

# Local Search $k$ -means++ with Foresight

Theo Conrads

Department of Computer Science, University of Cologne, Germany

Lukas Drexler  

Faculty of Mathematics and Natural Sciences, Department of Computer Science,  
Heinrich Heine University Düsseldorf, Germany

Joshua Könen  

Institute of Computer Science, University of Bonn, Germany

Daniel R. Schmidt  

Faculty of Mathematics and Natural Sciences, Department of Computer Science  
Heinrich Heine University Düsseldorf, Germany

Melanie Schmidt  

Faculty of Mathematics and Natural Sciences, Department of Computer Science,  
Heinrich Heine University Düsseldorf, Germany

---

## Abstract

Since its introduction in 1957, Lloyd’s algorithm for  $k$ -means clustering has been extensively studied and has undergone several improvements. While in its original form it does not guarantee any approximation factor at all, Arthur and Vassilvitskii (SODA 2007) proposed  $k$ -means++ which enhances Lloyd’s algorithm by a seeding method which guarantees a  $\mathcal{O}(\log k)$ -approximation in expectation. More recently, Lattanzi and Sohler (ICML 2019) proposed LS++ which further improves the solution quality of  $k$ -means++ by local search techniques to obtain a  $\mathcal{O}(1)$ -approximation. On the practical side, the *greedy* variant of  $k$ -means++ is often used although its worst-case behaviour is provably worse than for the standard  $k$ -means++ variant.

We investigate how to improve LS++ further in practice. We study two options for improving the practical performance: (a) Combining LS++ with greedy  $k$ -means++ instead of  $k$ -means++, and (b) Improving LS++ by better entangling it with Lloyd’s algorithm. Option (a) worsens the theoretical guarantees of  $k$ -means++ but improves the practical quality also in combination with LS++ as we confirm in our experiments. Option (b) is our new algorithm, Foresight LS++. We experimentally show that FLS++ improves upon the solution quality of LS++. It retains its asymptotic runtime and its worst-case approximation bounds.

**2012 ACM Subject Classification** Mathematics of computing → Combinatorial algorithms; Theory of computation → Facility location and clustering; Information systems → Clustering

**Keywords and phrases**  $k$ -means clustering,  $k$ means++, greedy, local search

**Digital Object Identifier** 10.4230/LIPIcs.SEA.2024.7

**Supplementary Material** *Software*: <https://github.com/algo-hhu/FLSpp> [11]  
archived at `swb:1:dir:39136777e542456572a51515813b6cb377ed0940`

**Funding** *Lukas Drexler*: Gefördert durch die Deutsche Forschungsgemeinschaft (DFG) – Projekt-nummer 456558332 – funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Projektnummer 456558332.

*Joshua Könen*: funded by Lamarr Institute for Machine Learning and Artificial Intelligence – lamarr-institute.org

*Melanie Schmidt*: Gefördert durch die Deutsche Forschungsgemeinschaft (DFG) – Projektnummer 456558332 – funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Projektnummer 456558332.



© Theo Conrads, Lukas Drexler, Joshua Könen, Daniel R. Schmidt, and Melanie Schmidt;  
licensed under Creative Commons License CC-BY 4.0

22nd International Symposium on Experimental Algorithms (SEA 2024).

Editor: Leo Liberti; Article No. 7; pp. 7:1–7:20



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

## 1 Introduction

In the vast area of clustering, one of the most popular approaches is min-sum-of-squared-error clustering, as modeled by the  $k$ -means cost function. Given a set of vectors  $X = \{x_1, \dots, x_n\} \subset \mathbb{R}^d$  and a number of clusters  $k \in \mathbb{N}$ , the aim of  $k$ -means clustering is to find a set  $C = \{c_1, \dots, c_k\}$  of  $k$  centers that minimizes  $\sum_{i=1}^n \min_{j=1, \dots, k} \|x_i - c_j\|^2$ , i.e., the sum of the squared distances of all points to their respective closest center in  $C$ . For half a century, the most known algorithm for minimizing the  $k$ -means cost function was a local search heuristic often called  $k$ -means algorithm or *Lloyd's algorithm*, developed in 1957. The main steps of the algorithm are as follows:

1. Find an initial set of  $k$  centers, e.g., randomly chosen from  $X$ .
2. Repeat a given number  $s$  times or until convergence:
  - a. For every  $x \in X$ , find a closest center  $\phi(x)$  (ties broken arbitrarily) and use this information to form clusters  $C_1, \dots, C_k$  by  $C_i = \phi^{-1}(c_i)$ .
  - b. For all  $C_i$ , compute the mean  $\mu(C_i) = \frac{1}{|C_i|} \sum_{x \in C_i} x$  and then replace  $C$  by  $C_{\text{new}} = \{\mu(C_1), \dots, \mu(C_k)\}$ .


Notice that this description ignores edge cases like clusters which “run empty” during step 2(a). We also do not discuss variations of other stopping criteria here. It is well documented [7, 15, 26] that the implementation of the first step of Lloyd's algorithm is crucial because bad initial centers can lead to bad local optima. It is also easy to find theoretical worst case examples where the local optimum is arbitrarily bad [20].

In 2007, Arthur and Vassilvitskii [4] proposed a method to significantly improve the first step of Lloyd's algorithm, leading to a new de-facto-standard algorithm for the  $k$ -means problem. The core idea is to choose the initial centers by an adaptive sampling procedure known as  $d^2$ -sampling initialization (also see Algorithm 1 below):

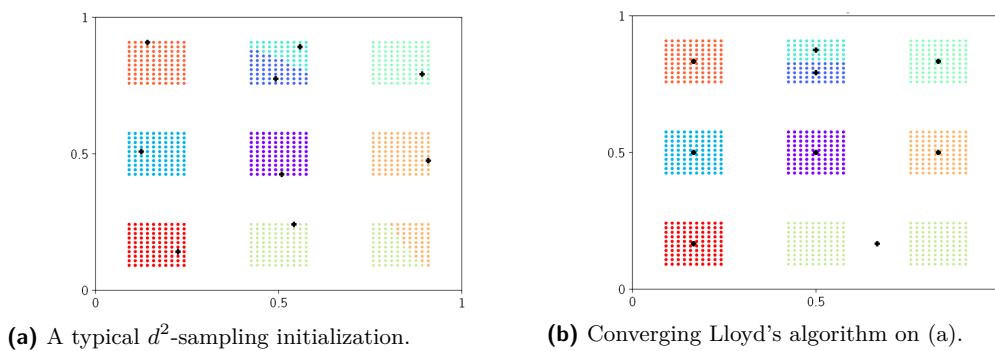
1. Choose  $c_1$  uniformly at random from  $X$ .
2. For  $i = 2, \dots, k$ :
  - For all  $x \in X$ , compute  $p(x) := \frac{\min_{c \in C} \|x - c\|^2}{\sum_{y \in X} \min_{c \in C} \|y - c\|^2}$
  - Sample a point  $c_i$  where every  $x \in X$  has probability  $p(x)$
3. Return  $C = \{c_1, \dots, c_k\}$

The algorithm proposed in [4] first computes  $C$  by this routine and then runs Lloyd's algorithm with  $C$  as the initial center set. This combination of  $d^2$ -sampling and Lloyd's algorithm is called  $k$ -means++.

Let us consider some background to better understand the advantages of  $k$ -means++ (and the subsequent improvements): There are two main reasons why a clustering obtained with Lloyd's algorithm may be bad. Firstly, it may be that the underlying structure of the data does not fit the  $k$ -means objective, e.g., because the points are not well-clusterable with spherical clusters, or because we chose the wrong  $k$ . In this case, we should choose a different clustering method like kernel  $k$ -means, hierarchical clustering methods or density based clustering. But secondly, it may be that although optimal clusters with respect to the  $k$ -means objective are indeed perfect for us, Lloyd's algorithm does not find them because it converges to a local optimum. Figure 1 gives a visual example for this: We see a  $k$ -means based image compression with four colors (The input pixels are represented by three-dimensional points based on the RGB-values, these points are clustered, and then the color of every pixel is replaced by the mean color of its cluster. Clusters are not contiguous within the picture). In this example, we can see how the difference in  $k$ -means cost is indeed reflected by the different quality of the compressed image.



■ **Figure 1** Compression of an image with  $k = 4$  centers (i.e. colors). Subfigure (a) shows the original image. Subfigure (b) shows a local optimum with a  $k$ -means cost of  $55.18 \cdot 10^8$ . We found this local optimum in runs of Lloyd with uniform initialization and in single runs of  $k$ -means++. Subfigure (c) shows a solution with a  $k$ -means cost of  $43.09 \cdot 10^8$  (for example found by FLS++).



■ **Figure 2** This data set is by Fritzke [16], the illustration by Conrads [9]. The left side shows nine centers sampled by one run of  $k$ -means++ and the corresponding induced clusters are illustrated by colors. The right side shows how the clusters and centers look after running Lloyd's algorithm to convergence with the nine centers from the left as input.

Provided that the cost can actually guide us to a good solution, we know that there are hidden optimal clusters which we desire to find, and that we should improve upon the  $k$ -means cost to find them. The idea of  $k$ -means++ is to find one point from each optimal cluster. This point should be “relatively” good, in the sense that it should allow Lloyd's algorithm to find a good center for the optimum cluster later. We call an optimum cluster *covered* when the  $d^2$ -sampling sampled one of its points. Now, observe that clusters that are *not* covered induce a high cost. By basing the  $d^2$ -sampling probabilities on the cost of points, the algorithm strives to find good points from uncovered optimum clusters. Indeed,  $k$ -means++ succeeds with this goal most of the time, but it can miss some of the optimum clusters – which then remain uncovered because the subsequent Lloyd steps are (in general) unable to shift a mean point from one optimum cluster to another. However,  $k$ -means++ still computes a  $\Theta(\log k)$ -approximation on expectation, and it improves the practical performance of Lloyd's algorithm vastly. Figure 2 illustrates how  $k$ -means++ can end up in a local optimum. We see that out of the nine natural clusters in the data set,  $k$ -means++ has covered eight – the upper middle cluster was hit by two samples, and thus there was no center left in the end to cover the lower right cluster. Running Lloyd's algorithm on this center set then results in nicely placed centers which capture the natural clusters perfectly in seven of nine cases. However, Lloyd's algorithm cannot repair the mistake of not finding the lower right cluster.




(a) Initial converged clustering (without the circle-shaped green point). The circle-shaped green point is sampled, the crossed out center is removed.

(b) After performing the center swap, Lloyd's iterations will converge to the optimal clustering.

■ **Figure 3** Improving solutions with local search steps.

It is now around fifteen years after the development of  $k$ -means++, and there has been increasing interest in the question of improving upon its theoretical and practical properties. There are two main ways to improve the behaviour explained in Figure 2. Option A is to try to avoid missing clusters. It was already proposed in [4] to improve  $k$ -means++ by adding a greedy procedure: When choosing  $c_i$ , compute  $\ell$  candidates, all chosen according to the same probability distribution, and then out of these, pick the point which decreases the current cost the most. This variant of  $k$ -means++ is called *greedy  $k$ -means++*. Although it has been observed as performing better in practice in some instances (as in [4] and in [7]), greedy  $k$ -means++ actually has a *worse* worst-case behavior than  $k$ -means++. Bhattacharya et al. [6] give a family of point sets where the expected quality of solutions computed by greedy  $k$ -means++ is lower bounded by  $\Omega(\ell \log k)$ . This is because points which look beneficial to the greedy procedure may actually be very bad centers, for instance, if points lie “on the boundary” between clusters. If two clusters are uncovered and there is a point in the middle, then choosing this point decreases the overall cost more than any point closer to the mean of the two clusters. However, a point in the middle of two clusters and far away from their means is certainly a bad choice for a cluster center. Recently, [17] showed a better lower bound of  $\Omega(\ell^3 \log^3 k / \log^2(\ell \log k))$  while also giving an upper bound of  $\mathcal{O}(\ell^3 \log^3 k)$ . Still, greedy  $k$ -means++ is for example the default initialization method in Python's scikit learn package [25] and our experiments confirm that this is a justified choice.

Another option is to perform improvement steps on the  $k$ -means++ solution before running Lloyd's algorithm. A prominent technique in the theoretical analysis of the  $k$ -means problem (and other clustering problems) is to do local search by swapping in and out centers. A solution set  $C$  is improved by taking out a constant number  $t$  of centers and replacing them by  $t$  points from  $X$ . [20] show that this approach with sufficiently large  $t$  leads to a  $(9 + \epsilon)$ -approximation for the  $k$ -means problem. However, the approach is fairly slow since even for  $t = 1$ , checking if there is an improving swap takes a lot of time. Lattanzi and Sohler [21] combine local search for  $t = 1$  with adaptive sampling. Instead of checking all  $x \in X$  for improving swaps, they (multiple times) sample a point from  $X$  by adaptive sampling and check if exchanging it for a center in  $C$  improves the solution. They show that starting from an initial  $k$ -means++ solution, this yields a  $\mathcal{O}(1)$ -approximation after  $\mathcal{O}(k \log \log k)$  steps. The resulting algorithm is called LS++. The analysis has since been improved by Choo et al. in [8], where it is shown that, for any  $\epsilon > 0$ , performing  $\epsilon k$  steps yields a  $\mathcal{O}(1/\epsilon^3)$ -approximation. Figure 3 shows an idealized visualization of how improvement steps shall work: In a solution where one optimum cluster got covered twice and other clusters have been clustered together, one wants to find a point from an uncovered optimal cluster with  $d^2$ -sampling and then converge to an optimum solution with Lloyd's algorithm.



■ **Figure 4** An example with eight optimal clusters (green). When swapping in a new center, the best center to delete is the one in the middle: without it, Lloyd’s algorithm can repair the solution.

Our aim in this paper is to improve the practical quality of both greedy  $k$ -means++ and LS++. We start with LS++, but our idea is to *combine* it with the strength of Lloyd’s algorithm instead of simply running local improvement steps ahead of it. We propose *Foresight-LS++* (*FLS++*) combining steps of Lloyd’s algorithm and local search: When exchanging centers, it is not at all clear which centers should be swapped because we do not know the effect of the swap when Lloyd’s algorithm runs later. Does a particular swap really enable Lloyd’s algorithm to find good centers? The best way to answer this question is to actually perform one step of Lloyd’s algorithm for a potential swap. This way, instead of comparing the cost of the solution before and after the swap directly, we add some *foresight*: We perform one step of Lloyd’s algorithm and see how the solution develops depending on the swaps. The candidates for new centers are also found with  $d^2$ -sampling, so the centers we sample follow the same distribution. With our procedure, we may avoid putting in unfavorable centers. However, the biggest impact of our change is finding the best center to *remove* in a more informed way. More precisely, some of the swaps take out superfluous centers in clusters that were hit multiple times. However, there are also beneficial swaps of a different type: Here, we swap out a center that is actually the only one covering a cluster, just to better distribute the centers in that area. Figure 4 shows a schematic illustration of situations where this can be indeed helpful. By running one step of Lloyd’s algorithm, FLS++ can identify both (and potentially more) types of beneficial swaps.

**Further related work.** The  $k$ -means problem is NP-hard even for  $k = 2$  [10] and  $d = 2$  [23], but there exist different constant-factor approximation algorithms. The currently best approximation ratio is 6.357, achieved by Ahmadian et al. [2] with an LP rounding based approach. Prior to their work, the best known approximation ratio of  $9 + \epsilon$  was achieved via a local search approach analyzed by Kanungo et al. [20]. Awasthi et al. [5] show that for general  $k$  and  $d$  there exists some  $c > 0$  such that it is NP-hard to approximate  $k$ -means with factor  $c$ . The constant  $c$  is at least 1.0013 as shown by Lee et al. [22]. The problem gets easier when *bicriterial* solutions are allowed, i.e., when we may select more than  $k$  centers and still compare to the cost of a solution with  $k$  centers. There are multiple bicriterial approximations, and Wei [27] indeed shows that for any constant  $\beta > 1$ , selecting  $\beta k$  centers with  $D^2$ -sampling gives a constant factor approximation (where the constant depends on  $\beta$ ).

There is a vast literature on Lloyd’s algorithm and various speed-up techniques. A good survey on initialization methods is due to Celebi et al. [7] who give an extensive evaluation of different heuristics to initialize Lloyd’s algorithm, including  $k$ -means++ and a greedy variant of it. LS++ was developed after that survey, so it is not included in the comparison. Examples for speed-up techniques are the works by Elkan [12] and Hamerly [18]. These speed up the execution of Lloyd’s algorithm but do not change the outcome.

**Notation.** The distance of a point  $p$  to a set of centers  $C$  is  $d(p, C) = \min_{c \in C} \|p - c\|$ , and for a set of points  $P$  we define the distance to a set of centers as  $d(P, C) = \sum_{p \in P} d(p, C) = \sum_{p \in P} \min_{c \in C} \|p - c\|$ . We refer to the *cost* of  $P$  with respect to center set  $C$  as  $\Phi(P, C) = d^2(P, C) = \sum_{p \in P} d^2(p, C) = \sum_{p \in P} \min_{c \in C} \|p - c\|^2$ . For an arbitrary assignment  $\phi' : P \rightarrow C$  we define  $\Phi(P, C, \phi') = \sum_{p \in P} d^2(p, \phi'(p))$  as the cost of a clustering  $C$  using assignment  $\phi'$ .

## 2 Algorithms

We first describe the  $k$ -means++ algorithm. It is stated in Algorithm 2 and explained below.

### Algorithm 1 $d^2$ -sampling-init.

---

**Input:** Point set  $P \subset \mathbb{R}^d$ , number  $k \in \mathbb{N}$   
**Output:** Center set  $C \subset \mathbb{R}^d$

- 1  $C = \emptyset$
- 2 **for**  $i = 1$  **to**  $k$  **do**
- 3      $c = \text{SampleCenter}(P, C)$
- 4      $C = C \cup \{c\}$
- 5 **end**
- 6 **return**  $C$

---

### Algorithm 2 $k$ -means++.

---

**Input:** Point set  $P \subset \mathbb{R}^d$ ,  
 numbers  $k, s \in \mathbb{N}$   
**Output:** Center set  $C \subset \mathbb{R}^d$

- 1  $C = d^2\text{-sampling-init}(P, k)$
- 2  $C, \phi = \text{Lloyd}(P, C, s)$
- 3 **return**  $C$

---

The  $k$ -means++ algorithm consists of sampling  $k$  initial centers with  $d^2$ -sampling and then heuristically refining the solution by Lloyd's method. For ease of notation in Algorithm 3, the method `LLOYD`( $P, C, s$ ) also outputs the optimal assignment of the points to the centers before the final centroids are computed, see the description of Lloyd's algorithm at the beginning of the introduction. For the sampling part, we assume that we have access to a function `SAMPLECENTER`( $P, C$ ) which, given a point set  $P$  and center set  $C$ , returns a point  $p$  from  $P$  that is sampled with probability  $d^2(p, C)/d^2(P, C)$  (unless  $C = \emptyset$ , in which case the probability is  $1/|P|$ ). This is the original  $k$ -means++ sampling, also referred to as  $d^2$ -sampling. Notice that a sample can be obtained in time  $O(n)$  if all distances  $d^2(p, C)$  are already known: First, add  $d^2(p, C)$  for all  $p \in P$  to obtain the sum  $S$ . Second, draw a uniform sample  $r$  from  $[0, 1]$ . Third, iterate through  $P$  and add up the distances again until the sum exceeds  $r \cdot S$  for the first time; take the point before the point for which this happened. It is important that one needs all values of  $d^2(p, C)$  to do so; one can compute them in time  $O(ndk)$  when needed but it is beneficial to store values. Repeating `SAMPLECENTER`( $P, C$ ) for  $k$  iterations yields the initialization part of  $k$ -means++, see Algorithm 1. We also assume that we have access to a subroutine for Lloyd's method. Given a set  $P$  of input points, a set  $C$  of initial centers and a maximum number of steps  $s$ , the method `LLOYD`( $P, C, s$ ) runs the loop in Step 2 of Lloyd's algorithm as described on page 2 to compute and return a set  $C_{\text{out}}$  of centers. We define  $\Phi(\text{LLOYD}(P, C, s)) = \Phi(P, C_{\text{out}})$ . Given these subroutines,  $k$ -means++ can be described by the pseudo code in Algorithm 2.

**LS++ and FLS++.** Algorithm 3 shows the pseudo code of FLS++. Without Lines 2, 5 and 7 and with directly comparing  $\Phi(P, (C \setminus \{c_{\text{old}}\}) \cup \{c_{\text{new}}\})$  to  $\Phi(P, C^{\text{min}})$  in Line 8, the code describes LS++. The core idea of LS++ is to further improve the initial solution by local search steps before running Lloyd's algorithm. LS++ first samples  $k$  centers with  $d^2\text{-sampling-init}(P, k)$ . Then it does the following  $Z$  times: Pick another point  $c_{\text{new}}$  with  $d^2$ -sampling, then find the best way to swap out a center  $c_{\text{old}}$  and replace it with  $c_{\text{new}}$ . The

---

**Algorithm 3** Foresight-LS++.

---

```

Input: Point set  $P \subset \mathbb{R}^d$ , numbers  $k, s \in \mathbb{N}$ 
Parameters:  $Z \in \mathbb{N}_0$ 
Output: Center set  $C \subset \mathbb{R}^d$ 

1  $C = d^2\text{-sampling}(P, k)$ 
2  $C, \phi = \text{Lloyd}(P, C, 1)$ 
3 for  $i = 1$  to  $Z$  do
4    $c_{\text{new}} = \text{SampleCenter}(P, C)$ 
5    $C^{\text{min}}, \phi^{\text{min}} = \text{Lloyd}(P, C, 1)$ 
6   for  $c_{\text{old}} \in C$  do
7      $C', \phi' = \text{Lloyd}(P, (C \setminus \{c_{\text{old}}\}) \cup \{c_{\text{new}}\}, 1)$ 
8     if  $\Phi(P, C', \phi') < \Phi(P, C^{\text{min}}, \phi^{\text{min}})$  then
9        $C^{\text{min}}, \phi^{\text{min}} = C', \phi'$ 
10    end
11  end
12   $C = C^{\text{min}}$ 
13 end
14  $C, \phi = \text{Lloyd}(P, C, s)$ 
15 return  $C$ 

```

---

parameter  $Z$  is assumed to be at least  $100000k \log \log k$  in the theoretical analysis in [21], but set to values  $\leq 25$  in the practical evaluation [21]. After  $Z$  improvement steps, LS++ calls Lloyd's algorithm.

LS++ does indeed improve the solution quality. But the strength in the original  $k$ -means++ algorithm lies in the fact that it uses the power of Lloyd's to quickly do local refinements:  $d^2$ -sampling itself would not be competitive, it is the combination of finding clusters and refining the centers that makes  $k$ -means++ such a success in practical applications. Consequently, we believe that increasing the symbiosis of sampling and Lloyd's algorithm is the key to improving LS++.

The additional lines of FLS++ call Lloyd's algorithm for one iteration, which means that all centers are replaced by the centroids of their clusters. Line 2 uses this to refine the initial sampling solution, Line 5 makes sure that we always keep centroids (instead of mixing in input points into the center set), and Line 7 is the step that gave the algorithm the prefix *foresight*: Instead of checking the cost of  $(C \setminus \{c_{\text{old}}\}) \cup \{c_{\text{new}}\}$ , we run a Lloyd's step to evaluate how the solution cost will develop after updating the centers. This makes it much easier for FLS++ to identify swaps that remove a superfluous center like the extra center in the upper cluster in Figure 2a.

One can view FLS++ either as a modified LS++ algorithm or as a modified Lloyd's algorithm. Viewed as a modified Lloyd's algorithm one can think of FLS++ as performing Lloyd's algorithm but doing a center swap in every iteration to avoid running into local optima. Viewed as a LS++ modification, one can think of it as running Lloyd's algorithm for 1 step between the local search improvement steps.

Why do we run exactly one step of Lloyd's algorithm? The reason is that this can be done very efficiently. Running multiple Lloyd's steps for *each*  $c_{\text{old}}$  would incur a running time of  $O(ndk^2)$  for *each* local search step. But recomputing the centroids can be done cleverly for all points together such that one local search step in Algorithm 3 takes  $\mathcal{O}(ndk)$  just as for LS++.

► **Lemma 1.** *One iteration of the For-Loop in Lines 3-13 of Algorithm 3 can be implemented such to run in time  $\mathcal{O}(ndk)$ .*

**Proof.** First we compute the distances of every point  $p$  to its closest center  $\phi_1(p)$  and second-closest center  $\phi_2(p)$  in time  $\mathcal{O}(ndk)$ . With this information, we can sample a point  $c_{\text{new}}$  with  $d^2$ -sampling in time  $\mathcal{O}(n)$ . Next, we compute the distances of all points to the sample point  $c_{\text{new}}$ . Then we check all  $k$  current centers whether we want to swap it out. For one candidate  $c_{\text{old}}$ , we run through all  $n$  points  $p$ , and determine which is its closest center. We can do this in time  $\mathcal{O}(1)$  in the following way: if  $\phi_1(p) \neq c_{\text{new}}$ , then we compare  $d(p, \phi_1(p))$  with  $d(p, c_{\text{new}})$ . If  $\phi_1(p) = c_{\text{new}}$ , then we compare  $d(p, \phi_2(p))$  with  $d(p, c_{\text{new}})$ . After doing this for all  $c_{\text{old}}$  and all  $p$ , we have determined the cluster membership of all points in time  $\mathcal{O}(kn)$ . Finally, we want to recompute the centroids and costs of all clusters which we can do in time  $\mathcal{O}(nd)$ . It is important to note that we do not recompute the reassignment of all points to the new set of centers as this would increase the overall running time  $\mathcal{O}(ndk^2)$ . This is delayed until the next iteration and thus only done for the chosen  $C^{\min}$ . ◀

The approximation guarantee is also not affected.

► **Corollary 2.** *FLS++ with  $Z \geq 100000k \log \log k$  computes an  $\mathcal{O}(1)$ -approximation.*

**Proof.** [21] shows that for any  $C$  with cost larger than  $500 \cdot \text{OPT}$  and a  $c_{\text{new}}$  sampled by  $d^2$ -sampling, there is a  $c_{\text{old}}$  such that the cost of  $(C \setminus \{c_{\text{old}}\}) \cup \{c_{\text{new}}\}$  compared to the cost of  $C$  is smaller than that of  $C$  by a fraction of  $(1/100k)$  with probability  $1/1000$  (Lemma 3 in [21]). The analysis of the algorithm follows from this fact. Lloyd’s steps are always improving, since the centroid is always the best center for a cluster. Thus Lemma 3 in [21] still holds and thus does the approximation guarantee. ◀

### 3 Experimental Results

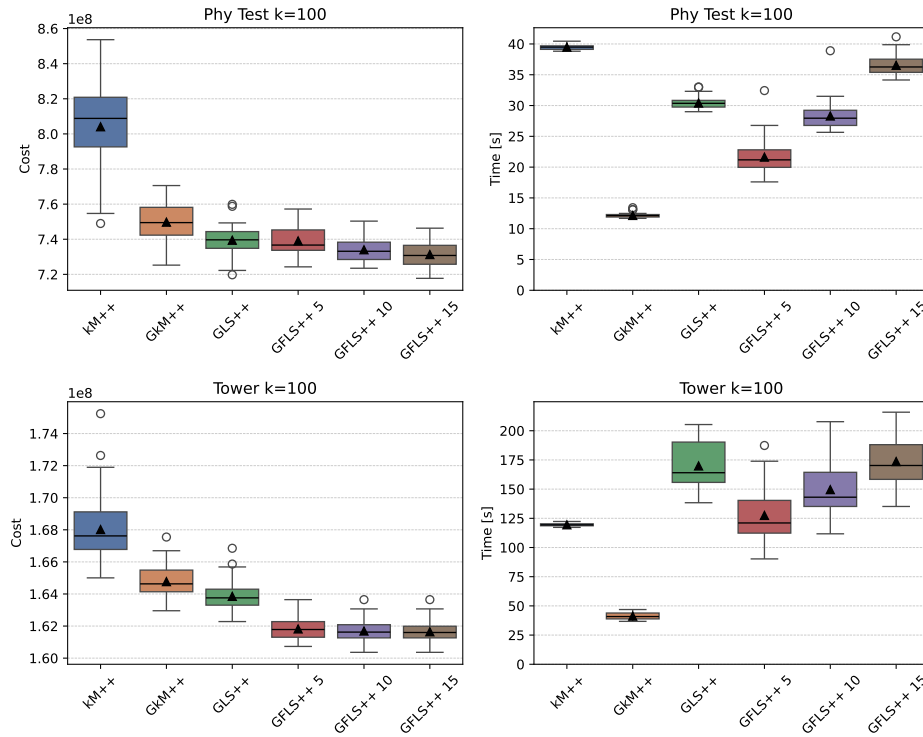
**Setup.** The computations were performed on an Intel(R) Xeon(R) E3-1240 running at 3.7Ghz and 8 cores. The code is written in C++. The code and datasets used can be found at [https://github.com/lukasdrexl/flspp\\_code](https://github.com/lukasdrexl/flspp_code). In the following we abbreviate  $k$ -means++ with KM++. We refer to the greedy variant of KM++ as GKM++ (analogously for (G)LS++ and (G)FLS++).

**Datasets.** Our experiments are based on datasets used in [21], image RGB-data used in [13], some datasets where the optimal solution is known for specific values of  $k$ , and some new synthetically generated datasets similar to *rectangles*.

dataset	number of points	dimension $d$	source	dataset	number of points	dimension $d$	source
<i>rectangles</i>	1296	2	[14]	<i>Tower</i>	4915200	3	[13]
<i>circles</i>	10000	2	-	<i>Clegg</i>	716320	3	[13]
<i>close circles</i>	10000	2	-	<i>Frymire</i>	1235390	3	[13]
<i>pr91</i>	2392	2	[24]	<i>Body measurements</i>	507	5 (reduced)	[19]
<i>D31</i>	3100	2	[14]	<i>Concrete strength</i>	1030	9	[28]
<i>s3</i>	5000	2	[14]	<i>KDD Bio Train</i>	145751	74	[1]
<i>A3</i>	7500	2	[14]	<i>KDD Phy Test</i>	100000	78	[1]

For some datasets, the choice of  $k$  is obvious, as there exists a clear ground truth clustering: *rectangles* from Figure 2 consist of  $k = 36$  clusters; likewise, the two generated datasets *circles* and *close circles* consist of  $k = 100$  separated clusters as visualized in Figure 9 in the Appendix. For other datasets, we evaluated several values of  $k$ : e.g. on *pr91* where the optimum  $k$ -means cost is known for several values of  $k$  (see [3]).





■ **Figure 5** Comparison on two large datasets, for  $R = 50$  repetitions. GLS++ always performs 25 local search steps. For GFLS++, we display the results for performing 5, 10 and 15 such steps.

**Overall performance on big data sets.** To get an idea of the algorithms' strengths and weaknesses, we compare the results of all four algorithms on three large data sets, *KDD Bio Train*, *KDD Phy Test*, and *Tower*. In Figure 5, we compare KM++, GKM++, GLS++ with 25 local search steps (which is the number chosen in [21]) and GFLS++ with 5, 10, and 15 steps. Since GFLS++ performs one Lloyd iteration in every local search step, 25 such steps take longer than 25 steps in GLS++. Hence, we would like to compare how GFLS++ performs when using *fewer* local search steps. We take a more detailed look at the trade-off between runtime and cost further below. We perform 50 *runs* of each configuration, i.e., each algorithm is called 50 times with a specific  $k$  on each dataset. The initial centers for GLS++ and GFLS++ are chosen by greedy  $d^2$ -sampling. We execute GKM++, GLS++ and GFLS++ with the same initial set of centers.

As is to be expected, a larger value of  $Z$  increases runtime but decreases cost. On both *Phy Test* and *Tower*, GFLS++ with  $Z = 15$  obtains the overall minimum cost at the expense of a longer runtime, while outperforming GLS++ slightly at  $Z = 10$  in terms of cost *and* runtime. Table 1 shows the maximum, average and minimum cost that the algorithms produced on each data set with  $k = 100$ . While the overall differences are not too big, GFLS++ consistently obtains the lowest values. Runtimes and performance comparisons for different values of  $k$  and for the *KDD Bio Train* dataset can be seen in Section C in the appendix.

**Best of repeated runs within a time limit.** As we have seen in the previous section, solutions computed by GFLS++ have a smaller cost than GKM++ or GLS++ on average, but need more time to terminate. In the following we want to analyze if running KM++

or LS++ multiple times can yield better results than GFLS++ in the same time. Thus, we repeat each algorithm until a time limit is reached and for each algorithm only the best solution is returned.

We choose some iteration value  $B \in \mathbb{N}$  and report the best solution found by exactly  $B$  repetitions of GFLS++. We then repeat (G)KM++ and GLS++ as long as their respective elapsed time is less or equal to the time used by FLS++; again, we return the best solution found.

The result of this comparison is shown in Table 2. It shows the numerical values of the final average cost of each algorithm. Additionally we evaluated the relative cost difference when compared to GFLS++. For two algorithms  $A_1, A_2 \in \{\text{KM++}, \text{LS++}, \text{FLS++}\}$  and their respective average costs  $c_{A_1}, c_{A_2} \in \mathbb{R}_{>0}$  we define their respective percentage cost difference as  $C(A_1, A_2) := (1 - c_{A_1}/c_{A_2}) \cdot 100\%$ . On data set *circles* the cost difference was the highest with approximately 22.37% and 8.8% when we compare the final average cost of FLS++ to KM++ and LS++. This large difference also shows that even if clusters are well separated, LS++ might still fail to find an improvement if the optimal center centroids are not close to the actual data points. In contrast, FLS++ can more often find an improvement because it evaluates the actual optimal centers for a new choice of centers, which brings the new centers closer to the actual optimal centroids. For the other datasets the relative difference in cost is not large, but a positive amount which, if we take into account the number of times each algorithms returned the smallest cost, can be achieved with large success probability.

**Performance over time.** We compare the best solutions found by all three algorithms and their average cost progression over the time. We test this procedure  $R \in \mathbb{N}$  times, where the  $r$ -th run being defined as:

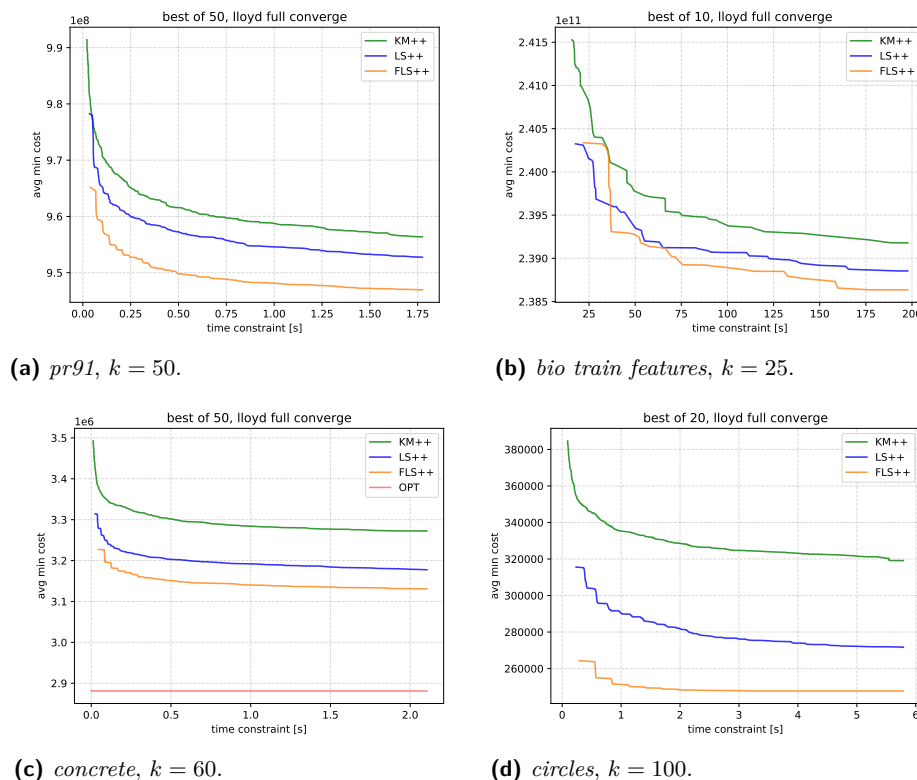
- Run GFLS++  $B$  times, let  $t_B^r$  be the used time until termination in run  $r$ .
  - Repeat GKM++ and GLS++ each as long as their respective elapsed time is at most  $t_B^r$ .
- For run  $r \in [R]$  let  ${}^r c_A^t$  for  $A \in \{\text{KM++}, \text{LS++}, \text{FLS++}\}$  be the current minimum cost of algorithm  $A$  at time  $t \in \mathbb{R}_{\geq 0}$  and run  $r$ . Let  $t_{\min}$  be the first point in time where  $A$  did terminate in every run  $r \in [R]$ , i.e.,  ${}^r c_A^t$  is always defined for every time  $t \geq t_{\min}$ . The average cost of any algorithm  $A$  after  $R$  runs in time step  $t \geq t_{\min}$  is defined as  $\text{AVG}(A, t) := \frac{1}{R} \sum_{r=1}^R {}^r c_A^t$ . As before, each algorithm starts in the  $r$ -th run with the same set of centers from the initialization process.

■ **Table 1** Average cost comparison on large data sets with  $k = 100$  and 50 runs.

Datasets	Algorithms						
	KM++	GKM++	LS++	FLS++ 5	FLS++ 10	FLS++ 15	
Bio Train	<b>Min</b>	1.6292E+11	1.6258E+11	1.6244E+11	1.6116E+11	1.6116E+11	<b>1.6110E+11</b>
	<b>Mean</b>	1.6487E+11	1.6311E+11	1.6302E+11	1.6162E+11	1.6162E+11	<b>1.6160E+11</b>
	<b>Max</b>	1.6937E+11	1.6379E+11	1.6357E+11	1.6240E+11	<b>1.6219E+11</b>	<b>1.6219E+11</b>
Phy Test	<b>Min</b>	7.4898E+08	7.2525E+08	7.1975E+08	7.2427E+08	7.2348E+08	<b>7.1773E+08</b>
	<b>Mean</b>	8.0406E+08	7.4972E+08	7.3942E+08	7.3905E+08	7.3400E+08	<b>7.3128E+08</b>
	<b>Max</b>	8.5367E+08	7.7058E+08	7.5988E+08	7.5721E+08	7.5032E+08	<b>7.4630E+08</b>
Tower	<b>Min</b>	1.6500E+08	1.6296E+08	1.6228E+08	1.6073E+08	<b>1.6036E+08</b>	<b>1.6036E+08</b>
	<b>Mean</b>	1.6802E+08	1.6477E+08	1.6385E+08	1.6182E+08	1.6169E+08	<b>1.6164E+08</b>
	<b>Max</b>	1.7524E+08	1.6755E+08	1.6685E+08	<b>1.6365E+08</b>	<b>1.6365E+08</b>	<b>1.6365E+08</b>

■ **Table 2** Average cost. Last two columns show  $(1 - c_{\text{FLS}++}/c_B) \cdot 100\%$  for  $B \in \{\text{KM}++, \text{LS}++\}$ .

data set	$c_{\text{KM}++}$	$c_{\text{LS}++}$	$c_{\text{FLS}++}$	$C(\text{FLS}++, \text{KM}++)$	$C(\text{FLS}++, \text{LS}++)$
<i>pr91</i>	9.5637E+08	9.5276E+08	9.4696E+08	0.98%	0.61%
<i>bio train features</i>	2.3918E+11	2.3885E+11	2.3864E+11	0.23%	0.09%
<i>concrete</i>	3.2724E+06	3.1775E+06	3.1308E+06	4.33%	1.47%
<i>circles</i>	3.1913E+05	2.7164E+05	2.4773E+05	22.37%	8.8%



■ **Figure 6** Comparison of average cost decrease depending on the runtime of FLS++.

Figure 6 shows the progress of each algorithm over time. For LS++, we use  $Z = 25$  as number of repeat steps, as in the original experiments in [21]. We use the same value for FLS++ in the following experiments if not specified otherwise. We repeat Lloyd steps, i.e., computing memberships of all points to their closest center and recomputing the new optimal centers for this clustering, until two consecutive iterations produce solutions  $C_1, C_2$  with  $1 - \frac{\Phi(P, C_2)}{\Phi(P, C_1)} < 0.0001$ . All plots are averaged over 100 iterations, i.e.,  $R = 100$ , except for (b) which is averaged over 20 iterations. In all plots, FLS++ achieves the smallest final average cost and often beats the other algorithms over the entire time frame. Similarly, LS++ is consistently better than KM++. In (c) we also compare with the known optimal solution OPT: Despite the FLS++ improvements, there is still a significant gap.

**Number of best solutions found.** Table 3 shows how many times KM++, LS++ or FLS++ computed the overall best solution for a given dataset within the time limit; most of the time FLS++ did return the smallest solution. Despite FLS++ having the fewest average number of iterations due to its longer running time, we can see that most of the time it

■ **Table 3** Respective number of times each algorithm returned the smallest solution for the respective image in Figure 6 and how many iterations we could repeat the algorithm until the time limit was reached. We only count a win if the respective algorithm returned the strictly smallest cost compared to the other two algorithms in this run.

data set	#Wins KM++	#Avg. iterations KM++	#Wins LS++	#Avg. iterations LS++	#Wins FLS++	#Avg iterations FLS++
<i>pr91</i>	5	160.49	8	64.42	87	50
<i>bio train features</i>	2	19.3	6	13.7	12	10
<i>concrete</i>	0	357.02	4	103.22	96	50
<i>circles</i>	0	103.73	0	29.13	94	20

did report the best solution. In the remaining cases, the best solution was mostly found by LS++, while KM++, having the most number of average iterations for each dataset, could only report the smallest solution twice, namely in the large dataset *bio train features*.

### 3.1 Greedy vs. non-greedy $d^2$ -sampling

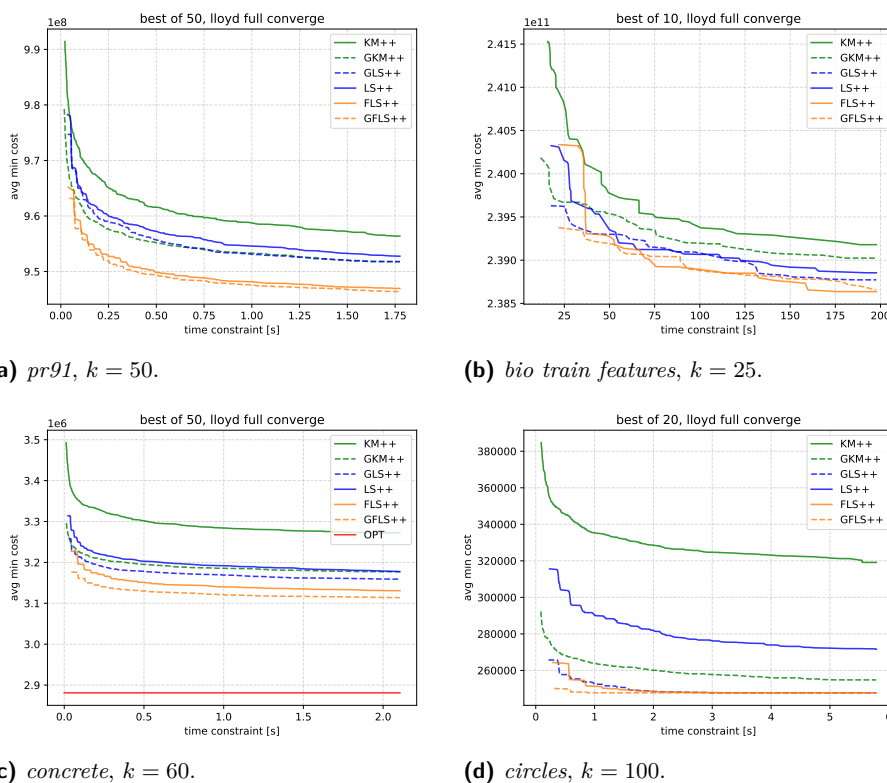
Although greedy  $d^2$ -sampling yields a bad theoretical approximation ratio [6], it is commonly used in practice, e.g. it is the default in the scikit-learn-library with parameter  $\ell = 2 + \lfloor \ln(k) \rfloor$ . In this section, we investigate whether the practical choice is justified and how LS++/FLS++ changes when using greedy sampling.

In Figure 7 we compare the average performance of all algorithms when using *greedy  $d^2$ -sampling* vs. *standard  $d^2$ -sampling*. Here, on almost all evaluated data sets, the average cost decrease is larger when using *greedy  $d^2$ -sampling* regardless of which algorithm is chosen afterwards. Especially KM++ benefits from using *greedy  $d^2$ -sampling*. LS++ and FLS++ benefit from *greedy  $d^2$ -sampling* most of the time, and by roughly the same amount. Using *greedy  $d^2$ -sampling* on dataset *concrete* in (c) brings each algorithm closer to the optimal solution but a gap remains. An analysis of the greedy initialization presented in the same way as in Section 3 as well as an analysis how the initialization process did improve the average cost of the returned solution can be found in section B in the Appendix.

**Varying the number of local search steps.** Running LS++ and FLS++ with the same number of local search steps expectedly leads to FLS++ performing better in terms of cost, but worse in terms of runtime. But Figure 5 already hints at the possibility that using fewer local search steps in FLS++ might lead to better cost *and* better runtime. To this end, we fix the number of local search iterations of LS++ to 25 (guided by the experiments in [21]), while trying out different values for FLS++.

More precisely, we let the algorithms perform 50 runs in total, where one run consists of the following calls. First, we call LS++ with  $Z = 25$ , and afterwards FLS++ for all  $Z \in \{5, \dots, 20\}$ . Importantly, we again ensure that each call in the run starts with the same set of initial centers. After completing all 50 runs, we compute the average cost and runtime of the LS++ calls across all runs, and similarly the averages for all FLS++ calls. The results are shown in Figure 8. The red line indicates the average cost (resp. time) of LS++ with  $Z = 25$  over all 50 runs. Each point on the blue curve corresponds to the average cost (resp. time) of 50 calls to FLS++ with a certain number of local search iterations. We see that for FLS++ even very small values of  $Z$  already suffice to beat LS++ in terms of average cost. At the same time, these small  $Z$  lead to a shorter runtime, making it possible to beat LS++ both in terms of cost and time simultaneously.

To ensure that LS++ could not theoretically improve by, given some time limit, rerunning it multiple times we use a similar setup as in Section 3 to analyze the average performance for some given time bound in the Appendix, see Figure 12. In this case we use the time limit







■ **Figure 7** Comparison of average cost decrease depending on the runtime of FLS++. Dotted lines corresponds to the greedy variant of the original algorithm.

given by FLS++ for  $B = 25$  for all algorithms. We also fixed  $Z = 25$  for LS++ and GLS++ but varied the number of local search steps for FLS++ and GFLS++ as  $Z \in \{5, 10, 15, 20\}$ . Even with fewer local search steps in FLS++ or GLS++ we get on average a smaller cost compared to LS++ or GLS++. For two considered datasets *concrete* and *pr91*, choosing  $Z = 15$  was always enough to get approximately the same cost. For *pr91*, even choosing  $Z = 5$  did result in a better solution than LS++ or GLS++ on average for the specified time limit.

## 4 Conclusion

We propose the new algorithm FLS++ for the  $k$ -means problem: The algorithm uses *foresight* by combining local search with  $d^2$ -sampling and can outperform the established methods of  $k$ -means++ and LS++ in terms of both quality and running time. Even though FLS++ only differs slightly from LS++, our experiments suggest that performing just one Lloyd iteration in each local search step often leads to better performance in terms of cost as well as runtime. Additionally, we investigate if the popularity of *greedy  $d^2$ -initialization* in practice is justified even though it performs poorly in theory. Surprisingly, it turns out that on most data sets considered here, all standard algorithms (as well as our new algorithm) perform substantially better when using greedy initialization.

On the other hand, running our algorithm *without* greedy initialization offers a robust way to recover the quality of the greedy initialization without sacrificing running time.

(a) KDD Phy Test Cost,  $k = 50$ .(b) KDD Phy Test Time,  $k = 50$ .(c) Tower Cost,  $k = 40$ .(d) Tower Time,  $k = 50$ .

■ **Figure 8** Impact on cost and time of increasing the number of local search steps in FLS++. The dashed red line shows average cost (resp. time) of 50 runs of LS++ with 25 local search steps.

## References


- 1 Available at <https://www.kdd.org/kdd-cup/view/kdd-cup-2004/data>.
- 2 Sara Ahmadian, Ashkan Norouzi-Fard, Ola Svensson, and Justin Ward. Better guarantees for k-means and Eucl. k-median by primal-dual alg. *SIAM J. Comput.*, 49(4), 2020.
- 3 Daniel Aloise, Pierre Hansen, and Leo Liberti. An improved column generation algorithm for minimum sum-of-squares clustering. *Mathematical Programming*, 131:195–220, 2012.
- 4 David Arthur and Sergei Vassilvitskii. K-means++: The advantages of careful seeding. In *Proceedings of the 18th SODA*, pages 1027–1035, USA, 2007.
- 5 Pranjal Awasthi, Moses Charikar, Ravishankar Krishnaswamy, and Ali Kemal Sinop. The hardness of approximation of Euclidean k-means. In Lars Arge and János Pach, editors, *Proc. of the 31st SoCG*, volume 34 of *LIPICs*, pages 754–767, 2015.
- 6 Anup Bhattacharya, Jan Eube, Heiko Röglin, and Melanie Schmidt. Noisy, greedy and not so greedy k-means++. In *Proc. of the 28th ESA*, 2020.
- 7 M. Emre Celebi, Hassan A. Kingravi, and Patricio A. Vela. A comparative study of efficient initialization methods for the k-means clustering algorithm. *Expert Syst. Appl.*, 40(1):200–210, 2013.
- 8 Davin Choo, Christoph Grunau, Julian Portmann, and Václav Rozhon. k-means++: few more steps yield constant approximation. In *International Conference on Machine Learning*, pages 1909–1917. PMLR, 2020.
- 9 Theo Conrads. Lokale Such- und Samplingmethoden für das k-Means- und k-Median-Problem. Master’s thesis, Universität zu Köln, 2021.
- 10 Sanjoy Dasgupta. The hardness of k-means clustering, 2008. Technical report.
- 11 Lukas Drexler, Joshua Könen, Daniel R. Schmidt, Melanie Schmidt, and Giulia Baldini. algo-hhu/FLSpp. Software, swuid: `swh:1:dir:39136777e542456572a51515813b6cb377ed0940` (visited on 2024-06-27). URL: <https://github.com/algo-hhu/FLSpp>.

- 12 Charles Elkan. Using the triangle inequality to accelerate k-means. In Tom Fawcett and Nina Mishra, editors, *Proc. of the 20th ICML*, pages 147–153, 2003.
- 13 Gereon Frahling and Christian Sohler. A fast k-means implementation using coresets. In *Proc. of the 22nd SoCG*, pages 135–143, 2006.
- 14 Pasi Fränti and Sami Sieranoja. K-means properties on six clustering benchmark datasets. *Appl. Intell.*, 48(12):4743–4759, 2018.
- 15 Pasi Fränti and Sami Sieranoja. How much can k-means be improved by using better initialization and repeats? *Pattern Recognition*, 93:95–112, 2019.
- 16 Bernd Fritzsche. The k-means-u\* algorithm: non-local jumps and greedy retries improve k-means++ clustering. *CoRR*, abs/1706.09059, 2017.
- 17 Christoph Grunau, Ahmet Alper Özüdoğru, Václav Rozhoň, and Jakub Tětek. A nearly tight analysis of greedy k-means++. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1012–1070. SIAM, 2023.
- 18 Greg Hamerly. Making k-means even faster. In *SDM*, pages 130–140. SIAM, 2010.
- 19 Grete Heinz, Louis J. Peterson, Roger W. Johnson, and Carter J. Kerk. Exploring relationships in body dimensions. *Journal of Statistics Education*, 11(2), 2003.
- 20 Tapas Kanungo, David M. Mount, Nathan S. Netanyahu, Christine D. Piatko, Ruth Silverman, and Angela Y. Wu. A local search approximation algorithm for k-means clustering. *Comput. Geom.*, 28(2-3):89–112, 2004.
- 21 Silvio Lattanzi and Christian Sohler. A better k-means++ algorithm via local search. In *Proc. of the 36th ICML*, volume 97 of *Proceedings of Machine Learning Research*, pages 3662–3671. PMLR, 09–15 June 2019.
- 22 Euiwoong Lee, Melanie Schmidt, and John Wright. Improved and simplified inapproximability for k-means. *Inf. Process. Lett.*, 120:40–43, 2017.
- 23 Meena Mahajan, Prajakta Nimbhorkar, and Kasturi R. Varadarajan. The planar k-means problem is np-hard. *Theor. Comput. Sci.*, 442:13–21, 2012.
- 24 Manfred Padberg and Giovanni Rinaldi. A branch-and-cut algorithm for the resolution of large-scale symmetric traveling salesman problems. *SIAM Review*, 33(1):60–100, 1991.
- 25 F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- 26 J.M Peña, J.A Lozano, and P Larrañaga. An empirical comparison of four initialization methods for the k-means algorithm. *Pattern Recognition Letters*, 20(10):1027–1040, 1999.
- 27 Dennis Wei. A constant-factor bi-criteria approximation guarantee for k-means++. In *Advances in Neural Information Processing Systems*, volume 29, 2016.
- 28 I.-C. Yeh. Modeling of strength of high-performance concrete using artificial neural networks. *Cement and Concrete Research*, 28(12):1797–1808, 1998.

## Appendix

### A Further datasets and cumulative costs

In the following we also show the performance of each algorithm with *normal  $d^2$ -sampling* and *greedy  $d^2$ -sampling* for some other datasets. Like in sections 3 and 3.1 we evaluate the performances of each algorithm using only *normal  $d^2$ -sampling* and then with *greedy  $d^2$ -sampling* by checking f.e. which algorithm did return the best solution in every run or the cost factors between the algorithms. Figure 10 shows the development of the average found minimum cost for each algorithm for 6 other datasets. Each dataset was averaged over 100 iterations, except for the dataset *frymire*, which is averaged over 20 iterations.



(a) Circles dataset.

(b) Close circles dataset.

■ **Figure 9** Two artificially generated datasets also used in the experiments.

■ **Table 4** Respective number of times each algorithm returned the smallest solution for the respective image in Figure 10 and how many iterations we could repeat the algorithm until the time limit was reached. We only count a win if the respective algorithm returned the strictly smallest cost compared to the other two algorithms in this run.

data set	#Wins KM++	#Avg. iterations KM++	#Wins LS++	#Avg. iterations LS++	#Wins FLS++	#Avg iterations FLS++
<i>close circles</i>	0	134.27	0	57.68	100	50
<i>D31</i>	0	229.98	2	73.55	3	50
<i>frymire</i>	0	42.6	5	14.05	15	10
<i>rectangles</i>	0	453.67	0	88.66	0	50
<i>s3</i>	0	127.15	3	58.58	97	50
<i>Body measurements</i>	0	401.23	0	108.15	100	50

We can see in Figure 10 almost the same image as for the previous cases in section 3. For datasets *D31* and *rectangles* we additionally see that all algorithms except for KM++ almost never fail to find the optimal solution in a short amount of time.

## B Cost improvement through greedy initialization

Now we want to consider how using *greedy d<sup>2</sup>-sampling* did improve the average costs after reaching the time limits in section 3.1.


For  $A \in \{\text{FLS++}, \text{LS++}, \text{KM++}\}$  we define their average costs  $c_A, c_A^G$  where the superscript G indicates if we use *greedy d<sup>2</sup>-sampling*. In our case we compare both average costs over all runs  $r \in [R]$  when using time limit  $t_B^r$ . Lastly, similar to the analysis above, we define the improvement factor as  $(1 - \frac{c_A^G}{c_A}) \cdot 100\%$ . As we can see in Table 7 for most data sets and algorithms the improvement is not below 0% and most of the time larger than 0%.

At last we analyse the number of times each algorithm did report the strictly best solution, the number of iterations, average costs and cost factor when compared to GFLS++. Like in the case when using *normal d<sup>2</sup>-sampling* we can see in Tables 5 and 8 that GFLS++ still on average computes the solution with the smallest cost and does so with large probability. In case of the dataset *circles*, now GLS++ and GFLS++ computed the optimal solution after the timelimit in every round. But we can see in Figure 7 that GFLS++ does so faster than GLS++.




■ **Table 5** Respective number of times each algorithm returned the smallest solution for the respective image in Figure 7 and how many iterations we could repeat the algorithm until the time limit was reached. We only count a win if the respective algorithm returned the strictly smallest cost compared to the other two algorithms in this run.


data set	#Wins GKM++	#Avg. iterations GKM++	#Wins GLS++	#Avg. iterations GLS++	#Wins GFLS++	#Avg. iterations GFLS++
<i>pr91</i>	5	160.39	8	62.4	87	48.14
<i>bio train features</i>	2	20.9	6	12.95	12	10.1
<i>concrete</i>	0	295.71	0	95.44	100	48.25
<i>circles</i>	0	109.96	0	29.33	0	19.86
<i>close circles</i>	3	127.14	2	55.47	95	48.46
<i>D31</i>	0	250.55	1	69.75	0	47.3
<i>frymire</i>	5	41.15	2	13.55	13	10
<i>rectangles</i>	0	442.31	0	83.12	0	48.12
<i>s3</i>	4	131	6	57.38	89	48.38
<i>Body measurements</i>	0	336.77	0	101.81	100	47.72




(a) *close circles*,  $k = 100$ .




(b) *D31*,  $k = 31$ .




(c) *frymire*,  $k = 20$ .



(d) *rectangles*,  $k = 36$ .



(e) *s3*,  $k = 50$ .



(f) *Body measurements*,  $k = 50$ .

■ **Figure 10** Comparison of average cost decrease depending on the runtime of FLS++. Dotted lines corresponds to the greedy variant of the original algorithm.

■ **Table 6** Respective average cost of each algorithm. Last two columns show  $(1 - c_{FLS++}/c_A) \cdot 100\%$  for  $A \in \{KM++, LS++\}$ .

data set	$c_{KM++}$	$c_{LS++}$	$c_{FLS++}$	$C(FLS++, KM++)$	$C(FLS++, LS++)$
<i>close circles</i>	6.0468E+05	6.0434E+05	5.9341E+05	1.86%	1.81%
<i>D31</i>	3447.57	3393.26	3393.26	1.58%	0%
<i>frymire</i>	1.3378E+09	1.3287E+09	1.3222E+09	1.17%	0.49%
<i>rectangles</i>	1.5	1.46	1.46	2.56%	0%
<i>s3</i>	6.2325E+12	6.1887E+12	6.1397E+12	1.49%	0.79%
<i>Body measurements</i>	14894.11	14652.64	14427.42	3.13%	1.54%


■ **Table 8** Respective average cost of each algorithm. Last two columns show  $(1 - c_{GFLS++}/c_A) \cdot 100\%$  for  $A \in \{GKM++, GLS++\}$ .

data set	$c_{GKM++}$	$c_{GLS++}$	$c_{GFLS++}$	$C(GFLS++, GKM++)$	$C(GFLS++, GLS++)$
<i>pr91</i>	9.518E+08	9.5177E+08	9.4638E+08	0.57%	0.57%
<i>bio train features</i>	2.3902E+11	2.3877E+11	2.3866E+11	0.15%	0.05%
<i>concrete</i>	3.1772E+06	3.1591E+06	3.1138E+06	2%	1.43%
<i>circles</i>	2.5481E+05	2.4773E+05	2.4773E+05	2.78%	0%
<i>close circles</i>	6.0113E+05	6.0244E+05	5.9485E+05	1.04%	1.26%
<i>D31</i>	3393.26	3393.26	3393.26	0%	-0%
<i>frymire</i>	1.33E+09	1.3321E+09	1.3244E+09	0.42%	0.58%
<i>rectangles</i>	1.46	1.46	1.46	0%	0%
<i>s3</i>	6.1786E+12	6.1742E+12	6.137E+12	0.67%	0.6%
<i>Body measurements</i>	14637.51	14584.68	14357.35	1.91%	1.56%


■ **Table 7** Average cost decrease by using *greedy d<sup>2</sup>-sampling*.

data set	$C(\text{GKM}++, \text{KM}++)$	$C(\text{GLS}++, \text{LS}++)$	$C(\text{GFLS}++, \text{FLS}++)$
<i>pr91</i>	0.48%	0.1%	0.06%
<i>bio train features</i>	0.07%	0.03%	-0.01%
<i>concrete</i>	2.91%	0.58%	0.54%
<i>circles</i>	20.16%	8.8%	0%
<i>close circles</i>	0.59%	0.32%	-0.24%
<i>D31</i>	1.58%	0%	0%
<i>frymire</i>	0.59%	-0.26%	-0.17%
<i>rectangles</i>	2.56%	0%	0%
<i>s3</i>	0.87%	0.23%	0.04%
<i>Body measurements</i>	1.72%	0.46%	0.49%


**C** Local search remaining Figures




(a) cost developments LS++.



(b) cost developments GLS++.




(c) cost developments FLS++.




(d) cost developments GFLS++.

■ **Figure 11** Dataset *concrete*. Comparison of the average cost decrease for (G)LS++ and (G)FLS++ for  $k = 60$ ,  $R = 100$  and varying values of  $Z$ .

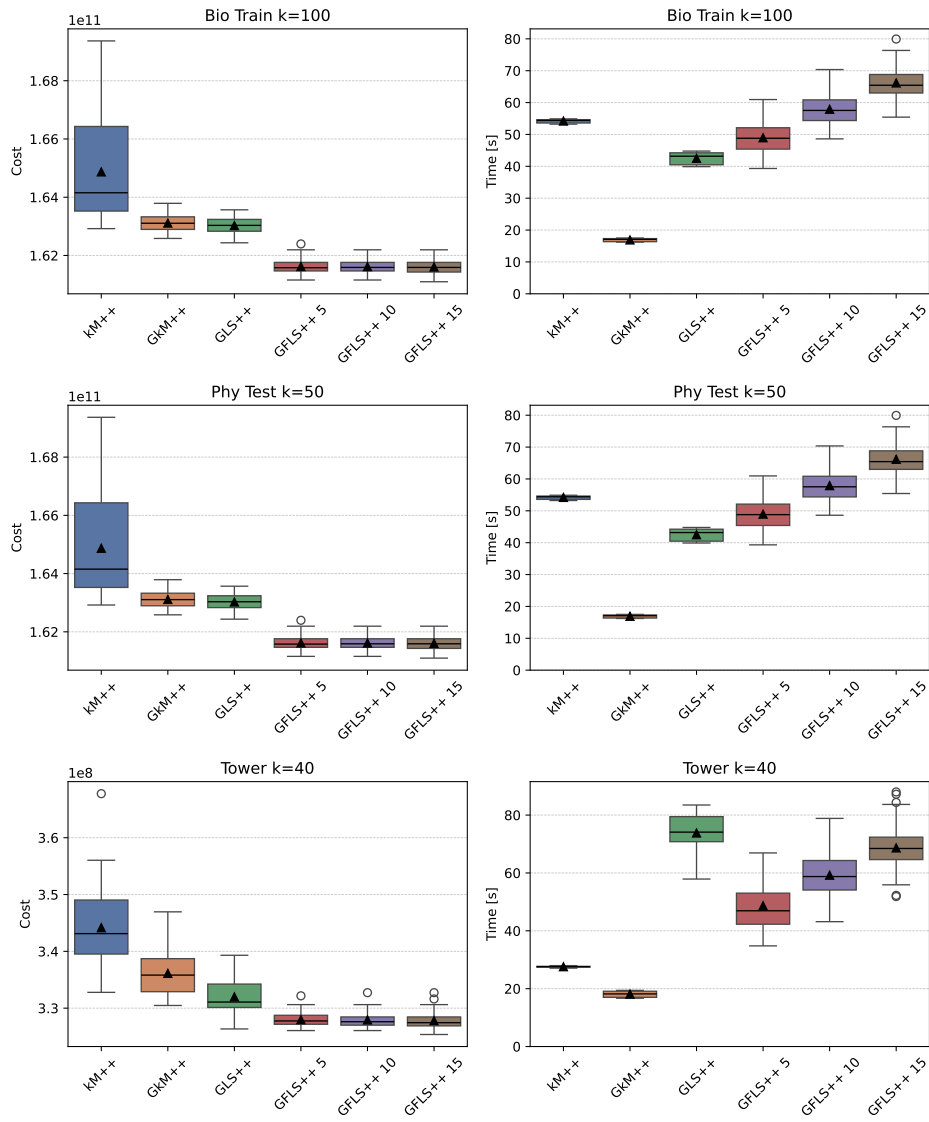


(a) Dataset *concrete*,  $k = 60$ .



(b) Dataset *pr91*,  $k = 50$ .

■ **Figure 12** Fixing number of local search steps for LS++ and GLS++ while using a variable number for FLS++ and GFLS++. The number after either FLS++ or GFLS++ represents the value  $Z$  for this algorithm.



■ **Figure 13** Performance boxplots for different values of  $k$  and one additional dataset.