follows: It initializes $\Theta(\log(k/\delta))$ copies of the procedure `SelectProc`, each with a different set of input parameters, where we define $\phi_\alpha := 150 \log(32k/\delta)/\alpha$ for $\alpha > 0$. It then iteratively reads the points from the input stream one by one, and feeds each point to each of the copies of `SelectProc`, so that each copy sequentially observes the entire input stream. Whenever any copy of `SelectProc` selects a point as a center, this point is selected by MUNSC as a center.

The internal procedure `SelectProc` uses a robust cost estimation approach to select a small number of centers. The procedure gets as input several technical parameters, in addition to the input parameters of MUNSC. We explain the procedure and the meaning of the technical parameters in Section 4.2. The values of the technical parameters set by MUNSC for each of the copies of `SelectProc` ensure that MUNSC select a small number of centers, while obtaining a constant approximation factor. This is done using a multiscale approach. We discuss this in detail in Section 4.3.

Algorithm 2 Internal procedure: SelectProc

input $\delta \in (0, 1)$, $k \in \mathbb{N}$, $n \in \mathbb{N}$, \mathcal{A} (an offline k -median algorithm),
sequential access to the input stream.
Technical parameters: $\alpha \in (0, \frac{1}{6}]$, $\gamma \in (\alpha, 1 - 2\alpha]$, $M \in \mathbb{N}$, $\tau > 0$.

- 1: Δ **Phase 1** (Calculate an initial clustering T_α):
- 2: $P_1 \leftarrow$ the first αn points from the input stream.
- 3: $T_\alpha := \{c_1, \dots, c_{k_+}\} \leftarrow \mathcal{A}(P_1, k_+)$ Δ Only calculation, no centers are actually selected here.
- 4: Δ **Phase 2** (Estimate the cost of T_α on large optimal clusters):
- 5: $P_2 \leftarrow$ next αn points from the input stream.
- 6: $\psi \leftarrow \frac{1}{3\alpha} \text{cost}_{2\alpha(k+1)\phi_\alpha}(P_2, T_\alpha)$ Δ ψ estimates the cost of T_α on large optimal clusters.
- 7: Δ **Phase 3** (Select centers):
- 8: $\forall i \in [k_+]$, $n_i \leftarrow 0$, $\text{Near}_i \leftarrow \text{FALSE}$. Δ n_i counts selected points associated with c_i .
 Near_i indicates if a point close to c_i was selected.
- 9: **for** γn iterations **do**
- 10: Read the next point x from the input stream.
- 11: $i \leftarrow \text{argmin}_{i \in [k_+]} \rho(x, c_i)$. Δ Find the closet center to x in T_α .
- 12: **if** $\rho(x, c_i) > \frac{\psi}{k\tau}$ or $\neg \text{Near}_i$ or $n_i \leq M$ **then**
- 13: **Select** x as a center. Δ This is the only line in which actual selections occur.
- 14: $n_i \leftarrow n_i + 1$.
- 15: **If** $\rho(x, c_i) \leq \frac{\psi}{k\tau}$ **then** $\text{Near}_i \leftarrow \text{TRUE}$.
- 16: **end if**
- 17: **end for**

4.2 The internal procedure: SelectProc

In this section, we present the internal procedure, SelectProc. It has been observed in previous work [Liberty et al., 2016] that knowing in advance the optimal cost on the entire stream allows a successful center selection. Under a random arrival order, one may hope that obtaining a good estimate for the optimal cost would be straight-forward. However, since the metric space is unbounded, even a small number of outliers can bias the cost estimate considerably. We overcome this challenge by showing that it suffices to estimate a version of the optimal cost that ignores outliers, and that this version can be estimated reliably from a small prefix of the stream. Our analysis is based on a distinction between small and large optimal clusters, which we formally define in Section 5 below.

SelectProc, listed in Alg. 2, uses the following notation. For two sets $S, T \subseteq X$ and an integer r , denote by $\text{far}_r(S, T) \subseteq S$ the set of r points in S that are the furthest from T according to the metric ρ . If $|S| < r$, we define $\text{far}_r(S, T) := S$ and call $\text{far}_r(S, T)$ a *trivial far set*. Denote by $\text{cost}_r(S, T) := \text{cost}(S \setminus \text{far}_r(S, T), T)$ the cost of T on S after discounting the r points that incur the most cost. Let $k_+ := k + 38 \log(32k/\delta)$. SelectProc receives the same input parameters as MUNSC. In addition, it requires several technical parameters, denoted α, γ, M and τ . We explain the meaning of these parameters in the proper context below.

SelectProc works in three phases. For $i \in [3]$, we denote by P_i the set of points that are read in phase i . In each of the first two phases, an α fraction of the points in the stream is read. In the last phase, a γ fraction of the points in the stream is read. Note that depending on the values of α and γ , the procedure may ignore a suffix of the stream. The first two phases are used for calculations. Centers are selected only during the third phase.

In the first phase, a reference clustering T_α is calculated using the input offline algorithm \mathcal{A} . Note that the centers in T_α cannot be selected as centers by SelectProc, since \mathcal{A} calculates T_α only after observing all the points in the phase. In the second phase, an estimate ψ of the cost of T_α is calculated. In the third phase, SelectProc observes points from the stream one by one, deciding for each one whether it should be selected as a center. For each observed point, SelectProc first finds the center $c_i \in T_\alpha$ which is closest to it. A point is then selected as a center if it satisfies one of three conditions: (1) its distance from c_i is more than $\psi/(k\tau)$ (where τ is one of the input technical parameters); (2) c_i does not yet have M associated points (where M is one of the technical parameters); or (3) no point close to c_i has been selected so far, as maintained by the Boolean variable Near_i . In Theorem 5.1 in Section 5, we provide a guarantee on the output of SelectProc, which upper bounds the number of

centers selected by `SelectProc` and the cost obtained by these centers, as a function of the technical parameters.

Challenges and solutions in `SelectProc`. We give here an informal explanation of the workings of `SelectProc`. The full analysis is provided in Section 5 and the supplementary. Consider the clusters induced by some fixed optimal k -median clustering on X . Call these clusters *optimal clusters*. Optimal clusters that are large (that is, include a sufficiently large fraction of X) are easy to identify from a small random subset of points. Therefore, approximate centers for these clusters will be identified in the first phase of `SelectProc` by the black box algorithm \mathcal{A} , and will be included in T_α . Thus, the cost of T_α on the subset of points that belong to large optimal clusters will be close to optimal. Note that T_α is a clustering with $k_+ > k$ centers. This overcomes the fact that the data set might include outliers which are not proportionally represented in each phase. Searching for a k -clustering in P_1 might thus lead to a clustering that is too biased towards such outliers. By increasing the size of the solution searched in the first phase to k_+ , this allows the solution to include the outliers as centers, while still choosing also centers that are close to the optimal centers of all large clusters. The selection conditions of the third phase further guarantee that at least M points are selected near each center $c_i \in T_\alpha$ which represents a large cluster, thus making sure that a center which is very close to each such c_i is selected. This allows bounding the obtained cost on points in large optimal clusters.

An additional challenge is to ensure that points in small optimal clusters are not too far from the set T_{out} of all selected centers. Since the metric space is unbounded, even a single such point can destroy the constant approximation factor. To overcome this, `SelectProc` selects a point near each center in T_α , as well as all the points that are far from T_α . The threshold $\psi/(k\tau)$, which defines a point as far, is set so that the number of points selected from large optimal clusters can be bounded. This bound crucially relies on the accuracy of ψ as an estimate for the cost of T_α on large clusters. The required accuracy is obtained by ignoring the points that are furthest from T_α when calculating the cost on P_2 (see line 6). For small clusters, they hold a small number of points by definition, thus the number of such points selected as centers is also bounded.

4.3 The clustering algorithm: MUNSC

As described above, MUNSC runs several copies of `SelectProc` on the same input stream: each point is read from the input stream and then fed to each of the copies of `SelectProc`. Each of these copies selects some centers, and the set of all selected centers is the solution of MUNSC. The difference between the copies is in the value of the technical input parameters. First, the value of α (the stream fraction for the first and second phases) is progressively doubled, starting with $\alpha_1 = \delta/(4k)$ and ending with $\alpha_{I+1} = \alpha_1 \cdot 2^I$, where I is set so that $\alpha_{I+1} \in (1/12, 1/6]$. The value of τ (which controls the selection threshold) is set to $\phi_{\alpha_{I+1}}$ in all copies, based on the largest value of α . In all but the last copy, γ (the stream fraction for the third phase) is set to $2\alpha_i$, where i is the copy index. This means that the first and second phase of copy i are each of size α_i , while the third phase is of size $2\alpha_i$ (the rest of the stream is ignored by copy i). Therefore, copy i observes a $4\alpha_i = 2\alpha_{i+1}$ fraction of the stream. Thus, the first two phases of copy $i + 1$ exactly overlap with the fraction of the stream observed by copy i . Figure 1 illustrates the overlap of phases in consecutive copies of `SelectProc`. In the last copy, indexed by $I + 1$, γ is set to $1 - 2\alpha_{I+1}$, thus this copy reads the entire stream. It can be seen that each point, except for the first $2n\alpha_1$ points, participates in the third phase of exactly one copy of `SelectProc`. As our analysis below shows, this overlap between the phases guarantees that the set of centers selected by the copies of `SelectProc` obtains a small cost on all the small optimal clusters. The last copy is slightly different: in addition to setting the length of the third phase γ to include all the points of the stream, it also sets the technical parameter M to a number larger than 1. We show below that in this way, the selected set of centers obtains a small cost on large optimal clusters as well.

The use of several copies of `SelectProc` with different phase sizes ensures that MUNSC selects only a quasi-linear number of centers, by overcoming the following challenge: On the one hand, the first phase and the second phase must be small, otherwise we might miss a point that needs to be selected, since no points are selected in these two phases. On the other hand, small first and second phases lead to a poor quality of the reference clustering T_α . The multiscale approach makes sure that almost all points participate in some selection phase, while at the same time improving the quality of T_α as




Figure 1: Illustrating the phases in copies 1, 2, 3 of `SelectProc` within `MUNSC`.

α grows. Our analysis below shows that in this way, the number of centers selected by each copy remains similar, even though larger values of α lead to a larger selection phase.

Using `MUNSC` in a bounded memory setting. `MUNSC` can be used in a bounded memory setting, by using a 1-pass streaming clustering algorithm as the black-box algorithm \mathcal{A} . For instance, the algorithm of Charikar et al. [2003] requires $O(k \cdot \text{polylog}(n))$ memory. `MUNSC` can be implemented to require additional memory of only $O(k \log(k/\delta))$ on top of the requirements of the black-box algorithm. We now explain in full how this can be done. First, consider the memory requirements of `SelectProc`: In phase 1, the points can be fed to \mathcal{A} one by one and do not need to be saved in memory. Only T_α , the $k_+ = O(k + \log(\frac{1}{\delta}))$ output points of \mathcal{A} , need to be saved in memory. In phase 2, calculating the truncated cost of T_α on P_2 can be done by saving at most $2\alpha \cdot (k+1)\phi_\alpha = O(k \log(\frac{k}{\delta}))$ at any time, so as to identify the farthest points from T_α , which are the ones to disregard. In phase 3, the points are inspected one by one and do not need to be saved in memory (if one wishes to save the selected centers in memory, although this is not in principle required in our setting, then an additional $O(k \log(\frac{k}{\delta}))$ factor is added). In total, `SelectProc` can be implemented using a memory size of $O(k \log(\frac{k}{\delta}))$ + the memory requirements of \mathcal{A} . `MUNSC` runs $O(\log(\frac{k}{\delta}))$ parallel runs of `SelectProc`. Thus, it requires at most a memory size of $O(\log(\frac{k}{\delta})(k \log(\frac{k}{\delta}) + \text{the memory requirements of } \mathcal{A}))$. In fact, this can be made even more memory-efficient, by using an algorithm \mathcal{A} which is a streaming anytime algorithm. In this case, it is possible to use a single copy of \mathcal{A} for all the parallel runs of `SelectProc`, since P_1 in the i 'th copy of `SelectProc` is the prefix of P_1 in the $i+1$ 'th copy (see Figure 1). That is, \mathcal{A} is run on the entire stream and whenever a phase 1 of any copy ends, its current set of centers is retrieved for this copy. In this case, the memory size required by `MUNSC` is $O(k \log^2(\frac{k}{\delta}))$ + the memory requirements of \mathcal{A} .

5 Analysis

In this section, we give an overview of the proof of Theorem 4.1. The supplementary provides the full proof. Denote the centers in the optimal solution `OPT` by $\{c_i^*\}_{i \in [k]}$, and the clusters induced by `OPT` (the *optimal clusters*) by $\{C_i^*\}_{i \in [k]}$. Formally, $C_i^* := \{x \in X \mid i = \text{argmin}_{j \in [k]} \rho(c_j^*, x)\}$. Optimal clusters with fewer than ϕ_α points are α -small optimal clusters, and the complement are α -large. Denote the indices of α -small and α -large clusters by $I_{\text{small}}^\alpha := \{i \in [k] \mid |C_i^*| < \phi_\alpha\}$ and $I_{\text{large}}^\alpha := [k] \setminus I_{\text{small}}^\alpha$, and the points in these clusters by $C_{\text{small}}^\alpha := \bigcup_{i \in I_{\text{small}}^\alpha} C_i^*$ and $C_{\text{large}}^\alpha := \bigcup_{i \in I_{\text{large}}^\alpha} C_i^*$. To prove Theorem 4.1, we provide the following guarantee on the centers selected by `SelectProc`.

Theorem 5.1. *Let $k, n \in \mathbb{N}$. Let $\delta \in (0, 1)$. Let (X, ρ) be a metric space of size n . Let T_{out} be the set of centers selected by `SelectProc` for the input parameters $\delta, k, n, \alpha \in (0, 1/6]$, $\gamma \in (\alpha, 1 - 2\alpha]$, $M \in \mathbb{N}$, $\tau > 0$. Suppose that the input stream is a random permutation of X , and that \mathcal{A} is a β -approximation offline k -median algorithm. Then, with a probability at least $1 - \delta/2$,*

1. $|T_{\text{out}}| = O(\frac{2}{\alpha} k \log(k/\delta) + k\tau + (k + \log(k/\delta))M)$;
2. For any $i \in [k]$, if $c_i^* \in P_3$, then

$$\text{cost}(C_i^*, T_{\text{out}}) \leq \text{cost}(C_i^*, \{c_i^*\}) + \frac{|C_i^*|}{k\tau} (36\beta + 20) \text{cost}(X, \text{OPT}); \quad (1)$$

3. If $\gamma = 1 - 2\alpha$, $\tau = \phi_\alpha$ and $M = \log(8k_+/\delta)$, then

$$\text{cost}(C_{\text{large}}^\alpha, T_{\text{out}}) \leq \text{cost}(C_{\text{large}}^\alpha, \text{OPT}) + (468\beta + 260) \text{cost}(X, \text{OPT}). \quad (2)$$

The proof of Theorem 4.1, provided in Appendix A, uses part 1 to bound the total number of centers, and parts 2 and 3 to bound the cost of small and large optimal clusters, respectively. The statement of Theorem 5.1 hints to the reason for the multiscale design of MUNSC: On the one hand, part 2 gives an approximation upper bound only if the cluster center appears in the third phase. For this to hold with a high probability, the third phase needs to be very large. On the other hand, for a quasi-linear bound on selected centers in part 1, the ratio between the third phase and the first phases (γ/α) needs to be small. The multiscale design overcomes this by keeping the ratio constant in each copy of `SelectProc`, while ensuring that the union of all third phases is almost the entire stream.

In the rest of the section, we give the proof of Theorem 5.1: In Section 5.1, we introduce the concept of *linear bin divisions*, and derive its properties. In Section 5.2, we give an overview of the proof of Theorem 5.1, using the results of Section 5.1. The full proof is provided in the supplementary.

5.1 Linear bin divisions

A main tool in the proof of Theorem 5.1 is partitioning sets of points from the input stream into subsets, such that each of the subsets is probably well represented in a relevant random subset of the stream. Meyerson et al. [2004] defined the concept of a *bin division* of X with respect to a clustering T , which is a partition of X into bins of equal size, where points are allocated to bins based on their distance from T . Here, we define the concept of a *linear bin division*, in which the bins linearly increase in size. This gradual increase allows keeping the size of the first bin independent of the stream length, while still proving that with a high probability, the overlap of each of the bins in the division with random subsets of the stream is close to expected. The fixed size of the first bin is crucial for deriving guarantees that are independent of the stream length. In addition, the ratio between adjacent bin sizes is kept bounded, which allows proving a bounded approximation factor.

Definition 5.2 (Linear bin division). *Let $W, T \subseteq X$ be finite sets, and $z \in \mathbb{N}$. A z -linear bin division of W with respect to T is a partition $\mathcal{B} = (\mathcal{B}(1), \dots, \mathcal{B}(L))$ of W (for an integer L) such that:*

1. *If $z \leq |W|$, then $\forall i \in [L], |\mathcal{B}(i)| \geq z \cdot (i+1)/2$. Otherwise, the bin division is called *trivial*, and defined as $\mathcal{B} := \mathcal{B}(1) := W$.*
2. *$|\mathcal{B}(1)| \leq \frac{5}{2}z$.*
3. *$\forall i \in [L-1], |\mathcal{B}(i+1)|/|\mathcal{B}(i)| \leq 3/2$.*
4. *$\forall i \in [L-1]$, and $\forall x \in \mathcal{B}(i), x' \in \mathcal{B}(i+1)$, it holds that $\rho(x, T) \geq \rho(x', T)$.*

A linear bin division exists for any size of W : For $|W| \leq z$, the conditions trivially hold. For $|W| \geq z$, the three first properties hold for the following allocation of bin sizes: Let L be the largest integer such that $B := \sum_{i \in [L]} z \cdot (i+1)/2 \leq |W|$. Set the size of $\mathcal{B}(i)$ to $z \cdot (i+1)/2 + (|W| - B)/L$. Property 1 clearly holds. Property 2 holds since $(|W| - B)/L \leq (z(L+2)/2)/L \leq 3z/2$. Property 3 holds since $(i+2+a)/(i+1+a) \leq 3/2$ for all $i \geq 1$ and any non-negative a . To satisfy property 4, allocate the elements of W into the bins in descending order of their distance from T .

We say that a set is *well-represented* in another set if the size of its overlap with the set is similar to expected. This extends naturally to bin divisions. Formally, this is defined as follows.

Definition 5.3 (Well-represented). *Let W be a finite set and $A, B \subseteq W$. We say that B is well-represented in A for W if $|B \cap A|/|B| \in [r/2, (3/2)r]$, where $r := |A|/|W|$. We say that a linear bin division \mathcal{B} of W is well represented in A for W if each bin in \mathcal{B} is well represented in A for W .*

The following lemma shows that if A is selected uniformly at random from W , then any fixed set B is well-represented in A with a high probability. Moreover, the same holds for any z -linear bin division of W with a sufficiently large z . The proof of the lemma is provided in Appendix B.

Lemma 5.4. *Let W be a finite set. Let $B \subseteq W$ be a set, and let \mathcal{B} be a z -linear bin division of some subset of W with respect to some T , for some integer z . Let $A \subseteq W$ be a set of size $r|W|$ selected uniformly at random from W . Then the following hold:*

1. *With a probability at least $1 - 2e^{-\frac{r|B|}{10}}$, B is well-represented in A for W .*
2. *If $z \geq 10 \log(4/\delta)/r$, then with a probability $1 - \delta$, \mathcal{B} is well represented in A for W .*

We now state and prove a main property of linear bin divisions, which will allow us to use sub-streams of the input stream to bound the cost of a clustering on X .

Lemma 5.5. *Let $W \subseteq X$ and let $z \in \mathbb{N}$. Let \mathcal{B} be a z -linear bin division of W with respect to some T . Let $A \subseteq W$ and $r \in (0, 1)$. If $\forall i \in [L]$, $|\mathcal{B}(i) \cap A|/|\mathcal{B}(i)| \leq r$, then $\text{cost}(A \setminus \mathcal{B}(1), T) \leq \frac{3}{2}r \text{cost}(W, T)$.*

Proof. Let L be the number of bins in \mathcal{B} . We first prove the following inequality, which relates the cost of the intersection of a bin with A to the cost of the preceding bin.

$$\forall i \in [L-1], \text{cost}(A \cap \mathcal{B}(i+1), T) \leq \frac{3}{2}r \cdot \text{cost}(\mathcal{B}(i), T). \quad (3)$$

To prove Eq. (3), fix $i \in [L-1]$, and denote $b := \max_{x \in \mathcal{B}(i+1)} \rho(x, T)$. By the assumptions, we have $|A \cap \mathcal{B}(i+1)| \leq r|\mathcal{B}(i+1)|$. Hence, $\text{cost}(A \cap \mathcal{B}(i+1), T) \leq r|\mathcal{B}(i+1)| \cdot b$. By property 3 of linear bin divisions, $|\mathcal{B}(i+1)| \leq \frac{3}{2}|\mathcal{B}(i)|$. Therefore, $\text{cost}(A \cap \mathcal{B}(i+1), T) \leq \frac{3}{2}r|\mathcal{B}(i)| \cdot b$. In addition, by property 4 of linear bin divisions and the definition of b , $\forall x \in \mathcal{B}(i)$, $b \leq \rho(x, T)$. Therefore, $\text{cost}(\mathcal{B}(i), T) \geq |\mathcal{B}(i)| \cdot b$. Combining the two inequalities, we get Eq. (3). It follows that:

$$\text{cost}(A \setminus \mathcal{B}(1), T) = \sum_{i=2}^L \text{cost}(A \cap \mathcal{B}(i), T) \leq \frac{3}{2}r \sum_{i=1}^{L-1} \text{cost}(\mathcal{B}(i), T) \leq \frac{3}{2}r \text{cost}(W, T).$$

This proves the statement of the lemma. \square

To prove Theorem 5.1, we define an event, denoted E , which holds with a high probability. We prove each of the claims in the theorem under the assumption that this event holds. To define the event, we first provide some necessary notation. Consider a run of `SelectProc` with some fixed set of input parameters, and assume that the input stream is a random permutation of the points in X . Recall that $T_\alpha = \{c_1, \dots, c_{k_+}\}$ is the clustering calculated in the first phase of `SelectProc`. Denote the clusters induced by T_α on X by C_1, \dots, C_{k_+} , where $C_i := \{x \in X \mid i = \text{argmin}_{j \in k_+} \rho(c_j, x)\}$. We define several sets and bin divisions, which will be used to define the required event. Let \mathcal{B}_a be a $(\phi_\alpha/15)$ -linear bin division of X with respect to `OPT`. Define $F_b := \text{far}_{k\phi_\alpha}(X \setminus P_1, T_\alpha)$, and let \mathcal{B}_b be a $(\phi_\alpha/3)$ -linear bin division of $X \setminus (P_1 \cup F_b)$ with respect to T_α . Define $F_c := \text{far}_{4(k+1)\phi_\alpha}(X \setminus P_1, T_\alpha)$, and let \mathcal{B}_c be a $(\phi_\alpha/3)$ -linear bin division of $X \setminus (P_1 \cup F_c)$ with respect to T_α . Lastly, define $F_d := \text{far}_{5(k+1)\phi_\alpha}(X \setminus P_1, T_\alpha)$. Note that each of these objects may be trivial if the stream is small.

The event E is defined as the following conjunction: (1) For each $i \in I_{\text{large}}^\alpha$, C_i^* is well-represented in P_1 and in $P_1 \cup P_2$ for X . (2) Each of F_b and F_c is trivial or well-represented in P_2 for $X \setminus P_1$; F_d is trivial or well-represented in P_3 for $X \setminus P_1$. (3) \mathcal{B}_a is trivial or well-represented in P_1 for X ; Each of \mathcal{B}_b and \mathcal{B}_c is trivial or well-represented in P_2 for $X \setminus P_1$, and (4) For each $i \in [k_+]$, one of the first $\log_2(8k_+/\delta)$ points observed from C_i in P_3 is closer to c_i than at least half of the points in $C_i \cap P_3$.

The following lemma, proved in Appendix C, shows that E holds with a high probability.

Lemma 5.6. *Consider a run of `SelectProc` with fixed input parameters. Assume that the input stream is a random permutation of the points of X . Then E holds with a probability at least $1 - \delta/2$.*

5.2 Theorem 5.1: proof overview

In this section, we give an overview of the proof of Theorem 5.1. The full proof is given in Appendix D. In the first phase of `SelectProc`, the clustering T_α is calculated using the offline algorithm \mathcal{A} . We first bound the cost of T_α on points in α -large optimal clusters.

Lemma 5.7. *If E holds, then $\text{cost}(C_{\text{large}}^\alpha, T_\alpha) \leq (18\beta + 10)\text{cost}(X, \text{OPT})$.*

In the second phase of `SelectProc`, the goal is to estimate a truncated version of the cost of T_α on X . This truncated version ignores the cost on the outliers, which are the $k\phi_\alpha$ furthest points from T_α in X . Setting the estimate to $\psi = \frac{1}{3\alpha} \text{cost}_{2\alpha(k+1)\phi_\alpha}(P_2, T_\alpha)$ ignores the furthest $2\alpha(k+1)\phi_\alpha$ points in P_2 , which with a high probability include all the $k\phi_\alpha$ outliers in X . The following lemma bounds ψ between two versions of a truncated cost of T_α .

Lemma 5.8. *If E holds, then $\frac{1}{9} \text{cost}_{5(k+1)\phi_\alpha}(X \setminus P_1, T_\alpha) \leq \psi \leq \text{cost}_{k\phi_\alpha}(X, T_\alpha)$.*

The proof of this lemma uses the linear bin divisions studied in Section 5.1. In particular, the upper bound uses Lemma 5.5, and the lower bound uses the specific construction of the bin sizes.

From the lemmas above, it can be seen that ψ is also upper bounded by the optimal cost, as follows. α -small optimal clusters are smaller than ϕ_α , thus their total size is less than $k\phi_\alpha$. It follows that

$$\text{cost}_{k\phi_\alpha}(X, T_\alpha) \leq \text{cost}(X \setminus C_{\text{small}}^\alpha, T_\alpha) = \text{cost}(C_{\text{large}}^\alpha, T_\alpha). \quad (4)$$

Combining Lemma 5.7, Eq. (4) and Lemma 5.8, we get that ψ is upper bounded by the optimal cost.

In the third phase, the points selected as centers include (among others) all the points with a distance of more than $\psi/(k\tau)$ from T_α . This allows bounding the cost of T_{out} on points in some of the optimal clusters by a linear expression in $\psi/(k\tau)$. Combined with the upper bound on ψ , this allows proving part 2 of Theorem 5.1. Part 3 bounds the cost of T_{out} on large optimal clusters, for specific settings of the technical parameters of `SELECTPROC`. To prove this part, we show that under these parameter settings, the centers selected based on the reference clustering T_α include sufficiently many points around each large optimal cluster. We then show that one of these points is close to the optimal center.

Lastly, part 1 of Theorem 5.1 upper bounds the number of centers selected by `SELECTPROC`. The following lemma upper-bounds the number of points that are selected because they are far from T_α .

Lemma 5.9. *Let $N := |\{x \in P_3 \mid \rho(x, T_\alpha) > \psi/(k\tau)\}|$. Under E , $N \leq 8(k+1)\phi_\alpha \frac{\gamma}{1-\alpha} + 9k\tau$.*

The proof Lemma 5.9 uses the lower bound on ψ shown in Lemma 5.8, to conclude that there cannot be too many points that are more than $\psi/(k\tau)$ far from T_α , since this would make the truncated cost in the lower bound larger than ψ . We note that the resulting upper bound is independent of the stream size, due to the specific construction of the bin sizes in the linear bin division. Using this lemma, we can prove part 1 of `SELECTPROC`. The full proof of Theorem 5.1 is given in Appendix D.

6 Discussion

In this work, we provided `MUNSC`, the first constant-approximation algorithm for sequential no-substitution k -median clustering that does not require structural assumptions on the input data. Moreover, the number of centers selected by `MUNSC` is only quasi-linear in k .

`MUNSC` can also be used when the stream length n is not known in advance, using a simple doubling trick. Moshkovitz [2021] showed that an unknown stream length with an approximation factor that does not depend on n require a logarithmic dependence on n in the number of selected centers. Indeed, a doubling trick for `MUNSC` would preserve the approximation factor, increasing the number of selected centers only by the unavoidable $\Theta(\log(n))$ factor.

Funding Transparency Statement

Sivan Sabato and Tom Hess were supported in part by the Israel Science Foundation (grant No. 555/15) and by the United-States-Israel Binational Science Foundation (BSF) (grant no. 2017641). Michal Moshkovitz has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 882396), by the Israel Science Foundation (grant number 993/17), Tel Aviv University Center for AI and Data Science (TAD), and the Yandex Initiative for Machine Learning at Tel Aviv University.

References

- Marcel R Ackermann, Marcus Märtens, Christoph Raupach, Kamil Swierkot, Christiane Lammersen, and Christian Sohler. StreamKM++: A clustering algorithm for data streams. *Journal of Experimental Algorithmics (JEA)*, 17:2–4, 2012.
- Charu C Aggarwal. A framework for diagnosing changes in evolving data streams. In *Proceedings of the 2003 ACM SIGMOD international conference on Management of data*, pages 575–586. ACM, 2003.
- Nir Ailon, Ragesh Jaiswal, and Claire Monteleoni. Streaming k-means approximation. In *Advances in neural information processing systems*, pages 10–18, 2009.

- Aditya Bhaskara and Aravinda Kanchana Rwanpathirana. Robust algorithms for online k -means clustering. In *Algorithmic Learning Theory*, pages 148–173, 2020.
- Robi Bhattacharjee and Michal Moshkovitz. No-substitution k -means clustering with adversarial order. In *Algorithmic Learning Theory*, pages 345–366. PMLR, 2021.
- Allan Borodin and Ran El-Yaniv. *Online computation and competitive analysis*. cambridge university press, 2005.
- Vladimir Braverman, Dan Feldman, and Harry Lang. New frameworks for offline and streaming coresets constructions. *arXiv preprint arXiv:1612.00889*, 2016.
- Moses Charikar, Sudipto Guha, Éva Tardos, and David B Shmoys. A constant-factor approximation algorithm for the k -median problem. *Journal of Computer and System Sciences*, 65(1):129–149, 2002.
- Moses Charikar, Liadan O’Callaghan, and Rina Panigrahy. Better streaming algorithms for clustering problems. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 30–39. ACM, 2003.
- Ke Chen. On coresets for k -median and k -means clustering in metric and euclidean spaces and their applications. *SIAM Journal on Computing*, 39(3):923–947, 2009.
- Devdatt P Dubhashi and Desh Ranjan. Balls and bins: A study in negative dependence. *BRICS Report Series*, 3(25), 1996.
- Sudipto Guha, Nina Mishra, Rajeev Motwani, and Liadan O’Callaghan. Clustering data streams. In *Foundations of computer science, 2000. proceedings. 41st annual symposium on*, pages 359–366. IEEE, 2000.
- Tom Hess and Sivan Sabato. Sequential no-substitution k -median-clustering. In *International Conference on Artificial Intelligence and Statistics*, pages 962–972. PMLR, 2020.
- Kumar Joag-Dev and Frank Proschan. Negative association of random variables with applications. *The Annals of Statistics*, pages 286–295, 1983.
- Kingsly Leung and Christopher Leckie. Unsupervised anomaly detection in network intrusion detection using clusters. In *Proceedings of the Twenty-eighth Australasian conference on Computer Science-Volume 38*, pages 333–342. Australian Computer Society, Inc., 2005.
- Shi Li and Ola Svensson. Approximating k -median via pseudo-approximation. *SIAM Journal on Computing*, 45(2):530–547, 2016.
- Edo Liberty, Ram Sriharsha, and Maxim Sviridenko. An algorithm for online k -means clustering. In *2016 Proceedings of the Eighteenth Workshop on Algorithm Engineering and Experiments (ALENEX)*, pages 81–89. SIAM, 2016.
- Adam Meyerson, Liadan O’callaghan, and Serge Plotkin. A k -median algorithm with running time independent of data size. *Machine Learning*, 56(1-3):61–87, 2004.
- Michal Moshkovitz. Unexpected effects of online no-substitution k -means clustering. In *Algorithmic Learning Theory*, pages 892–930. PMLR, 2021.
- Olfa Nasraoui, Jeff Cerwinski, Carlos Rojas, and Fabio Gonzalez. Performance of recommendation systems in dynamic streaming environments. In *Proceedings of the 2007 SIAM International Conference on Data Mining*, pages 569–574. SIAM, 2007.
- Andriy Shepitsen, Jonathan Gemmell, Bamshad Mobasher, and Robin Burke. Personalized recommendation in social tagging systems using hierarchical clustering. In *Proceedings of the 2008 ACM conference on Recommender systems*, pages 259–266. ACM, 2008.
- Bichen Zheng, Sang Won Yoon, and Sarah S Lam. Breast cancer diagnosis based on feature extraction using a hybrid of k -means and support vector machine algorithms. *Expert Systems with Applications*, 41(4):1476–1482, 2014.

A Proof of Theorem 4.1

In this section, we use Theorem 5.1 to prove Theorem 4.1. First, define the following event, which we denote by G :

$$\forall i \in [k], c_i^* \in X \text{ is not observed in the first } n \cdot \delta / (2k) \text{ reads from the input stream.}$$

Observe that G holds with a probability at least $1 - \delta/2$: For each c_i^* , the probability that it is observed in the prefix is $\delta/(2k)$, and a union bound over k centers gives a probability of at most $\delta/2$ that the event does not hold. We use this observation in the proofs below. First, we prove the guarantees of `SelectProc`.

Proof of Theorem 4.1. First, we show that we can assume that Theorem 5.1 holds simultaneously for all copies of `SelectProc` run in `MUNSC`. The number of copies of `SelectProc` is $I + 1$. Since Alg. 1 calls `SelectProc` with the confidence parameter δ' , each of the executions of `SelectProc` satisfies the claims of Theorem 5.1 with a probability at least $1 - \delta'/2$. Since $\delta' = \delta/(I + 1)$, by a union bound, Theorem 5.1 holds for all executions simultaneously with a probability at least $1 - \delta/2$. With a probability at least $1 - \delta$, this holds simultaneously with the event G defined above. We henceforth assume that all these events hold. Note that since $\alpha_1 = \delta/(4k)$, we have

$$I + 1 = \lfloor \log_2(1/(6\alpha_1)) \rfloor + 1 = O(\log(k/\delta)).$$

Therefore, $\delta' = \Omega(\delta/\log(k/\delta))$ and $\log(k/\delta') = O(\log(k/\delta))$.

To prove the first part of Theorem 4.1, which upper bounds $|T_{\text{out}}|$, we apply the first claim of Theorem 5.1 (with δ'), which states that $|T_{\text{out}}| = O(\frac{2}{\alpha}k \log(k/\delta') + k\tau + (k + \log(k/\delta'))M)$, to each of the copies of `SelectProc`. In all the copies of `SelectProc`, we have

$$\tau := \phi_{\alpha_{I+1}} = O(\log(k/\delta')) = O(\log(k/\delta)).$$

In addition, in all but the last copy of `SelectProc`, we have $M := 1$ and $\gamma/\alpha := 2$. Thus, the number of centers selected in each of these copies of `SelectProc` is $O(k \log(k/\delta))$. There are $I = O(\log(k/\delta))$ such copies, hence the total number of centers selected in all but the last copy is $O(k \log^2(k/\delta))$. In the last copy, `SelectProc` is called with $M = O(\log(k_+/\delta')) = O(\log(k/\delta))$ and $\gamma = 1 - 2\alpha_{I+1}$. Since $\alpha_{I+1} > 1/12$, we have $\gamma/\alpha_{I+1} = O(1)$. Therefore, again by the first claim of Theorem 5.1, the number of centers selected by the last copy is $O(k \log(k/\delta) + \log^2(k/\delta))$. The overall number of centers selected in `MUNSC` is thus $O(k \log^2(k/\delta))$.

We now prove the second part of Theorem 4.1, which bounds the approximation factor of T_{out} on X . Since G holds, none of the centers c_i^* , for $i \in [k]$, appears in the first $2\alpha_1 n$ points read from the input stream. Now, each point observed after this prefix appears in the selection phase of some copy of `SelectProc` in `MUNSC`. It follows that for each $i \in [k]$, Eq. (1) in the second part of Theorem 5.1 holds for one of the copies.

We apply Eq. (1) noting that in all copies, we have $\tau := \phi_{\alpha_{I+1}}$. In addition, all copies get the full stream X as input. Moreover, the set of centers T_{out} selected by `MUNSC` is a superset of each of the sets selected in each of the calls to `SelectProc`, thus its cost cannot be larger. It follows that for all $i \in [k]$,

$$\text{cost}(C_i^*, T_{\text{out}}) \leq \text{cost}(C_i^*, \{c_i^*\}) + \frac{|C_i^*|}{k\phi_{\alpha_{I+1}}} (36\beta + 20) \text{cost}(X, \text{OPT}).$$

In particular, for $i \in I_{\text{small}}^{\alpha_{I+1}}$, we have $|C_i^*| \leq \phi_{\alpha_{I+1}}$. Therefore,

$$\text{cost}(C_{\text{small}}^{\alpha_{I+1}}, T_{\text{out}}) = \sum_{i \in C_{\text{small}}^{\alpha_{I+1}}} \text{cost}(C_i^*, T_{\text{out}}) \leq \text{cost}(C_{\text{small}}^{\alpha_{I+1}}, \text{OPT}) + (36\beta + 20) \text{cost}(X, \text{OPT}). \quad (5)$$

Now, consider the last copy of `SelectProc`. In this copy, the conditions of the third claim of Theorem 5.1 hold, and so Eq. (2) holds. Combining Eq. (2) with Eq. (5), and noting that $X = C_{\text{small}}^{\alpha_{I+1}} \cup C_{\text{large}}^{\alpha_{I+1}}$, the second part of Theorem 4.1 follows. \square

B Proof of Lemma 5.4

In this section, we prove Lemma 5.4, stated in Section 5.1, which provides guarantees for the property of being well-represented.

Proof of Lemma 5.4. For the first claim, we use a multiplicative concentration bound of Dubhashi and Ranjan [1996]. This bound states that for negatively associated random variables Z_1, \dots, Z_n taking values in $\{0, 1\}$, where $Z := \sum_{i \in [n]} Z_i$, it holds that

$$\mathbb{P}[|Z - \mathbb{E}[Z]| \geq \epsilon \mathbb{E}[Z]] \leq 2e^{-\frac{\epsilon^2 \mathbb{E}[Z]}{2+\epsilon}}.$$

Let Z_i be equal to 1 if the i 'th element of B is in A . By Joag-Dev and Proschan [1983], the random variables of a uniform sample without replacement are negatively associated, hence the inequality above holds for $\{Z_i\}$. In addition, $Z = |B \cap A|$ and $\mathbb{E}[Z] = r|B|$. The first claim follows by setting $\epsilon := 1/2$.

To prove the second claim, note that the first claim implies that each set $\mathcal{B}(i)$ in \mathcal{B} is well-represented in A for W with a probability at least $1 - 2 \exp(-r|\mathcal{B}(i)|/10)$. By property 1 of linear bin divisions, and the assumption that $z \geq 10 \log(4/\delta)/r$, we have

$$|\mathcal{B}(i)| \geq z(i+1)/2 \geq 5(i+1) \log(4/\delta)/r.$$

Therefore, $\exp(-r|\mathcal{B}(i)|/10) \leq (\delta/4)^{(i+1)/2}$.

By a union bound over all the bins in \mathcal{B} , the probability that at least one bin in \mathcal{B} is not well-represented in A for W is upper-bounded by $2 \sum_{i=1}^{\infty} (\delta/4)^{(i+1)/2} \leq \frac{\delta}{2-\sqrt{\delta}} \leq \delta$, which proves the claim. \square

C Proof of Lemma 5.6

In this section, we prove Lemma 5.6, which states that the event E holds with high probability.

Proof of Lemma 5.6. First, we show that part 1 of E holds with a probability at least $1 - \delta/32$. This is proved by applying the first part of Lemma 5.4 to each of the clusters C_i^* for $i \in I_{\text{large}}^\alpha$, and noting that their size is at least ϕ_α . Applying the lemma with the sets P_1, X and with the sets $P_1 \cup P_2, X$ and using a union bound over at most k large optimal clusters, it follows that the probability that part 1 does not hold is at most $4k \exp(-\alpha\phi_\alpha/10) = 4k \exp(15 \log(\delta/(32k))) = 4k \cdot (\delta/(32k))^{15} \leq \delta/32$.

For part 2, we also show that it holds with a probability at least $1 - \delta/32$. We similarly apply the first part of Lemma 5.4 to each of the required sets using the ratios $|P_2|/|X \setminus P_1| = \alpha/(1-\alpha) \geq \alpha$ and $|P_3|/|X \setminus P_1| = \gamma/(1-\alpha) \geq \alpha$ (since $\gamma \geq \alpha$). Since F_b, F_c, F_d , if not trivial, are each of size at least $k\phi_\alpha$, we get that the probability that any of these sets is not well represented is at most $6 \exp(-\alpha\phi_\alpha/10) \leq 6 \exp(15 \log(\delta/(32k))) \leq \delta/32$.

Next, we show that part 3 holds with a probability at least $1 - \delta/4$. For \mathcal{B}_a , if it is not trivial, then the conditions of the second part of Lemma 5.4 hold with confidence parameter $\delta/8$, since it is a z -linear bin division with $z = \phi_\alpha/15 \geq 10 \log(32/\delta)/\alpha$, and $\alpha = r$, where r is the ratio used in Lemma 5.4. For \mathcal{B}_b and \mathcal{B}_c , if they are not trivial, the conditions of the second part of Lemma 5.4 with confidence parameter $\delta/16$ hold, since they are z -linear bin divisions with $z = \phi_\alpha/3 = 50 \log(32k/\delta)/\alpha \geq 10 \log(4 \cdot 16/\delta)/r$.

Lastly, we show that part 4 holds with a probability at least $1 - \delta/8$. Fix $i \in [k_+]$. If $|C_i \cap P_3| \leq \log_2(8k_+/\delta)$, then the statement of part 4 trivially holds. Assume that $|C_i \cap P_3| > \log_2(8k_+/\delta)$. Points from $C_i \cap P_3$ are observed in a random order in the stream. The fraction of points in $C_1 \cap P_3$ that are closer to c_i than at least half of the points in $C_1 \cap P_3$ is at least half. Therefore, every draw from $C_i \cap P_3$ has a probability of at least half of satisfying this requirement, conditioned on all previous points not satisfying this requirement. The probability that none of the points out of the first $\log_2(8k_+/\delta)$ satisfies this requirement is thus at most $2^{-\lceil \log_2(8k_+/\delta) \rceil} \leq \delta/(8k_+)$. A union bound over $i \in [k_+]$ gives the desired bound.

By a union bound over all the parts of the event, it hold with a probability at least $1 - \delta/2$. \square

D Proof of Theorem 5.1

In this section, we prove Theorem 5.1. First, we state two helpful lemmas. Then, we give the proof of the theorem. The proofs of the lemmas below and of the lemmas stated in Section 5.2, which are all used in the proof of the theorem, are provided in the subsections below.

The following two lemmas allow us to bound the cost of the set of selected centers T_{out} on points in all optimal clusters. The first lemma considers optimal clusters whose center is observed in the third phase.

Lemma D.1. *For any $i \in [k]$, if $c_i^* \in P_3$ then $\text{cost}(C_i^*, T_{\text{out}}) \leq \text{cost}(C_i^*, \{c_i^*\}) + 2|C_i^*|\psi/(k\tau)$.*

The second lemma bounds the cost of the set of selected centers on large optimal clusters.

Lemma D.2. *If E holds, $\gamma = 1 - 2\alpha$, $\tau = \phi_\alpha$ and $M = \log(8k_+/\delta)$, then*

$$\text{cost}(C_{\text{large}}^\alpha, T_{\text{out}}) \leq \text{cost}(C_{\text{large}}^\alpha, \text{OPT}) + 26\text{cost}(C_{\text{large}}^\alpha, T_\alpha).$$

Based on the lemmas above and the lemmas stated in Section 5.2, the proof of Theorem 5.1 can be given.

Proof of Theorem 5.1. By Lemma 5.6, E holds with probability $1 - \delta/2$. Assume that E holds. First, we upper bound $|T_{\text{out}}|$. Note that `SelectProc` selects a point only if one of the three conditions in line 12 hold.

By Lemma 5.9, the number of points selected by `SelectProc` due to the first condition is

$$N \leq 8(k+1)\phi_\alpha \frac{\gamma}{1-\alpha} + 9k\tau.$$

Since $\phi_\alpha = O(\log(k/\delta)/\alpha)$ and $\alpha \leq 1/6$, we have $N = O((\frac{\gamma}{\alpha}k \log(k/\delta) + k\tau)$. The rest of the points selected by `SelectProc` are those satisfying one of the other conditions of line 12. This allows at most M additional points for each $i \in [k_+]$. Thus, the total number of points selected by `SelectProc` is $O(\frac{\gamma}{\alpha}k \log(k/\delta) + k\tau + k_+ \cdot M)$. Since by definition, $k_+ = O(k + \log(k/\delta))$, part 1 of Theorem 5.1 immediately follows.

Part 2 of Theorem 5.1 follows immediately, by combining Lemma D.1, Eq. (4), Lemma 5.8 and Lemma 5.7. Part 3 is also immediate, by combining Lemma 5.7 and Lemma D.2. This completes the proof of the theorem. \square

In the next subsections, we prove the lemmas used in the proof of Theorem 5.1.

D.1 Bounding the cost of T_α on large clusters

In this section, we prove Lemma 5.7, which bounds $\text{cost}(C_{\text{large}}^\alpha, T_\alpha)$.

For a given $x \in X$, denote the center closest to it in T_α by $c_\alpha(x)$. Recall that \mathcal{B}_a is a $(\phi_\alpha/15)$ -linear bin division of X with respect to OPT . To prove Lemma 5.7, we first bound the cost of T_α on large clusters, using an expression that depends on the optimal cost and on the points in the first phase of `SelectProc`.

Lemma D.3. *If E holds and $n \geq \phi_\alpha/15$, then*

$$\text{cost}(C_{\text{large}}^\alpha, T_\alpha) \leq \text{cost}(C_{\text{large}}^\alpha, \text{OPT}) + \frac{4}{\alpha}(\text{cost}(P_1 \setminus \mathcal{B}_a(1), \text{OPT}) + \text{cost}(P_1, T_\alpha)).$$

Proof. For any $x \in X$ and any $i \in [k]$, we have

$$\rho(x, c_\alpha(x)) \leq \rho(x, c_\alpha(c_i^*)) \leq \rho(x, c_i^*) + \rho(c_i^*, c_\alpha(c_i^*)).$$

Therefore,

$$\begin{aligned} \text{cost}(C_{\text{large}}^\alpha, T_\alpha) &= \sum_{i \in I_{\text{large}}^\alpha} \sum_{x \in C_i^*} \rho(x, c_\alpha(x)) \leq \sum_{i \in I_{\text{large}}^\alpha} \sum_{x \in C_i^*} \rho(x, c_i^*) + \sum_{i \in I_{\text{large}}^\alpha} \sum_{x \in C_i^*} \rho(c_i^*, c_\alpha(c_i^*)) \\ &\leq \text{cost}(C_{\text{large}}^\alpha, \text{OPT}) + \sum_{i \in I_{\text{large}}^\alpha} |C_i^*| \cdot \rho(c_i^*, c_\alpha(c_i^*)). \end{aligned} \quad (6)$$

We now bound $\rho(c_i^*, c_\alpha(c_i^*))$. Define $A_i := C_i^* \cap P_1 \setminus \mathcal{B}_a(1)$ (we show below that $A_i \neq \emptyset$). From the definition of c_α , we have

$$\forall x \in A_i, \quad \rho(c_i^*, c_\alpha(c_i^*)) \leq \rho(c_i^*, c_\alpha(x)) \leq \rho(c_i^*, x) + \rho(x, c_\alpha(x)).$$

Hence,

$$\rho(c_i^*, c_\alpha(c_i^*)) \leq \frac{1}{|A_i|} \sum_{x \in A_i} (\rho(c_i^*, x) + \rho(x, c_\alpha(x))) = \frac{1}{|A_i|} (\text{cost}(A_i, \{c_i^*\}) + \text{cost}(A_i, T_\alpha)).$$

Denote $l_i := |C_i^*|$ and $l'_i := |A_i|$. Then, by Eq. (6),

$$\text{cost}(C_{\text{large}}^\alpha, T_\alpha) \leq \text{cost}(C_{\text{large}}^\alpha, \text{OPT}) + \sum_{i \in I_{\text{large}}^\alpha} \frac{l_i}{l'_i} (\text{cost}(A_i, \{c_i^*\}) + \text{cost}(A_i, T_\alpha)). \quad (7)$$

We now upper bound l_i/l'_i . Since $n \geq \phi_\alpha/15$, we have that $\mathcal{B}_a(1)$ is non-trivial. Since E holds, $\mathcal{B}_a(1)$ and $\{C_i^*\}_{i \in I_{\text{large}}^\alpha}$ are well-represented in P_1 for X . In addition, $|P_1| = \alpha|X|$. Hence,

$$l'_i = |C_i^* \cap P_1 \setminus \mathcal{B}_a(1)| \geq |C_i^* \cap P_1| - |\mathcal{B}_a(1) \cap P_1| \geq \alpha|C_i^*|/2 - (3/2)\alpha|\mathcal{B}_a(1)|. \quad (8)$$

By the definition of α -large optimal clusters, for all $i \in I_{\text{large}}^\alpha$, $\phi_\alpha \leq |C_i^*| = l_i$. Since \mathcal{B}_a is a $(\phi_\alpha/15)$ -linear bin division of X , by property 2 of linear bin divisions, we have $|\mathcal{B}_a(1)| \leq \frac{5}{2} \cdot \frac{\phi_\alpha}{15} = \phi_\alpha/6 \leq l_i/6$. Hence, by Eq. (8), $l'_i \geq \alpha l_i/2 - (3/2)\alpha \cdot (l_i/6) \geq \alpha l_i/4$. Therefore, $l_i/l'_i \leq 4/\alpha$. Note that this also implies that A_i is non-empty. Thus, from Eq. (7), we have

$$\begin{aligned} \text{cost}(C_{\text{large}}^\alpha, T_\alpha) &\leq \text{cost}(C_{\text{large}}^\alpha, \text{OPT}) + \frac{4}{\alpha} \left(\text{cost}(P_1 \cap C_{\text{large}}^\alpha \setminus \mathcal{B}_a(1), \text{OPT}) \right. \\ &\quad \left. + \text{cost}(P_1 \cap C_{\text{large}}^\alpha \setminus \mathcal{B}_a(1), T_\alpha) \right). \end{aligned}$$

Since $P_1 \cap C_{\text{large}}^\alpha \setminus \mathcal{B}_a(1) \subseteq P_1 \cap \mathcal{B}_a(1)$, this leads to the statement of the lemma. \square

In the next lemma, we bound $\text{cost}(P_1, T_\alpha)$, which appears in the RHS of the inequality given in Lemma D.3. In the proof of this lemma, the importance of selecting k_+ centers for T_α instead of only k becomes apparent, as it allows deriving an upper bound on the cost of T_α that disregards the outliers in $\mathcal{B}(1)$.

Lemma D.4. *If E holds and $n \geq \phi_\alpha/15$, then $\text{cost}(P_1, T_\alpha) \leq 2\beta \text{cost}(P_1 \setminus \mathcal{B}_a(1), \text{OPT})$.*

Proof. Denote the optimal k_+ -median solution for P_1 with centers from the entire X by

$$\text{OPT}_{P_1}^{k_+} := \underset{T \subseteq X, |T| \leq k_+}{\text{argmin}} \text{cost}(P_1, T).$$

Denote the optimal k_+ -median solution for P_1 with centers from P_1 by

$$\widetilde{\text{OPT}}_{P_1}^{k_+} := \underset{T \subseteq P_1, |T| \leq k_+}{\text{argmin}} \text{cost}(P_1, T).$$

It is well known [see, e.g., Guha et al., 2000] that for any set P_1 in any metric space (X, ρ) , the cost ratio between the solution that uses only centers from P_1 and the solution that uses the entire metric space is bounded by 2: $\text{cost}(P_1, \widetilde{\text{OPT}}_{P_1}^{k_+}) \leq 2\text{cost}(P_1, \text{OPT}_{P_1}^{k_+})$. Now, the black-box algorithm \mathcal{A} which outputs T_α in the first phase of `SelectProc` is a β -approximation offline algorithm. Therefore, we have

$$\text{cost}(P_1, T_\alpha) \leq \beta \text{cost}(P_1, \widetilde{\text{OPT}}_{P_1}^{k_+}) \leq 2\beta \text{cost}(P_1, \text{OPT}_{P_1}^{k_+}). \quad (9)$$

We bound $\text{cost}(P_1, \text{OPT}_{P_1}^{k_+})$ by the cost of the clustering defined by `OPT`, with additional centers that are the furthest from `OPT` in P_1 . Recall that \mathcal{B}_a is a $(\phi_\alpha/15)$ -linear bin division of X with respect to `OPT`. Define $A := \mathcal{B}_a(1) \cap P_1$. By property 4 of linear bin divisions, the points in A are the furthest from `OPT` in P_1 . We now bound the size of A . Since E holds and $n \geq \phi_\alpha/15$, $\mathcal{B}_a(1)$ is

non-trivial and well represented in P_1 for X . Therefore, $|A| \leq \frac{3}{2}\alpha|\mathcal{B}_a(1)|$. In addition, by property 2 of linear bin divisions, $|\mathcal{B}_a(1)| \leq \frac{5}{2} \cdot \frac{\phi_\alpha}{15} \leq \phi_\alpha/6$. Therefore, $|A| \leq \alpha\phi_\alpha/4 \leq 38 \log(32k/\delta)$.

It follows that $|\text{OPT} \cup A| \leq k + 38 \log(32k/\delta) = k_+$. Thus, by the definition of $\text{OPT}_{P_1}^{k_+}$,

$$\begin{aligned} \text{cost}(P_1, \text{OPT}_{P_1}^{k_+}) &\leq \text{cost}(P_1, \text{OPT} \cup A) = \text{cost}(P_1 \setminus A, \text{OPT} \cup A) \\ &= \text{cost}(P_1 \setminus \mathcal{B}_a(1), \text{OPT} \cup A) \leq \text{cost}(P_1 \setminus \mathcal{B}_a(1), \text{OPT}). \end{aligned}$$

Combining the inequality above with Eq. (9), we get the statement of the lemma. \square

Lemma D.3 and Lemma D.4 can be combined to get

$$\text{cost}(C_{\text{large}}^\alpha, T_\alpha) \leq \text{cost}(C_{\text{large}}^\alpha, \text{OPT}) + \frac{4}{\alpha}(2\beta + 1)\text{cost}(P_1 \setminus \mathcal{B}_a(1), \text{OPT}). \quad (10)$$

The proof of Lemma 5.7 is now almost immediate.

Proof of Lemma 5.7. We distinguish between two cases based on the stream length n . First, consider the case $n < \phi_\alpha/15$. In this case, all optimal clusters are smaller than $\phi_\alpha/15$, thus they are all α -small optimal clusters. Hence, $C_{\text{large}}^\alpha = \emptyset$, and so $\text{cost}(C_{\text{large}}^\alpha, T_\alpha) = 0$, which trivially satisfies the required upper bound.

Now, suppose that $n \geq \phi_\alpha/15$. Thus, \mathcal{B}_a is non-trivial. Since E holds, it follows that \mathcal{B}_a is well-represented in P_1 for X . Hence, $\forall i \in [L], |P_1 \cap \mathcal{B}(i)| \leq \frac{3}{2}\alpha|\mathcal{B}(i)|$. We now apply Lemma 5.5 with $W := X, T := \text{OPT}, A := P_1, \mathcal{B} := \mathcal{B}_a$, and $r := \frac{3}{2}\alpha$, we have

$$\text{cost}(P_1 \setminus \mathcal{B}_a(1), \text{OPT}) \leq \frac{9}{4}\alpha\text{cost}(X, \text{OPT}).$$

Combining this with Eq. (10), we get

$$\text{cost}(C_{\text{large}}^\alpha, T_\alpha) \leq \text{cost}(X, \text{OPT}) + 9(2\beta + 1)\text{cost}(X, \text{OPT}) \leq (18\beta + 10)\text{cost}(X, \text{OPT}),$$

As claimed. \square

D.2 Upper-bounding the cost estimate ψ

In Lemma 5.7, we showed that the cost of T_α on large optimal clusters is bounded by a constant factor over the optimal cost. In the second phase, `SelectProc` calculates an estimate of this cost, defined as $\psi := \frac{1}{3\alpha}\text{cost}_{2\alpha(k+1)\phi_\alpha}(P_2, T_\alpha)$. In this section, we prove the upper bound in Lemma 5.8, which shows that the removal of $2\alpha(k+1)\phi_\alpha$ outliers from the cost estimate indeed guarantees that ψ is upper bounded by the cost of T_α on large clusters. The proof of the lower bound in Lemma 5.8 is provided in Section D.3.

Proof of Lemma 5.8 (Upper bound). We distinguish between two cases, based on the size of the stream. If $n < 2(k+1)\phi_\alpha$, then $|P_2| < 2\alpha(k+1)\phi_\alpha$. In this case, $\text{cost}_{2\alpha(k+1)\phi_\alpha}(P_2, T_\alpha) = 0$, thus $\psi = 0$ and the upper bound trivially holds.

Now, consider the case $n \geq 2(k+1)\phi_\alpha$. In this case, $|X \setminus P_1| \geq (1-\alpha)n \geq (k+1)\phi_\alpha$. Recall that $F_b = \text{far}_{k\phi_\alpha}(X \setminus P_1, T_\alpha)$, and \mathcal{B}_b is a $(\phi_\alpha/3)$ -linear bin division of $(X \setminus P_1) \setminus F_b$ with respect to T_α . F_b is non-trivial, since $|X \setminus P_1| \geq k\phi_\alpha$. Similarly, \mathcal{B}_b is non-trivial, since $|(X \setminus P_1) \setminus F_b| \geq (k+1)\phi_\alpha - k\phi_\alpha \geq \phi_\alpha/3$.

As shown in Section 5.2, $\text{cost}_{k\phi_\alpha}(X, T_\alpha) \leq \text{cost}(C_{\text{large}}^\alpha, T_\alpha)$. Hence, to prove the lemma, it suffices to prove that $\psi \leq \text{cost}_{k\phi_\alpha}(X, T_\alpha)$. In the rest of the proof we derive this inequality.

We first show that $(F_b \cup \mathcal{B}_b(1)) \cap P_2 \subseteq \text{far}_{2\alpha(k+1)\phi_\alpha}(P_2, T_\alpha)$. This would imply that the points in $F_b \cup \mathcal{B}_b(1)$ do not contribute to the cost $\text{cost}_{2\alpha(k+1)\phi_\alpha}(P_2, T_\alpha)$, which is used in the calculation of ψ . Since E holds, F_b and $\mathcal{B}_b(1)$ (which are both non-trivial) are well-represented in P_2 for $X \setminus P_1$. We also have $|P_2|/|X \setminus P_1| = \alpha/(1-\alpha) \leq \frac{4}{3}\alpha$, where the last inequality follows since $\alpha \leq 1/6$. Therefore, $|F_b \cap P_2| \leq \frac{3}{2} \cdot \frac{4}{3}\alpha|F_b| = 2\alpha k\phi_\alpha$, and similarly, $|\mathcal{B}_b(1) \cap P_2| \leq 2\alpha|\mathcal{B}_b(1)| \leq 2\alpha\phi_\alpha$,

where the last inequality follows from property 2 of linear bin divisions, since $(5/2) \cdot (\phi_\alpha/3) \leq \phi_\alpha$. It follows that

$$|(F_b \cup \mathcal{B}_b(1)) \cap P_2| \leq |F_b \cap P_2| + |\mathcal{B}_b(1) \cap P_2| \leq 2\alpha(k+1)\phi_\alpha.$$

By the definitions of F_b and \mathcal{B}_b , and by property 4 of linear bin division, we have that the points in $F_b \cup \mathcal{B}_b(1)$ are the furthest points from T_α in $X \setminus P_1$. Thus, $(F_b \cup \mathcal{B}_b(1)) \cap P_2$ are the furthest points from T_α in P_2 . Therefore, $(F_b \cup \mathcal{B}_b(1)) \cap P_2 \subseteq \text{far}_{2\alpha(k+1)\phi_\alpha}(P_2, T_\alpha)$, as we wished to show. This implies that

$$\text{cost}_{2\alpha(k+1)\phi_\alpha}(P_2, T_\alpha) \leq \text{cost}(P_2 \setminus (F_b \cup \mathcal{B}_b(1)), T_\alpha). \quad (11)$$

Now, since \mathcal{B}_b is well-represented in P_2 for $X \setminus P_1$, denoting the number of bins in \mathcal{B}_b by L , we have that for all $i \in [L]$, $|\mathcal{B}_b(i) \cap P_2|/|\mathcal{B}_b(i)| \leq 2\alpha$. Since \mathcal{B}_b is a partition of $(X \setminus P_1) \setminus F_b$, we have that $P_2 \cap \mathcal{B}_b(i) = (P_2 \setminus F_b) \cap \mathcal{B}_b(i)$. Therefore, we can apply Lemma 5.5 with $W := (X \setminus P_1) \setminus F_b$, $\mathcal{B} := \mathcal{B}_b$, $T := T_\alpha$, $A := P_2 \setminus F_b$, and $r := 2\alpha$. It follows that

$$\text{cost}((P_2 \setminus F_b) \setminus \mathcal{B}_b(1), T_\alpha) \leq 3\alpha \text{cost}((X \setminus P_1) \setminus F_b, T_\alpha) = 3\alpha \text{cost}_{k\phi_\alpha}(X \setminus P_1, T_\alpha). \quad (12)$$

The last equality follows from the definition of F_b . By combining Eq. (11) and Eq. (12), we get that $\text{cost}_{2\alpha(k+1)\phi_\alpha}(P_2, T_\alpha) \leq 3\alpha \text{cost}_{k\phi_\alpha}(X \setminus P_1, T_\alpha)$. Plugging in the definition of ψ , this implies that $\psi \leq \text{cost}_{k\phi_\alpha}(X \setminus P_1, T_\alpha) \leq \text{cost}_{k\phi_\alpha}(X, T_\alpha)$, as required. \square

D.3 Bounding the number of selected far points

In this section, we prove Lemma 5.9, which bounds the number of points that are selected as centers and their distance from T_α is more than $\psi/(k\tau)$. We first prove an auxiliary lemma, and provide the proof for the lower bound in Lemma 5.8, which is used in the proof of Lemma 5.9. The auxiliary lemma provides a relationship between the costs of consecutive bins.

Lemma D.5. *Let L be the number of bins in \mathcal{B}_c . If E holds, then*

$$\forall i \in [L-1], \quad \text{cost}(\mathcal{B}_c(i) \cap P_2, T_\alpha) \geq \frac{\alpha}{3} \text{cost}(\mathcal{B}_c(i+1), T_\alpha).$$

Proof. Fix $i \in [L-1]$, and define $b := \max_{x \in \mathcal{B}_c(i+1)} \rho(x, T_\alpha)$. By property 4 of linear bin divisions, $\rho(x, T_\alpha) \geq b$ for all $x \in \mathcal{B}_c(i)$. Therefore,

$$\begin{aligned} \text{cost}(\mathcal{B}_c(i) \cap P_2, T_\alpha) &\geq |\mathcal{B}_c(i) \cap P_2| \cdot b, \quad \text{and} \\ \text{cost}(\mathcal{B}_c(i+1), T_\alpha) &\leq |\mathcal{B}_c(i+1)| \cdot b. \end{aligned}$$

It follows that

$$\frac{\text{cost}(\mathcal{B}_c(i) \cap P_2, T_\alpha)}{\text{cost}(\mathcal{B}_c(i+1), T_\alpha)} \geq \frac{|\mathcal{B}_c(i) \cap P_2|}{|\mathcal{B}_c(i+1)|}.$$

Thus, to prove the claim, it suffices to lower-bound the RHS by $\alpha/3$. Since E holds, \mathcal{B}_c is well-represented in P_2 for $X \setminus P_1$. Since $|P_2|/|X \setminus P_1| = \alpha/(1-\alpha) \geq \alpha$, it follows that $|\mathcal{B}_c(i) \cap P_2| \geq (\alpha/2)|\mathcal{B}_c(i)|$. By property 3 of linear bin divisions, $|\mathcal{B}_c(i)| \geq \frac{2}{3}|\mathcal{B}_c(i+1)|$. Therefore, $|\mathcal{B}_c(i) \cap P_2| \geq (\alpha/3)|\mathcal{B}_c(i+1)|$, which proves the claim. \square

The next step is to prove the lower bound in Lemma 5.8, which bounds from below the cost estimate ψ calculated in `SelectProc`. Recall that $\psi := \frac{1}{3\alpha} \text{cost}_{2\alpha(k+1)\phi_\alpha}(P_2, T_\alpha)$.

Proof of Lemma 5.8 (Lower bound). First, note that if $|X \setminus P_1| \leq 5(k+1)\phi_\alpha$, then $\text{cost}_{5(k+1)\phi_\alpha}(X \setminus P_1, T_\alpha) = 0$, and so the lemma trivially holds. We henceforth assume that $|X \setminus P_1| > 5(k+1)\phi_\alpha$. Recall that $F_c = \text{far}_{4(k+1)\phi_\alpha}(X \setminus P_1, T_\alpha)$, and \mathcal{B}_c is a $(\phi_\alpha/3)$ -linear bin division of $(X \setminus P_1) \setminus F_c$ with respect to T_α . Under the assumption, $|X \setminus P_1| > 4(k+1)\phi_\alpha$, thus F_c is non-trivial. Similarly, \mathcal{B}_c is non-trivial, since $|(X \setminus P_1) \setminus F_c| \geq 5(k+1)\phi_\alpha - 4(k+1)\phi_\alpha > \phi_\alpha/3$.

To prove the lemma, we show the following inequality:

$$\text{cost}_{2\alpha(k+1)\phi_\alpha}(P_2, T_\alpha) \geq \frac{\alpha}{3} \text{cost}_{5(k+1)\phi_\alpha}(X \setminus P_1, T_\alpha). \quad (13)$$

This would immediately imply that $\psi = \frac{1}{3\alpha} \text{cost}_{2\alpha(k+1)\phi_\alpha}(P_2, T_\alpha) \geq \frac{1}{9} \text{cost}_{5(k+1)\phi_\alpha}(X \setminus P_1, T_\alpha)$, as required. Denote by L the number of bins in \mathcal{B}_c . Define

$$\tilde{X} := \bigcup_{i=2}^L \mathcal{B}_c(i) = (X \setminus P_1) \setminus (F_c \cup \mathcal{B}_c(1)).$$

See Figure 2 for an illustration of the sets used in this proof. To prove Eq. (13), we show the following string of inequalities:

$$\text{cost}_{2\alpha(k+1)\phi_\alpha}(P_2, T_\alpha) \geq \text{cost}(P_2 \setminus F_c, T_\alpha) \geq \frac{\alpha}{3} \text{cost}(\tilde{X}, T_\alpha) \geq \frac{\alpha}{3} \text{cost}_{5(k+1)\phi_\alpha}(X \setminus P_1, T_\alpha). \quad (14)$$




Figure 2: The sets used in the proof of the lower bound in Lemma 5.8. The horizontal axis represents the distance of the points in the set from T_α , in decreasing order.

First, we prove the first inequality in Eq. (14) by removing far points from P_2 . Recall that $F_c := \text{far}_{4(k+1)\phi_\alpha}(X \setminus P_1, T_\alpha)$ are the furthest points from T_α in $X \setminus P_1$. Since E holds, F_c is well-represented in P_2 for $X \setminus P_1$. Since $|P_2|/|X \setminus P_1| = \alpha/(1-\alpha) \geq \alpha$, this implies that

$$|P_2 \cap F_c| \geq \frac{\alpha}{2} |F_c| = 2(k+1)\phi_\alpha\alpha.$$

Hence, the $2(k+1)\phi_\alpha\alpha$ furthest points from T_α in P_2 are in F_c . Formally, $\text{far}_{2(k+1)\phi_\alpha\alpha}(P_2, T_\alpha) \subseteq F_c$. Therefore, $\text{cost}_{2(k+1)\phi_\alpha\alpha}(P_2, T_\alpha) \geq \text{cost}(P_2 \setminus F_c, T_\alpha)$, which proves the first inequality of Eq. (14).

Next, we prove the last inequality of Eq. (14). By property 2 of linear bin divisions, we have $|\mathcal{B}_c(1)| \leq \frac{5}{2} \cdot \frac{\phi_\alpha}{3} \leq \phi_\alpha$. Therefore,

$$|F_c \cup \mathcal{B}_c(1)| \leq |F_c| + |\mathcal{B}_c(1)| \leq 4(k+1)\phi_\alpha + \phi_\alpha \leq 5(k+1)\phi_\alpha.$$

In addition, by the definition of F_c and property 4 of linear bin divisions, the points in $F_c \cup \mathcal{B}_c(1)$ are the furthest from T_α out of the points in $X \setminus P_1$. Hence, $F_c \cup \mathcal{B}_c(1) \subseteq \text{far}_{5(k+1)\phi_\alpha}(X \setminus P_1, T_\alpha)$. It follows that $\tilde{X} \supseteq (X \setminus P_1) \setminus \text{far}_{5(k+1)\phi_\alpha}(X \setminus P_1, T_\alpha)$. Therefore, $\text{cost}(\tilde{X}, T_\alpha) \geq \text{cost}_{5(k+1)\phi_\alpha}(X \setminus P_1, T_\alpha)$, which proves the last inequality in Eq. (14).

To complete the proof, we have left to prove the second inequality of Eq. (14), which states that $\text{cost}(P_2 \setminus F_c, T_\alpha) \geq \frac{\alpha}{3} \text{cost}(\tilde{X}, T_\alpha)$. We have

$$\text{cost}(P_2 \setminus F_c, T_\alpha) = \sum_{i \in [L]} \text{cost}(\mathcal{B}_c(i) \cap (P_2 \setminus F_c), T_\alpha) = \sum_{i \in [L]} \text{cost}(\mathcal{B}_c(i) \cap P_2, T_\alpha). \quad (15)$$

The last equality follows since \mathcal{B} is a partition of $(X \setminus P_1) \setminus F_c$, thus $\mathcal{B}(i) \cap F_c = \emptyset$. By Lemma D.5, since E holds, $\forall i \in [L-1]$, $\text{cost}(\mathcal{B}_c(i) \cap P_2, T_\alpha) \geq \frac{\alpha}{3} \text{cost}(\mathcal{B}_c(i+1), T_\alpha)$. It follows that

$$\text{cost}(P_2 \setminus F_c, T_\alpha) \geq \sum_{i \in [L-1]} \text{cost}(\mathcal{B}_c(i) \cap P_2, T_\alpha) \geq \frac{\alpha}{3} \sum_{i=2}^L \text{cost}(\mathcal{B}_c(i), T_\alpha) = \frac{\alpha}{3} \text{cost}(\tilde{X}, T_\alpha).$$

The last equality follows from the definition of \tilde{X} . This completes the proof of the second inequality of Eq. (14), thus proving the lemma. \square

Based on the lemmas above, Lemma 5.9 can now be proved.

Proof of Lemma 5.9. Denote the set of points in P_3 that exceed the distance threshold used in `SelectProc` by

$$\mathcal{T} := \{x \in P_3 \mid \rho(x, T_\alpha) > \psi/(k\tau)\}.$$

Then $N = |\mathcal{T}|$. Recall that $F_d := \text{far}_{5(k+1)\phi_\alpha}(X \setminus P_1, T_\alpha)$. We have $N \leq |\mathcal{T} \setminus F_d| + |P_3 \cap F_d|$. To upper bound $|\mathcal{T} \setminus F_d|$, note that from the definition of F_d and \mathcal{T} ,

$$\text{cost}_{5(k+1)\phi_\alpha}(X \setminus P_1, T_\alpha) = \text{cost}((X \setminus P_1) \setminus F_d, T_\alpha) \geq \text{cost}(\mathcal{T} \setminus F_d, T_\alpha) \geq |\mathcal{T} \setminus F_d| \cdot \psi/(k\tau).$$

On the other hand, by Lemma 5.8, $\psi \geq \frac{1}{9} \text{cost}_{5(k+1)\phi_\alpha}(X \setminus P_1, T_\alpha)$. It follows that $|\mathcal{T} \setminus F_d| \leq 9k\tau$.

To upper bound $|P_3 \cap F_d|$, consider first the case where F_d is trivial. In this case, $|X \setminus P_1| < 5(k+1)\phi_\alpha$, thus

$$|P_3 \cap F_d| \leq |P_3| = |X \setminus P_1| \cdot (|P_3|/|X \setminus P_1|) \leq 5(k+1)\phi_\alpha \cdot \frac{\gamma}{1-\alpha}.$$

If F_d is non-trivial, then since E holds, F_d is well-represented in P_3 for $X \setminus P_1$. Since we have $|P_3|/|X \setminus P_1| = \frac{\gamma}{1-\alpha}$, it follows that

$$|P_3 \cap F_d| \leq \frac{3}{2} \frac{\gamma}{1-\alpha} |F_d| = \frac{3}{2} \frac{\gamma}{1-\alpha} \cdot 5(k+1)\phi_\alpha < 8(k+1)\phi_\alpha \frac{\gamma}{1-\alpha}.$$

Therefore, in both cases, $|P_3 \cap F_d| \leq 8(k+1)\phi_\alpha \frac{\gamma}{1-\alpha}$. Combining this with the upper bound on $|\mathcal{T} \setminus F_d|$, the statement of the lemma follows. \square

D.4 Cost approximation for small and large optimal clusters

In this section, we prove Lemma D.1 and Lemma D.2. The former lemma upper bounds the cost of the selected centers on optimal clusters whose center appears in the third stage, and is used for bounding the cost of points in small optimal clusters. The latter lemma upper bounds the cost of the selected centers on large optimal clusters, under an appropriate setting of algorithm parameters.

First, we provide an auxiliary lemma, which will be used to prove both of these lemmas. It shows that all points appearing in the last phase are close to some center in the set of selected centers.

Lemma D.6. *For all $x \in P_3$, $\rho(x, T_{\text{out}}) \leq 2\psi/(k\tau)$.*

Proof. If x is selected by Alg. 2, then $x \in T_{\text{out}}$, so $\rho(x, T_{\text{out}}) = 0$. Thus, in this case, the lemma trivially holds. Now, suppose that x is not selected as a center. Recall that $T_\alpha = \{c_1, \dots, c_{k_+}\}$, and denote $i := \text{argmin}_{i \in [k_+]} \rho(x, c_i)$. Due to the selection conditions on line 12 of Alg. 2, we have that, since x was not selected, $\rho(x, c_i) \leq \psi/(k\tau)$ and $\text{Near}_i = \text{TRUE}$ when x is observed. From the latter, it follows that some $y \in C_i$ was selected such that $\rho(c_i, y) \leq \psi/(k\tau)$. Therefore, $\rho(x, T_{\text{out}}) \leq \rho(x, y) \leq \rho(x, c_i) + \rho(c_i, y) \leq 2\psi/(k\tau)$. \square

The proof of Lemma D.1 is now almost immediate.

Proof of Lemma D.1. Fix $i \in [k]$. Observe that by the triangle inequality,

$$\text{cost}(C_i^*, T_{\text{out}}) \leq \text{cost}(C_i^*, \{c_i^*\}) + |C_i^*| \rho(c_i^*, T_{\text{out}}).$$

By the assumption of the lemma, $c_i^* \in P_3$. Applying Lemma D.6 to $x := c_i^*$, we get that $\rho(c_i^*, T_{\text{out}}) \leq 2\psi/(k\tau)$. Combined with the inequality above, this proves the lemma. \square

We now turn to prove Lemma D.2, which upper bounds the cost of the selected centers on large optimal clusters. First, we relate the cost of points from α -large optimal clusters in the third phase to the cost of the clustering T_α calculated in the first phase.

Lemma D.7. *If E holds, $\tau = \phi_\alpha$ and $M = \log(8k_+/\delta)$, then $\text{cost}(C_{\text{large}}^\alpha \cap P_3, T_{\text{out}}) \leq 13 \text{cost}(C_{\text{large}}^\alpha, T_\alpha)$.*

Proof. Recall that C_1, \dots, C_{k_+} are the clusters induced by the centers in T_α . We distinguish between *light clusters* and *heavy clusters*. Light clusters are those that include many points from small optimal clusters in the third phase. Formally, the set of light clusters is defined as

$$\text{Lt} := \{i \in [k_+] \mid |C_i \cap P_3 \cap C_{\text{small}}^\alpha| / |C_i \cap P_3| \geq 1/4\}.$$

Heavy clusters are non-light clusters, denoted $\text{Hv} := [k_+] \setminus \text{Lt}$. We have

$$\text{cost}(C_{\text{large}}^\alpha \cap P_3, T_{\text{out}}) \leq \sum_{i \in \text{Lt}} \text{cost}(C_i \cap P_3, T_{\text{out}}) + \sum_{i \in \text{Hv}} \text{cost}(C_i \cap P_3 \cap C_{\text{large}}^\alpha, T_{\text{out}}). \quad (16)$$

We bound each of these terms separately.

For light clusters, we show that these clusters are relatively small, and bound the cost contributed by each point in the cluster. The total number of points in light clusters is

$$\sum_{i \in \text{Lt}} |C_i \cap P_3| \leq 4 \sum_{i \in \text{Lt}} |C_i \cap P_3 \cap C_{\text{small}}^\alpha| \leq 4|C_{\text{small}}^\alpha| \leq 4k\phi_\alpha. \quad (17)$$

Moreover, for each $x \in C_i \cap P_3$ for $i \in \text{Lt}$, we have by Lemma D.6 that $\rho(x, T_{\text{out}}) \leq 2\psi/(k\tau)$. By the assumptions of the lemma, $\tau := \phi_\alpha$. Hence, combining with Eq. (17), we get

$$\sum_{i \in \text{Lt}} \text{cost}(C_i \cap P_3, T_{\text{out}}) \leq \frac{2\psi}{k\phi_\alpha} \sum_{i \in \text{Lt}} |C_i \cap P_3| \leq 8\psi \leq 8\text{cost}(C_{\text{large}}^\alpha, T_\alpha). \quad (18)$$

The last inequality follows from Lemma 5.8.

For heavy clusters, let $i \in \text{Hv}$. Denote $r_i := \rho(c_i, T_{\text{out}})$. By the triangle inequality,

$$\text{cost}(C_i \cap P_3 \cap C_{\text{large}}^\alpha, T_{\text{out}}) \leq \text{cost}(C_i \cap P_3 \cap C_{\text{large}}^\alpha, \{c_i\}) + |C_i \cap P_3 \cap C_{\text{large}}^\alpha| \cdot r_i.$$

Let $\text{cost}_i := \{x \in C_i \cap P_3 \mid \rho(x, c_i) \geq r_i\}$. To upper bound r_i , note that

$$\text{cost}(C_i \cap P_3 \cap C_{\text{large}}^\alpha, \{c_i\}) \geq |\text{cost}_i \cap C_{\text{large}}^\alpha| \cdot r_i.$$

Therefore,

$$r_i \leq \frac{\text{cost}(C_i \cap P_3 \cap C_{\text{large}}^\alpha, \{c_i\})}{|\text{cost}_i \cap C_{\text{large}}^\alpha|}.$$

It follows that

$$\text{cost}(C_i \cap P_3 \cap C_{\text{large}}^\alpha, T_{\text{out}}) \leq \text{cost}(C_i \cap P_3 \cap C_{\text{large}}^\alpha, \{c_i\}) \cdot \left(1 + \frac{|C_i \cap P_3 \cap C_{\text{large}}^\alpha|}{|\text{cost}_i \cap C_{\text{large}}^\alpha|}\right). \quad (19)$$

We now lower-bound $|\text{cost}_i \cap C_{\text{large}}^\alpha|$. Alg. 2 selects the first $M = \log(8k_+/\delta)$ points from cluster C_i observed in the third phase. From part 4 of the event E_i , we have that at least one of those points, call it x^* , is closer to c_i than at least half of the points in $C_i \cap P_3$. Formally, $\rho(c_i, x^*) \leq \rho(c_i, x)$ for at least half of the points $x \in C_i \cap P_3$. Since $x^* \in T_{\text{out}}$, we have $r_i \leq \rho(c_i, x^*)$, and so $r_i \leq \rho(c_i, x)$ for at least half of the points $x \in C_i \cap P_3$. It follows that $|\text{cost}_i| \geq |C_i \cap P_3|/2$.

On the other hand, since $i \in \text{Hv}$, we have $|\text{cost}_i \cap C_{\text{small}}^\alpha| \leq |C_i \cap P_3 \cap C_{\text{small}}^\alpha| \leq |C_i \cap P_3|/4$. Therefore, $|\text{cost}_i \cap C_{\text{large}}^\alpha| = |\text{cost}_i| - |\text{cost}_i \cap C_{\text{small}}^\alpha| \geq |C_i \cap P_3|/4 \geq |C_i \cap P_3 \cap C_{\text{large}}^\alpha|/4$. Combined with Eq. (19), we get

$$\forall i \in \text{Hv}, \quad \text{cost}(C_i \cap P_3 \cap C_{\text{large}}^\alpha, T_{\text{out}}) \leq 5\text{cost}(C_i \cap P_3 \cap C_{\text{large}}^\alpha, c_i).$$

It follows that $\sum_{i \in \text{Hv}} \text{cost}(C_i \cap P_3 \cap C_{\text{large}}^\alpha, T_{\text{out}}) \leq 5\text{cost}(C_{\text{large}}^\alpha, T_\alpha)$. Combining this with Eq. (16) and Eq. (18), we get the statement of the lemma. \square

Based on the lemma above, the proof of Lemma D.2 can be now provided.

Proof of Lemma D.2. To bound $\text{cost}(C_{\text{large}}^\alpha, T_{\text{out}})$, we bound $\text{cost}(C_{\text{large}}^\alpha \cap P_3, T_{\text{out}})$ and $\text{cost}(C_{\text{large}}^\alpha \cap (P_1 \cup P_2), T_{\text{out}})$ separately. The former is bounded using Lemma D.7, which gives $\text{cost}(C_{\text{large}}^\alpha \cap P_3, T_{\text{out}}) \leq 13\text{cost}(C_{\text{large}}^\alpha, T_\alpha)$.

To bound $\text{cost}(C_{\text{large}}^\alpha \cap (P_1 \cup P_2), T_{\text{out}})$, we first show that at least half of the points in each α -large optimal cluster C_i^* are in P_3 . Fix $i \in I_{\text{large}}^\alpha$. Since E holds, C_i^* is well-represented in $P_1 \cup P_2$ for X . Since $|P_1 \cup P_2|/|X| = 2\alpha$ and $\alpha \leq 1/6$, it follows that $|C_i^* \cap (P_1 \cup P_2)|/|C_i^*| \leq \frac{3}{2} \cdot 2\alpha \leq 3\alpha \leq 1/2$. Since $\gamma = 1 - 2\alpha$, we have $P_1 \cup P_2 \cup P_3 = X$, hence $|C_i^* \cap P_3|/|C_i^*| \geq 1/2$. Now, define an injection $\mu : P_1 \cup P_2 \rightarrow P_3$ such that $x \in C_i^*$ if and only if $\mu(x) \in C_i^*$. Such an injection exists

since most of the points in each α -large optimal cluster are in P_3 . For each $i \in I_{\text{large}}^\alpha$ and each $x \in C_i^* \cap (P_1 \cup P_2)$, we have

$$\rho(x, T_{\text{out}}) \leq \rho(x, c_i^*) + \rho(c_i^*, \mu(x)) + \rho(\mu(x), T_{\text{out}}). \quad (20)$$

Now, since μ is an injection, we have that for any $T \subseteq X$,

$$\text{cost}(\{\mu(x) \mid x \in C_i^* \cap (P_1 \cup P_2)\}, T) \leq \text{cost}(C_i^* \cap P_3, T).$$

Summing Eq. (20) over $x \in C_i^* \cap (P_1 \cup P_2)$, it follows that

$$\begin{aligned} \text{cost}(C_i^* \cap (P_1 \cup P_2), T_{\text{out}}) &\leq \text{cost}(C_i^* \cap (P_1 \cup P_2), \{c_i^*\}) + \text{cost}(C_i^* \cap P_3, \{c_i^*\}) + \text{cost}(C_i^* \cap P_3, T_{\text{out}}) \\ &= \text{cost}(C_i^*, \text{OPT}) + \text{cost}(C_i^* \cap P_3, T_{\text{out}}). \end{aligned}$$

Summing over $i \in I_{\text{large}}^\alpha$, we get

$$\text{cost}(C_{\text{large}}^\alpha \cap (P_1 \cup P_2), T_{\text{out}}) \leq \text{cost}(C_{\text{large}}^\alpha, \text{OPT}) + \text{cost}(C_{\text{large}}^\alpha \cap P_3, T_{\text{out}}).$$

It follows that

$$\begin{aligned} \text{cost}(C_{\text{large}}^\alpha, T_{\text{out}}) &\leq \text{cost}(C_{\text{large}}^\alpha \cap (P_1 \cup P_2), T_{\text{out}}) + \text{cost}(C_{\text{large}}^\alpha \cap P_3, T_{\text{out}}) \\ &\leq \text{cost}(C_{\text{large}}^\alpha, \text{OPT}) + 2\text{cost}(C_{\text{large}}^\alpha \cap P_3, T_{\text{out}}). \end{aligned}$$

Combined with the upper bound on $\text{cost}(C_{\text{large}}^\alpha \cap P_3, T_{\text{out}})$ from Lemma D.7, we get the statement of the lemma. \square

This is the final lemma required for proving Theorem 5.1. Thus, this completes the proof of the main theorem, Theorem 4.1.