

# Net and Prune: A Linear Time Algorithm for Euclidean Distance Problems\*

Sariel Har-Peled<sup>†</sup>      Benjamin Raichel<sup>‡</sup>

September 18, 2015

## Abstract

We provide a general framework for getting expected linear time constant factor approximations (and in many cases FPTAS's) to several well known problems in Computational Geometry, such as  $k$ -center clustering and farthest nearest neighbor. The new approach is robust to variations in the input problem, and yet it is simple, elegant, and practical. In particular, many of these well studied problems which fit easily into our framework, either previously had no linear time approximation algorithm, or required rather involved algorithms and analysis. A short list of the problems we consider include farthest nearest neighbor,  $k$ -center clustering, smallest disk enclosing  $k$  points,  $k$ th largest distance,  $k$ th smallest  $m$ -nearest neighbor distance,  $k$ th heaviest edge in the MST and other spanning forest type problems, problems involving upward closed set systems, and more. Finally, we show how to extend our framework such that the linear running time bound holds with high probability.

## 1. Introduction

In many optimization problems one is given a geometric input and the goal is to compute some real valued function defined over it. For example, for a finite point set, the function might be the distance between the closest pair of points. For such a function, there is a search space associated with the function, encoding possible candidate values, and the function returns the minimum value over all possible solutions. For the closest pair problem, the search space is the set of all pairs of input points and their associated distances, and the target function computes the minimum pairwise distance in this space. As such, computing the value of such a function requires searching for the optimal solution, and returning its value; that is, *optimization*. Other examples of such functions include:

- (A) The minimum cost of a clustering of the input. Here, the search space is the possible partitions of the data into clusters together with their associated prices.
- (B) The price of a minimum spanning tree. Here, the search space is the set of all possible spanning trees of the input (and their associated prices).

---

\*Work on this paper was partially supported by NSF AF awards CCF-0915984 and CCF-1217462. The full updated version of this paper is available online [HR14]. A preliminary version of this paper appeared in STOC 2012 [HR13].

<sup>†</sup>Department of Computer Science; University of Illinois; 201 N. Goodwin Avenue; Urbana, IL, 61801, USA; sariel@uiuc.edu; <http://illinois.edu/~sariel>.

<sup>‡</sup>Department of Computer Science; University of Illinois; 201 N. Goodwin Avenue; Urbana, IL, 61801, USA; raichel2@uiuc.edu; <http://illinois.edu/~raichel2>.

- (C) The radius of the smallest enclosing ball. Here, the search space is the set of all possible balls that enclose the input and their radii (i.e., it is not necessarily finite).

Naturally, many other problems can be described in this way.

It is often possible to construct a decision procedure, for such an optimization problem, which given a query value can decide whether the query is smaller or larger than the value of the function, without computing the function explicitly. Naturally, one would like to use this decider to compute the function itself by (say) performing a binary search, to compute the minimum value which is still feasible (which is the value of the function). However, often this is inherently not possible as the set of possible query values is a real interval (which is not finite). Instead, one first identifies a finite set of critical values which must contain the optimum value, and then searches over these values.

Naturally, searching over these values directly can be costly as often the number of such critical values is much larger than the desired running time. Instead one attempts to perform an implicit search.

One of the most powerful techniques to solve optimization problems efficiently in Computational Geometry, using such an implicit search, is the *parametric search* technique of Megiddo [Meg83]. It is relatively complicated, as it involves implicitly extracting values from a simulation of a parallel decision procedure (often a parallel sorting algorithm). For this reason it is inherently not possible for parametric search to lead to algorithms which run faster than  $O(n \log n)$  time. Nevertheless, it is widely used in designing efficient geometric optimization algorithms, see [AST94, Sal97]. Luckily, in many cases one can replace parametric search by simpler techniques (see prune-and-search below for example) and in particular, it can be replaced by randomization, see the work by van Oostrum and Veltkamp [vOV04]. Another example of replacing parametric search by randomization is the new simplified algorithm for the Fréchet distance [HR11]. Surprisingly, sometimes these alternative techniques can actually lead to linear time algorithms.

**Linear time algorithms** There seems to be three main ways to get linear time algorithms for geometric optimization problems (exact or approximate):

- (A) **Coreset/sketch.** One can quickly extract a compact sketch of the input that contains the desired quantity (either exactly or approximately). As an easy example, consider the problem of computing the axis parallel bounding box of a set of points - an easy linear scan suffices. There is by now a robust theory of what quantities one can extract a coreset of small size for, such that one can do the (approximate) calculation on the coreset, where usually the coreset size depends only on the desired approximation quality. This leads to many linear time algorithms, from shape fitting [AHV04], to  $(1 + \varepsilon)$ -approximate  $k$ -center/median/mean clustering [Har01, Har04a, HM04, HK07] in constant dimension, and many other problems [AHV05]. The running times of the resulting algorithms are usually  $O(n + \text{func}(\text{sketch size}))$ . The limitation of this technique is that there are problems for which there is no small sketch, from clustering when the number of clusters is large, to problems where there is no sketch at all [Har04b] – for example, for finding the closest pair of points one needs all the given input and no sketching is possible.
- (B) **Prune and search.** Here one prunes away a constant fraction of the input, and continues the search recursively on the remaining input. The paramount example of such an algorithm is the linear time median finding algorithm, however, one can interpret many of the randomized algorithms in Computational Geometry as pruning algorithms [Cla88, CS89, Mul94]. For example, linear programming in constant dimension in linear time [Meg84], and its extension to LP-type problems [SW92, MSW96]. Intuitively, LP-type problems include low dimensional convex programming (a standard example is the smallest enclosing ball of a point set in constant dimension). However, surprisingly, such problems also include problems that are not convex in

nature – for example, deciding if a set of (axis parallel) rectangles can be pierced by three points is an LP-type problem. Other examples of prune-and-search algorithms that work in linear time include (i) computing an ear in a triangulation of a polygon [EET93], (ii) searching in sorted matrices [FJ84], and (iii) ham-sandwich cuts in two dimensions [LMS94]. Of course, there are many other examples of using prune and search with running time that is super-linear.

- (C) **Grids.** Rabin [Rab76] used randomization, the floor function, and hashing to compute the closest pair of a set of points in the plane, in expected linear time. Golin *et al.* [GRSS95] presented a simplified version of this algorithm, and Smid provides a survey of algorithms on closest pair problems [Smi00]. A prune and search variant of this algorithm was suggested by Khuller and Matias [KM95]. Of course, the usage of grids and hashing to perform approximate point-location is quite common in practice. By itself, this is already sufficient to break lower bounds in the comparison model, for example for  $k$ -center clustering [Har04a]. The only direct extension of Rabin’s algorithm the authors are aware of is the work by Har-Peled and Mazumdar [HM05] showing a linear time 2-approximation to the smallest ball containing  $k$  points (out of  $n$  given points).

There is some skepticism of algorithms using the floor function, since Schönhage [Sch79] showed how to solve a PSPACE complete problem, in polynomial time, using the floor function in the real RAM model – the trick being packing many numbers into a single word (which can be arbitrarily long in the RAM model, and still each operation on it takes only constant time). Note, that as Rabin’s algorithm does not do any such packing of numbers (i.e., its computation model is considerably more restricted), this criticism does not seem to be relevant in the case of this algorithm and its relatives.

In this paper, we present a new technique that combines together all of the above techniques to yield linear time approximation algorithms.

**Nets** Given a point set  $P$ , an  $r$ -net  $\mathcal{N}$  of  $P$  is a subset of  $P$  that represents well the structure of  $P$  in resolution  $r$ . Formally, we require that for any point in  $P$  there is a net point in distance at most  $r$  from it, and no two net points are closer than  $r$  to each other, see Section 2.2 for a formal definition. Thus, nets provide a sketch of the point set for distances that are  $r$  or larger. Nets are a useful tool in presenting point sets hierarchically. In particular, computing nets of different resolutions and linking between different levels, leads to a tree like data-structure that can be used to facilitate many tasks, see for example the net-tree [KL04, HM06] for such a data-structure for doubling metrics. Nets can be defined in any metric space, but in Euclidean space a grid can sometimes provide an equivalent representation. In particular, net-trees can be interpreted as an extension of (compressed) quadtrees to more abstract settings.

Computing nets is closely related to  $k$ -center clustering. Specifically, Gonzalez [Gon85] shows how to compute an approximate net that has  $k$  points in  $O(nk)$  time, which is also a 2-approximation to the  $k$ -center clustering. This was later improved to  $O(n)$  time, for low dimensional Euclidean space [Har04a], if  $k$  is sufficiently small (using grids and hashing). Har-Peled and Mendel showed how to preprocess a point set in a metric space with constant doubling dimension, in  $O(n \log n)$  time, such that an (approximate)  $r$ -net can be extracted in (roughly) linear time in the size of the net.

## Our contribution.

In this paper, we consider problems of the following form: Given a set  $P$  of weighted points in  $\mathbb{R}^d$ , one wishes to solve an optimization problem whose solution is one of the pairwise distances of  $P$  (or “close”

to one of these values). Problems of this kind include computing the optimal  $k$ -center clustering, or the length of the  $k$ th edge in the MST of  $P$ , and many others. Specifically, we are interested in problems for which there is a fast approximate decider. That is, given a value  $r > 0$  we can, in linear time, decide if the desired value is (approximately) smaller than  $r$  or larger than  $r$ . The goal is then to use this decider to approximate the optimum solution in linear time. As a first step towards a linear time approximation algorithm for such problems, we point out that one can compute nets in linear time in  $\mathbb{R}^d$ , see [Section 2.2](#).

However, even if we could implicitly search over the critical values (which we cannot) then we still would require a logarithmic number of calls to the decider which would yield a running time of  $O(n \log n)$ , as the number of critical values is at least linear (and usually polynomial). So instead we use the return values of the decision procedure as we search to thin out the data so that future calls to the decision procedure become successively cheaper. However, we still cannot search over the critical values (since there are too many of them) and so we also introduce random sampling in order to overcome this.

**Outline of the new technique** The new algorithm works by randomly sampling a point and computing the distance to its nearest neighbor. Let this distance be  $r$ . Next, we use the decision procedure to decide if we are in one of the following three cases.

- (A) **Net.** If  $r$  is too small then we zoom out to a resolution of  $r$  by computing an  $r$ -net and continuing the computation on the net instead of on the original point set. That is, we net the point set into a smaller point set, such that one can solve the original problem (approximately) on this smaller sketch of the input.
- (B) **Prune.** If  $r$  is too large then we remove all points whose nearest neighbor is further than  $r$  away (of course, this implies we should only consider problems for which such pruning does not affect the solution). That is, we isolate the optimal solution by pruning away irrelevant data – this is similar in nature to what is being done by prune-and-search algorithms.
- (C) **Done.** The value of  $r$  is the desired approximation.

We then continue recursively on the remaining data. In either case, the number of points being handled (in expectation) goes down by a constant factor and thus the overall expected running time is linear. In particular, getting this constant factor decrease is the reason we chose to sample from the set of nearest neighbor distances rather than from the set of all inter-point distances.

**Significance of results** Our basic framework is presented in a general enough manner to cover, and in many cases greatly simplify, many problems for which linear time algorithms had already been discovered. At the same time the framework provides new linear time algorithms for a large collection of problems, for which previously no linear time algorithm was known. The framework should also lead to algorithms for many other problems which are not mentioned.

At a conceptual level the basic algorithm is simple enough (with its basic building blocks already having efficient implementations) to be highly practical from an implementation standpoint. Perhaps more importantly, with increasing shifts toward large data sets, algorithms with super linear running time can be impractical. Additionally, our framework seems amenable to distributed implementation in frameworks like MapReduce. Indeed, every iteration of our algorithm breaks the data into grid cells, a step that is similar to the map phase. In addition, the aggressive thinning of the data by the algorithm guarantees that after the first few iterations the algorithm is resigned to working on only a tiny fraction of the data.

**Significance of the netting step** Without the net stage, our framework is already sufficient to solve the closest pair problem, and in this specific case, the algorithm boils down to the one provided by Khuller and Matias [KM95]. This seems to be the only application of the new framework where the netting is not necessary. In particular, the wide applicability of the framework seems to be the result of the netting step.

**Framework and results** We provide a framework that classifies which optimization problems can be solved using the new algorithm. We get the following new algorithms (all of them have an expected linear running time, for any fixed  $\varepsilon$ ):

- (A)  **$k$ -center clustering** (Section 4.1). We provide an algorithm that 2-approximates the optimal  $k$ -center clustering of a point set in  $\mathbb{R}^d$ . Unlike the previous algorithm [Har04a] that was restricted to  $k = O(n^{1/6})$ , the new algorithm works for any value of  $k$ . This new algorithm is also simpler.
- (B)  **$k$ th smallest distance** (Section 4.2). In the distance selection problem, given a set of points in  $\mathbb{R}^d$ , one would like to compute the  $k$ th smallest distance defined by a pair of points of  $P$ . It is believed that such exact distance selection requires  $\Omega(n^{4/3})$  time in the worst case [Eri95], even in the plane (in higher dimensions the bound deteriorates). We present an  $O(n/\varepsilon^d)$  time algorithm that  $(1+\varepsilon)$ -approximates the  $k$ th smallest distance. Previously, Bespamyatnikh and Segal [BS02] presented  $O(n \log n + n/\varepsilon^d)$  time algorithm using a well-separated pairs decomposition (see also [DHW12]).

Given two sets of points  $P$  and  $W$  with a total of  $n$  points, using the same approach, we can  $(1+\varepsilon)$ -approximate the  $k$ th smallest distance in a bichromatic set of distances  $X = \{d(p, q) \mid p \in P, q \in W\}$ . Intuitively, the mechanism behind distance selection underlines many optimization problems, as it (essentially) performs a binary search over the distances induced by a set of points. As such, being able to do approximate distance selection in linear time should lead to faster approximation algorithms that perform a similar search over such distances.

- (C) **The  $k$ th smallest  $m$ -nearest neighbor distance** (Section 4.3). For a set  $P = \{p_1, \dots, p_n\} \subseteq \mathbb{R}^d$  of  $n$  points, and a point  $p \in P$ , its  *$m$ th nearest neighbor* in  $P$  is the  $m$ th closest point to  $p$  in  $P \setminus \{p\}$ . In particular, let  $d_m(p, P)$  denote this distance. Here, consider the set of these distances defined for each point of  $P$ ; that is,  $X = \{d_m(p_1, P), \dots, d_m(p_n, P)\}$ . We can approximate the  $k$ th smallest number in this set in linear time.
- (D) **Exact nearest neighbor distances** (Section 4.3.1). For the special case where  $m = 1$ , one can turn the above approximation into an exact computation of the  $k$ th nearest neighbor distance with only a minor post processing grid computation.

As an application, when  $k = n$ , one can compute in linear time, exactly, the furthest nearest neighbor distance; that is, the nearest neighbor distance of the point whose nearest neighbor is furthest away. This measure can be useful, for example, in meshing applications where such a point is a candidate for a region where the local feature size is large and further refinement is needed.

We are unaware of any previous work directly on this problem, although one can compute this quantity by solving the all-nearest-neighbor problem, which can be done in  $O(n \log n)$  time [Cla83]. This is to some extent the “antithesis” to Rabin’s algorithm for the closest pair problem, and it is somewhat surprising that it can also be solved in linear time.

- (E) **The  $k$ th longest MST edge** (Section 4.4). Given a set  $P$  of  $n$  points in  $\mathbb{R}^d$ , we can  $(1+\varepsilon)$ -approximate, in  $O(n/\varepsilon^d)$  time, the  $k$ th longest edge in the MST of  $P$ .
- (F) **Smallest ball with a monotone property** (Section 4.5.2). Consider a property defined over a set of points,  $P$ , that is monotone; that is, if  $W \subseteq Q \subseteq P$  has this property then  $Q$  must also have

this property. Consider such a property that can be easily checked, for example, (i) whether the set contains  $k$  points, or (ii) if the points are colored, whether all colors are present in the given point set. Given a point set  $\mathbf{P}$ , one can  $(1 + \varepsilon)$ -approximate, in  $O(n/\varepsilon^d)$  time, the smallest radius ball  $\mathbf{b}$ , such that  $\mathbf{b} \cap \mathbf{P}$  has the desired property. For example, we get a linear time algorithm to approximate the smallest ball enclosing  $k$  points of  $\mathbf{P}$ . The previous algorithm for this problem [HM05] was significantly more complicated. Furthermore, we can approximate the smallest ball such that all colors appear in it (if the points are colored), or the smallest ball such that at least  $t$  different colors appear in it, etc.

More generally, the kind of monotone properties supported are *sketchable*; that is, properties for which there is a small summary of a point set that enables one to decide if the property holds, and furthermore, given summaries of two disjoint point sets, the summary of the union point set can be computed in constant time. We believe that formalizing this notion of sketchability is a useful abstraction. See Section 4.5.1 for details.

- (G) **Smallest connected component with a monotone property** (Section 4.5.3). Consider the connected components of the graph where two points are connected if they are distance at most  $r$  from each other. Using our techniques, one can approximate, in linear time, the smallest  $r$  such that there is a connected component of this graph for which a required sketchable property holds for the points in this connected component.

As an application, consider ad hoc wireless networks. Here, we have a set  $\mathbf{P}$  of  $n$  nodes and their locations (say in the plane), and each node can broadcast in a certain radius  $r$  (the larger the  $r$  the higher the energy required, so naturally we would like to minimize it). Assume there are two special nodes. It is natural to ask for the minimum  $r$ , such that there is a connected component of the above graph that contains both nodes. That is, these two special nodes can send message to each other, by message hopping (with distance at most  $r$  at each hop). We can approximate this connectivity radius in linear time.

- (H) **Clustering for a monotone property** (Section 4.6). Imagine that we want to break the given point set into clusters, such that the maximum diameter of a cluster is minimized (as in  $k$ -center clustering), and furthermore, the points assigned to each cluster comply with some sketchable monotone property. We present a  $(4 + \varepsilon)$ -approximation algorithm for these types of problems, that runs in  $O(n/\varepsilon^d)$  time. This includes lower bounded clustering (i.e., every cluster must contain at least  $\alpha$  points), for which the authors recently presented an  $O(n \log n)$  approximation algorithm [ERH12]. One can get a 2-approximation using network flow, but the running is significantly worse [APF<sup>+</sup>10]. See Section 4.6.1 for examples of clustering problems that can be approximated using this algorithm.

- (I) **Connectivity clustering for a monotone property** (Section 4.6.3). Consider the problem of computing the minimum  $r$ , such that each connected component (of the graph where points in distance at most  $r$  from each other are adjacent) has some sketchable monotone property. We approximate the minimum  $r$  for which this holds in linear time.

An application of this for ad hoc networks is the following – we have a set  $\mathbf{P}$  of  $n$  wireless clients, and some of them are base stations; that is, they are connected to the outside world. We would like to find the minimum  $r$ , such that each connected component of this graph contains a base station.

- (J) **Closest pair and smallest non-zero distance** (Section 4.7). Given a set of points in  $\mathbb{R}^d$ , consider the problem of finding the smallest non-zero distance defined by these points. This problem is an extension of the closest pair distance, as there might be many identical points in the given point set. We provide a linear time algorithm for computing this distance *exactly*, which follows easily from our framework.

**High probability** Finally, we show in [Section 5](#) how to modify our framework such that the linear running time holds with high probability. Since there is very little elbow room in the running time when committed to linear running time, this extension is quite challenging, and requires several new ideas and insights. See [Section 5](#) for details.

**What can be done without the net & prune framework** For the problems studied in this paper, an approximation algorithm with running time  $O(n \log n)$  follows by using a  $(1 + \varepsilon)$ -WSPD (well-separated pairs decomposition [[CK95](#), [Har11](#)]), for some fixed  $\varepsilon$ , to compute  $O(n)$  relevant resolutions, and then using binary search with the decider to find the right resolution. Even if the WSPD is given, it is not clear though how to get a linear time algorithm for the problems studied without using the net & prune framework. In particular, if the input has bounded spread, one can compute an  $O(1)$ -WSPD in linear time, and then use the WSPD inside the net & prune framework to get a deterministic linear time algorithm, see [Section 3.3.1](#).

**Paper organization** We describe how to compute nets in [Section 2](#), and how to remove faraway points efficiently in [Section 2.3](#). We define the abstract framework, and describe and analyze the new approximation algorithm, in [Section 3](#). We describe the applications in [Section 4](#). We show how to modify the framework to achieve linear running time with high probability in [Section 5](#). We conclude in [Section 6](#).

## 2. Preliminaries

### 2.1. Basic definitions

The *spread* of an interval  $\mathcal{I} = [x, y] \subseteq \mathbb{R}^+$  is  $\text{sprd}(\mathcal{I}) = y/x$ .


**Definition 2.1.1.** For a point set  $P$  in a metric space with a metric  $d$ , and a parameter  $r > 0$ , an  *$r$ -net* of  $P$  is a subset  $\mathcal{C} \subseteq P$ , such that (i) for every  $\mathbf{p}, \mathbf{q} \in \mathcal{C}$ ,  $\mathbf{p} \neq \mathbf{q}$ , we have that  $d(\mathbf{p}, \mathbf{q}) \geq r$ , and (ii) for all  $\mathbf{p} \in P$ , we have that  $\min_{\mathbf{q} \in \mathcal{C}} d(\mathbf{p}, \mathbf{q}) < r$ .

Intuitively, an  $r$ -net represents  $P$  in resolution  $r$ .

**Definition 2.1.2.** For a real positive number  $\Delta$  and a point  $\mathbf{p} = (p_1, \dots, p_d) \in \mathbb{R}^d$ , define  $G_\Delta(\mathbf{p})$  to be the grid point  $(\lfloor p_1/\Delta \rfloor \Delta, \dots, \lfloor p_d/\Delta \rfloor \Delta)$ .

The quantity  $\Delta$  is the *width* or *sidelength* of the *grid*  $G_\Delta$ . Observe that  $G_\Delta$  partitions  $\mathbb{R}^d$  into cubes, which are grid *cells*. The grid cell of  $\mathbf{p}$  is uniquely identified by the integer point  $\text{id}(\mathbf{p}) = (\lfloor p_1/\Delta \rfloor, \dots, \lfloor p_d/\Delta \rfloor)$ .

For a number  $r \geq 0$ , let  $N_{\leq r}(\mathbf{p})$  denote the set of grid cells in distance  $\leq r$  from  $\mathbf{p}$ , which is the *neighborhood* of  $\mathbf{p}$ . Note, that the neighborhood also includes the grid cell containing  $\mathbf{p}$  itself, and if  $\Delta = \Theta(r)$  then  $|N_{\leq r}(\mathbf{p})| = \Theta((2 + \lceil 2r/\Delta \rceil)^d) = \Theta(1)$ . See figure on the right.



### 2.2. Computing nets quickly for a point set in $\mathbb{R}^d$

There is a simple algorithm for computing  $r$ -nets. Namely, let all the points in  $P$  be initially unmarked. While there remains an unmarked point,  $\mathbf{p}$ , add  $\mathbf{p}$  to  $\mathcal{C}$ , and mark it and all other points in distance  $< r$

from  $\mathbf{p}$  (i.e. we are scooping away balls of radius  $r$ ). By using grids and hashing one can modify this algorithm to run in linear time. The following is implicit in previous work [Har04a], and we include it here for the sake of completeness\* – it was also described by the authors in [ERH12].

**Lemma 2.2.1.** *Given a point set  $P \subseteq \mathbb{R}^d$  of size  $n$  and a parameter  $r > 0$ , one can compute an  $r$ -net for  $P$  in  $O(n)$  time.*

*Proof:* Let  $G$  denote the grid in  $\mathbb{R}^d$  with side length  $\Delta = r/(2\sqrt{d})$ . First compute for every point  $\mathbf{p} \in P$  the grid cell in  $G$  that contains  $\mathbf{p}$ ; that is,  $\text{id}(\mathbf{p})$ . Let  $\mathcal{G}$  denote the set of grid cells of  $G$  that contain points of  $P$ . Similarly, for every cell  $\square \in \mathcal{G}$  we compute the set of points of  $P$  which it contains. This task can be performed in linear time using hashing and bucketing assuming the floor function can be computed in constant time. Specifically, store the  $\text{id}(\cdot)$  values in a hash table, and in constant time hash each point into its appropriate bin.

Scan the points of  $P$  one at a time, and let  $\mathbf{p}$  be the current point. If  $\mathbf{p}$  is marked then move on to the next point. Otherwise, add  $\mathbf{p}$  to the set of net points,  $\mathcal{C}$ , and mark it and each point  $\mathbf{q} \in P$  such that  $d(\mathbf{p}, \mathbf{q}) < r$ . Since the cells of  $N_{\leq r}(\mathbf{p})$  contain all such points, we only need to check the lists of points stored in these grid cells. At the end of this procedure every point is marked. Since a point can only be marked if it is in distance  $< r$  from some net point, and a net point is only created if it is unmarked when visited, this implies that  $\mathcal{C}$  is an  $r$ -net.

As for the running time, observe that a grid cell,  $c$ , has its list scanned only if  $c$  is in the neighborhood of some created net point. As  $\Delta = \Theta(r)$ , there are only  $O(1)$  cells which could contain a net point  $\mathbf{p}$  such that  $c \in N_{\leq r}(\mathbf{p})$ . Furthermore, at most one net point lies in a single cell since the diameter of a grid cell is strictly smaller than  $r$ . Therefore each grid cell had its list scanned  $O(1)$  times. Since the only real work done is in scanning the cell lists and since the cell lists are disjoint, this implies an  $O(n)$  running time overall. ■

Observe that the closest net point, for a point  $\mathbf{p} \in P$ , must be in one of its neighborhood's grid cells. Since every grid cell can contain only a single net point, it follows that in constant time per point of  $P$ , one can compute each point's nearest net point. We thus have the following.

**Corollary 2.2.2.** *For a set  $P \subseteq \mathbb{R}^d$  of  $n$  points, and a parameter  $r > 0$ , one can compute, in linear time, an  $r$ -net of  $P$ , and furthermore, for each net point the set of points of  $P$  for which it is the nearest net point.*

In the following, a **weighted point** is a point that is assigned a positive integer weight. For any subset  $S$  of a weighted point set  $P$ , let  $|S|$  denote the number of points in  $S$  and let  $\omega(S) = \sum_{\mathbf{p} \in S} \omega(\mathbf{p})$  denote the total weight of  $S$ .

In particular, **Corollary 2.2.2** implies that for a weighted point set one can compute the following quantity in linear time.

**Algorithm 2.2.3 (net).** *Given a weighted point set  $P \subseteq \mathbb{R}^d$ , let  $\text{net}(r, P)$  denote an  $r$ -net of  $P$ , where the weight of each net point  $\mathbf{p}$  is the total sum of the weights of the points assigned to it. We slightly abuse notation, and also use  $\text{net}(r, P)$  to designate the algorithm (of **Corollary 2.2.2**) for computing this net, which has linear running time.*

---

\*Specifically, the algorithm of Har-Peled [Har04a] is considerably more complicated than **Lemma 2.2.1**, and does not work in this settings, as the number of clusters it can handle is limited to  $O(n^{1/6})$ . **Lemma 2.2.1** has no such restriction.



## 2.3. Identifying close and far points

For a given point  $\mathbf{p} \in \mathbf{P}$ , let

$$\nu(\mathbf{p}, \mathbf{P}) = \arg \min_{\mathbf{q} \in \mathbf{P} \setminus \{\mathbf{p}\}} d(\mathbf{p}, \mathbf{q}) \quad \text{and} \quad d(\mathbf{p}, \mathbf{P}) = \min_{\mathbf{q} \in \mathbf{P} \setminus \{\mathbf{p}\}} d(\mathbf{p}, \mathbf{q}), \quad (2.1)$$

denote the nearest neighbor of  $\mathbf{p}$  in  $\mathbf{P} \setminus \{\mathbf{p}\}$ , and the distance to it, respectively. The quantity  $d(\mathbf{p}, \mathbf{P})$  can be computed (naively) in linear time by scanning the points. For a set of points  $\mathbf{P}$ , and a parameter  $r$ , let  $\mathbf{P}^{\geq r}$  denote the set of ***r-far*** points; that is, it is the set of all points  $\mathbf{p} \in \mathbf{P}$ , such that the nearest-neighbor of  $\mathbf{p}$  in  $\mathbf{P} \setminus \{\mathbf{p}\}$  is at least distance  $r$  away (i.e.,  $d(\mathbf{p}, \mathbf{P}) \geq r$ ). Similarly,  $\mathbf{P}^{< r}$  is the set of ***r-close*** points; that is, all points  $\mathbf{p} \in \mathbf{P}$ , such that  $d(\mathbf{p}, \mathbf{P}) < r$ .

**Lemma 2.3.1.** *Given a weighted set  $\mathbf{P} \subseteq \mathbb{R}^d$  of  $n$  points, and a distance  $\ell > 0$ , in  $O(n)$  time, one can compute the sets  $\mathbf{P}^{< \ell}$  and  $\mathbf{P}^{\geq \ell}$ . Let  $\mathbf{delFar}(\ell, \mathbf{P})$  denote the algorithm which computes these sets and returns  $\mathbf{P}^{< \ell}$ .*

*Proof:* Build a grid where every cell has diameter  $\ell/c$ , for some constant  $c > 1$ . Clearly any point  $\mathbf{p} \in \mathbf{P}$  such that  $d(\mathbf{p}, \mathbf{P}) \geq \ell$  (i.e., a far point) must be in a grid cell by itself. Therefore to determine all such “far” points only grid cells with singleton points in them need to be considered. For such a point  $\mathbf{q}$ , to determine if  $d(\mathbf{q}, \mathbf{P}) \geq \ell$ , one checks the points stored in all grid cells in distance  $\leq \ell$  from it. If  $\mathbf{q}$  has no such close neighbor, we mark  $\mathbf{q}$  for inclusion in  $\mathbf{P}^{\geq r}$ . By the same arguments as in [Lemma 2.2.1](#), the number of such cells is  $O(1)$ . Again by the arguments of [Lemma 2.2.1](#) every non-empty grid cell gets scanned  $O(1)$  times overall, and so the running time is  $O(n)$ . Finally, we copy all the marked (resp. unmarked) points to  $\mathbf{P}^{\geq \ell}$  (resp.  $\mathbf{P}^{< \ell}$ ). The algorithm  $\mathbf{delFar}(\ell, \mathbf{P})$  then simply returns  $\mathbf{P}^{< \ell}$ . ■

## 3. Approximating nice distance problems

### 3.1. Problem definition and an example

**Definition 3.1.1.** Let  $\mathbf{P}$  and  $\mathbf{Q}$  be two sets of weighted points in  $\mathbb{R}^d$  (of the same weight). The set  $\mathbf{Q}$  is a  ***$\Delta$ -drift*** of  $\mathbf{P}$ , if  $\mathbf{Q}$  can be constructed by moving each point of  $\mathbf{P}$  by distance at most  $\Delta$  (and not altering its weight). Formally, there is an onto mapping  $f : \mathbf{P} \rightarrow \mathbf{Q}$ , such that (i) For  $\mathbf{p} \in \mathbf{P}$ , we have that  $d(\mathbf{p}, f(\mathbf{p})) \leq \Delta$ , and (ii) for any  $\mathbf{q} \in \mathbf{Q}$ , we have that  $\omega(\mathbf{q})$  is the sum of the weights of all points  $\mathbf{p} \in \mathbf{P}$  such that  $f(\mathbf{p}) = \mathbf{q}$ .

Note that for a (potentially weighted) point set  $\mathbf{W}$ ,  $\mathbf{net}(\Delta, \mathbf{W})$  is a  $\Delta$ -drift of  $\mathbf{W}$ .

**Definition 3.1.2.** Given a set  $X$  and a function  $f : X \rightarrow \mathbb{R}$ , a procedure **decider** is a  ***$\varphi$ -decider*** for  $f$ , if for any  $x \in X$  and  $r > 0$ ,  $\mathbf{decider}(r, x)$  returns one of the following: (i)  $f(x) \in [\alpha, \varphi\alpha]$ , where  $\alpha$  is some real number, (ii)  $f(x) < r$ , or (iii)  $f(x) > r$ .

**Definition 3.1.3 ( $\varphi$ -NDP).** An instance of a  ***$\varphi$ -NiceDistanceProblem*** consists of a pair  $(\mathbf{W}, \Gamma)$ , where  $\mathbf{W} \subseteq \mathbb{R}^d$  is a set of  $n$  distinct weighted points<sup>†</sup>, and  $\Gamma$  is the ***context*** of the given instance (of size  $O(n)$ ) and consists of the relevant parameters for the problem<sup>‡</sup>. For a fixed  ***$\varphi$ -NDP***, the task is to evaluate a function  $f(\mathbf{W}, \Gamma) \rightarrow \mathbb{R}^+$ , defined over such input pairs, that has the following properties:

<sup>†</sup>The case when  $\mathbf{W}$  is a multiset can also be handled. See [Remark 3.2.2](#).

<sup>‡</sup>In almost all cases the context is a set of integers, usually of constant size.

- **Decider:** There exists an  $O(n)$  time  $\varphi$ -decider for  $f$ , for some *constant*  $\varphi \geq 1$ .
- **Lipschitz:** Let  $\mathbf{Q}$  be any  $\Delta$ -drift of  $\mathbf{W}$ . Then  $|f(\mathbf{W}, \Gamma) - f(\mathbf{Q}, \Gamma)| \leq 2\Delta$ .
- **Prune:** If  $f(\mathbf{W}, \Gamma) < r$  then  $f(\mathbf{W}, \Gamma) = f(\mathbf{W}^{<r}, \Gamma')$ , where  $\mathbf{W}^{<r}$  is the set of  $r$ -close points, and  $\Gamma'$  is an updated context which can be computed, in  $O(n)$  time, by a procedure denoted **contextUpdate**.

**Remark 3.1.4** ( $(1 + \varepsilon)$ -**NDP**). If we are interested in a  $(1 + \varepsilon)$ -approximation, then  $\varphi = 1 + \varepsilon$ . In this case, we require that the running time of the provided  $(1 + \varepsilon)$ -decider (i.e., property **Decider** above) is  $O(n/\varepsilon^c)$ , for some constant  $c \geq 1$ . This covers all the applications presented in this paper.

Our analysis still holds even if the decider has a different running time, but the overall running time might increase by a factor of  $\log(1/\varepsilon)$ , see [Lemma 3.3.1<sub>p15</sub>](#) for details.

### 3.1.1. An example – $k$ center clustering

As a concrete example of an **NDP**, consider the problem of  $k$ -center clustering.

**Problem 3.1.5 (kCenter).** Let  $\mathbf{W}$  be a set of points in  $\mathbb{R}^d$ , and let  $k > 0$  be an integer parameter. Find a set of **centers**  $C \subseteq \mathbf{W}$  such that the maximum distance of a point of  $\mathbf{W}$  to its nearest center in  $C$  is minimized, and  $|C| = k$ .

Specifically, the function of interest, denoted by  $f_{cen}(\mathbf{W}, k)$ , is the radius of the optimal  $k$ -center clustering of  $\mathbf{W}$ . We now show that **kCenter** satisfies the properties of an **NDP**.

**Remark 3.1.6.** Usually one defines the input for **kCenter** to be a set of unweighted points. However, since the definition of **NDP** requires a weighted point set, we naturally convert the initial input into a set of unit weight points,  $\mathbf{W}$ . Also, for simplicity we assume  $f_{cen}(\mathbf{W}, k) \neq 0$ , i.e.  $|\mathbf{W}| > k$  (which can easily be checked). Both assumptions would be used implicitly throughout the paper.

**Lemma 3.1.7.** The following relations between  $|\mathbf{net}(r, \mathbf{W})|$  and  $f_{cen}(\mathbf{W}, k)$  hold: (A) If we have  $|\mathbf{net}(r, \mathbf{W})| \leq k$  then  $f_{cen}(\mathbf{W}, k) < r$ . (B) If  $|\mathbf{net}(r, \mathbf{W})| > k$  then  $r \leq 2f_{cen}(\mathbf{W}, k)$ .

*Proof:* (A) Observe that if  $|\mathbf{net}(r, \mathbf{W})| \leq k$ , then, by definition, the set of net points of  $\mathbf{net}(r, \mathbf{W})$  are a set of  $\leq k$  centers such that all the points of  $\mathbf{W}$  are in distance strictly less than  $r$  from these centers.

(B) If  $|\mathbf{net}(r, \mathbf{W})| > k$  then  $\mathbf{W}$  must contain a set of  $k + 1$  points whose pairwise distances are all at least  $r$ . In particular any solution to **kCenter** with radius  $< r/2$  would not be able to cover all these  $k + 1$  points with only  $k$  centers. ■

**Lemma 3.1.8.** For any instance  $(\mathbf{W}, k)$  of **kCenter**,  $f_{cen}(\mathbf{W}, k)$  satisfies the properties of [Definition 3.1.3](#); that is, **kCenter** is a  $(4 + \varepsilon)$ -**NDP**, for any  $\varepsilon > 0$ .

*Proof:* We need to verify the required properties, see [Definition 3.1.3<sub>p9</sub>](#).

**Decider:** We need to describe a decision procedure for **kCenter** clustering. To this end, given a distance  $r$ , the decider first calls  $\mathbf{net}(r, \mathbf{W})$ , see [Algorithm 2.2.3<sub>p8</sub>](#). If we have  $|\mathbf{net}(r, \mathbf{W})| \leq k$ , then by [Lemma 3.1.7](#) the answer “ $f_{cen}(\mathbf{W}, k) < r$ ” can be returned. Otherwise, call  $\mathbf{net}((2 + \varepsilon/2)r, \mathbf{W})$ . If  $|\mathbf{net}((2 + \varepsilon/2)r, \mathbf{W})| \leq k$ , then, by [Lemma 3.1.7](#), we have  $r/2 \leq f_{cen}(\mathbf{W}, k) < (2 + \varepsilon/2)r$  and “ $f \in [r/2, (2 + \varepsilon/2)r]$ ” is returned. Otherwise  $|\mathbf{net}((2 + \varepsilon/2)r, \mathbf{W})| > k$ , and [Lemma 3.1.7](#) implies that the answer “ $r < f_{cen}(\mathbf{W}, k)$ ” can be returned by the decider.

```

ndpAlg( $W, \Gamma$ ):
  Let  $W_0 = W, \Gamma_0 = \Gamma$  and  $i = 1$ .
  while TRUE do
1:   Randomly pick a point  $p$  from  $W_{i-1}$ .
2:    $l_i \leftarrow d(p, W_{i-1})$ .
   // Next, estimate the value of  $f_{i-1} = f(W_{i-1}, \Gamma_{i-1})$ 
    $res_> = \mathbf{decider}(l_i, W_{i-1}, \Gamma_{i-1})$ 
    $res_< = \mathbf{decider}(c_{\text{net}}l_i, W_{i-1}, \Gamma_{i-1})$  //  $c_{\text{net}} = 37$ : Corollary 3.2.8.
3:   if  $res_> = "f_{i-1} \in [x, y]"$  then return " $f(W, \Gamma) \in [x/2, 2y]"$ .
4:   if  $res_< = "f_{i-1} \in [x', y']"$  then return " $f(W, \Gamma) \in [x'/2, 2y']"$ .
   if  $res_> = "l_i < f_{i-1}"$  and  $res_< = "f_{i-1} < c_{\text{net}}l_i"$  then
5:     return " $f(W, \Gamma) \in [l_i/2, 2c_{\text{net}}l_i]"$ .
   if  $res_> = "f_{i-1} < l_i"$  then
6:      $W_i = \mathbf{delFar}(l_i, W_{i-1})$  // Prune
      $\Gamma_i = \mathbf{contextUpdate}(W_{i-1}, \Gamma_{i-1}, W_i)$ 
   if  $res_< = "c_{\text{net}}l_i < f_{i-1}"$  then
7:      $W_i = \mathbf{net}(3l_i, W_{i-1})$  // Net
      $\Gamma_i = \Gamma_{i-1}$ 
    $i = i + 1$ 

```

Figure 3.1: The approximation algorithm. The implicit target value being approximated is  $f = f(W, \Gamma)$ , where  $f$  is a  $\varphi$ -NDP. That is, there is a  $\varphi$ -decider for  $f$ , denoted by **decider**, and the only access to the function  $f$  is via this procedure.

**Lipschitz:** Observe that if  $Q$  is a  $\Delta$ -drift of  $W$ , then a point and its respective center in a  $k$ -center clustering of  $W$  each move by distance at most  $\Delta$  in the transition from  $W$  to  $Q$ . As such, the distance between a point and its center changes by at most  $2\Delta$  by this process. This argument also works in the other direction, implying that the  $k$ -center clustering radius of  $W$  and  $Q$  are the same, up to an additive error of  $2\Delta$ .

**Prune:** It suffices to show that if  $f_{cen}(W, k) < r$  then  $f_{cen}(W, k) = f_{cen}(W^{<r}, k - |W^{\geq r}|)$ . Now, if  $f_{cen}(W, k) < r$  then any point of  $W$  whose neighbors are all  $\geq r$  away must be a center by itself in the optimal  $k$ -center solution, as otherwise it would be assigned to a center  $\geq r > f_{cen}(W, k)$  away. Similarly, any point assigned to it would be assigned to a point  $\geq r > f_{cen}(W, k)$  away. Therefore, for any point  $p \in W$  whose nearest neighbor is at least distance  $r$  away,  $f_{cen}(W, k) = f_{cen}(W \setminus \{p\}, k - 1)$ . Repeating this observation implies the desired result. The context update (and computing  $W^{<r}$  and  $W^{\geq r}$ ) can be done in linear time using **delFar**, see Lemma 2.3.1<sub>p9</sub>. ■

### 3.2. A linear time approximation algorithm

We now describe the general algorithm which given a  $\varphi$ -NDP,  $(W, \Gamma)$ , and an associated target function,  $f$ , computes in linear time a  $O(\varphi)$  spread interval containing  $f(W, \Gamma)$ . In the following, let **decider** denote the given  $\varphi$ -decider, where  $\varphi \geq 1$  is some parameter. Also, let **contextUpdate** denote the context updater associated with the given NDP problem, see Definition 3.1.3. Both **decider** and **contextUpdate** run in linear time. The algorithm for bounding the optimum value of an NDP is shown in Figure 3.1.

**Remark 3.2.1.** For the sake of simplicity of exposition we assume that  $f(W, \Gamma) \neq 0$ . Since all our applications have  $f(W, \Gamma) \neq 0$  we choose to make this simplifying assumption. In particular, we later show in [Corollary 3.2.8](#) that if initially  $f(W, \Gamma) \neq 0$ , then **ndpAlg** preserves this property.

**Remark 3.2.2.** Note that the algorithm of [Figure 3.1](#) can be modified to handle inputs where  $W$  is a multiset (namely, two points can occupy the same location), and we still assume (as done above) that  $f(W, \Gamma) \neq 0$ . Specifically, it must be ensured that the distance computed in [Line 2](#) is not zero, as this is required for the call to **decider**. This can be remedied (in linear time) by first grouping all the duplicates of the point sampled in [Line 1](#) into a single weighted point (with the sum of the weights) before calling [Line 2](#). This can be done by computing a net for a radius that is smaller than (say) half the smallest non-zero distance. How to compute this distance is described in [Section 4.7](#).

### 3.2.1. Analysis

A *net iteration* of the algorithm is an iteration where **net** gets called. A *prune iteration* is one where **delFar** gets called. Note that the only other type of iteration is the one where the algorithm returns. In the following,  $P^{\leq \ell}$  is defined analogously to  $P^{< \ell}$  (see [Section 2.3](#)).

**Lemma 3.2.3.** *Let  $P$  be a point set. A  $(2 + \varepsilon)\ell$ -net of  $P$ , for any  $\varepsilon > 0$ , can contain at most half the points of  $P^{\leq \ell}$ .*

*Proof:* Consider any point  $p$  in  $P^{\leq \ell}$  which became one of the net points. Since  $p \in P^{\leq \ell}$ , a disk of radius  $\ell$  centered at  $p$  must contain another point  $q$  from  $P^{\leq \ell}$  (indeed,  $p \in P^{\leq \ell}$  only if its distance from its nearest neighbor in  $P$  is at most  $\ell$ ). Moreover,  $q$  cannot become a net point since it is too close to  $p$ . Now if we place a ball of radius  $\ell$  centered at each point of  $P^{\leq \ell}$  which became a net point, then these balls will be disjoint because the pairwise distance between net points is  $\geq (2 + \varepsilon)\ell$ . Therefore each point of  $P^{\leq \ell}$  which became a net point can charge to at least one point of  $P^{\leq \ell}$  which did not make it into the net, such that no point gets charged twice. ■

**Lemma 3.2.4.** *Given an instance  $(W, \Gamma)$  of an **NDP**, the algorithm **ndpAlg** $(W, \Gamma)$  runs in expected  $O(n)$  time.*

*Proof:* In each iteration of the while loop the only non-trivial work done is in computing  $\ell_i$ , the two calls to **decider**, and the one call to either **net** or **delFar**. It has already been shown that all of these can be computed in  $O(|W_{i-1}|)$  time. Hence the total running time for the algorithm is  $O\left(\sum_{i=0}^{i=k-1} |W_i|\right)$ , where  $k$  denotes the last (incomplete) iteration of the while loop.

So consider the beginning of iteration  $i < k$  of the while loop. Let the points in  $W_{i-1}$  be labeled  $p_1, p_2, \dots, p_m$  in increasing order of their nearest neighbor distance in  $W_{i-1}$ . Let  $j$  be the index of the point chosen in [Line 1](#) and let  $(W_{i-1})^{\geq j}$  and  $(W_{i-1})^{\leq j}$  be the subset of the points with index  $\geq j$  and index  $\leq j$ , respectively. Now since a point is randomly picked in [Line 1](#), with probability  $\geq 1/2$ ,  $j \in [m/4, 3m/4]$ . Lets call this event a *successful iteration*. We have  $\min\left(|(W_{i-1})^{\geq j}|, |(W_{i-1})^{\leq j}|\right) \geq |W_{i-1}|/4$  for a successful iteration.

Since  $i < k$  is not the last iteration of the while loop, either **delFar** or **net** must get called. If **delFar** $(\ell_i, W_{i-1})$  gets called (i.e. [Line 6](#)) then by [Lemma 2.3.1](#), all of  $(W_{i-1})^{\geq j}$  gets removed. So suppose **net** gets called (i.e. [Line 7](#)). In this case [Lemma 3.2.3](#) implies that the call to **net** removes at least  $\left|(W_{i-1})^{\leq j}\right|/2$  points.

Therefore, for any iteration  $i < k$ , at least  $\nu_i = \min(|(\mathbf{W}_{i-1})^{\geq j}|, |(\mathbf{W}_{i-1})^{\leq j}|/2)$  points get removed. If an iteration is successful then  $\nu_i \geq |\mathbf{W}_{i-1}|/8$ . In particular,  $\mathbf{E}[\nu_i | |\mathbf{W}_{i-1}|] \geq |\mathbf{W}_{i-1}|/16$ . Now,  $|\mathbf{W}_i| \leq |\mathbf{W}_{i-1}| - \nu_i$  and as such  $\mathbf{E}[|\mathbf{W}_i| | |\mathbf{W}_{i-1}|] \leq (15/16)|\mathbf{W}_{i-1}|$ .

Therefore, for  $0 < i < k$ ,

$$\mathbf{E}[|\mathbf{W}_i|] = \mathbf{E}[\mathbf{E}[|\mathbf{W}_i| | \mathbf{W}_{i-1}]] \leq \mathbf{E}\left[\frac{15}{16}|\mathbf{W}_{i-1}|\right] = \frac{15}{16}\mathbf{E}[|\mathbf{W}_{i-1}|].$$

Hence, by induction on  $i$ ,  $\mathbf{E}[|\mathbf{W}_i|] \leq (15/16)^i |\mathbf{W}_0|$  and so, in expectation, the running time is bounded by

$$O\left(\mathbf{E}\left[\sum_{i=0}^{i=k-1} |\mathbf{W}_i|\right]\right) = O\left(\sum_{i=0}^{i=k-1} \mathbf{E}[|\mathbf{W}_i|]\right) = O\left(\sum_{i=0}^{i=k-1} (15/16)^i |\mathbf{W}_0|\right) = O(|\mathbf{W}_0|).$$

■

**Correctness** The formal proof of correctness is somewhat tedious, but here is the basic idea: At every iteration, either far points are being thrown away (and this does not effect the optimal value), or we net the points. However, the net radius being used is always significantly smaller than the optimal value, and throughout the algorithm execution the radii of the nets being used grow exponentially. As such, the accumulated error in the end is proportional to the radius of the last net computation before termination, which itself is also much smaller than the optimal value.

Before proving that **ndpAlg** returns a bounded spread interval containing  $f(\mathbf{W}_0, \Gamma_0)$ , several helper lemmas will be needed. For notational ease, in the rest of this section, we use  $f(\mathbf{W}_i)$  as shorthand for  $f(\mathbf{W}_i, \Gamma_i)$ .

**Lemma 3.2.5.** *Suppose that **net** is called in iteration  $i$  of the while loop (i.e., [Line 7](#) in [Figure 3.1<sub>p11</sub>](#)). Then for any iteration  $j > i$  we have,  $\ell_j \geq 3\ell_i$ .*

*Proof:* Consider the beginning of iteration  $j$  of the while loop. The current set of points,  $\mathbf{W}_{j-1}$ , are a subset of the net points of a  $3\ell_i$ -net (it is a subset since [Line 6](#) and [Line 7](#) might have been executed in between rounds  $i$  and  $j$ ). Therefore, being a net, the distance between any two points of  $\mathbf{W}_{j-1}$  is  $\geq 3\ell_i$ , see [Definition 2.1.1<sub>p7</sub>](#). In particular, this means that for any point  $\mathbf{p}$  of  $\mathbf{W}_{j-1}$ , we have  $d(\mathbf{p}, \mathbf{W}_{j-1}) \geq 3\ell_i$ . ■

**Lemma 3.2.6.** *For  $i = 1, \dots, k$ , we have  $|f(\mathbf{W}_i) - f(\mathbf{W}_0)| \leq 9\ell_i$ .*

*Proof:* Let  $I$  be the set of indices of the net iterations up to (and including) the  $i$ th iteration. Similarly, let  $\bar{I}$  be the set of iterations where **delFar** get called.

If **net** was called in the  $j$ th iteration, then  $\mathbf{W}_j$  is a  $3\ell_j$ -drift of  $\mathbf{W}_{j-1}$  and so by the Lipschitz property,  $|f(\mathbf{W}_j) - f(\mathbf{W}_{j-1})| \leq 6\ell_j$ . On the other hand, if **delFar** gets called in the  $j$ th iteration, then  $f(\mathbf{W}_j) = f(\mathbf{W}_{j-1})$  by the Prune property. Let  $m = \max I$ , we have that

$$\begin{aligned} |f(\mathbf{W}_i) - f(\mathbf{W}_0)| &\leq \sum_{j=1}^i \left| f(\mathbf{W}_j) - f(\mathbf{W}_{j-1}) \right| \\ &= \sum_{j \in I} \left| f(\mathbf{W}_j) - f(\mathbf{W}_{j-1}) \right| + \sum_{j \in \bar{I}} \left| f(\mathbf{W}_j) - f(\mathbf{W}_{j-1}) \right| \\ &\leq \sum_{j \in I} 6\ell_j + \sum_{j \in \bar{I}} 0 \leq 6\ell_m \sum_{j=0}^{\infty} \frac{1}{3^j} \leq 9\ell_m \leq 9\ell_i, \end{aligned}$$

by [Lemma 3.2.5](#). ■

The following lemma testifies that the radii of the nets computed by the algorithm are always significantly smaller than the value we are trying to approximate.

**Lemma 3.2.7.** *For any iteration  $i$  of the while loop such that **net** gets called, we have  $\ell_i \leq f(W_0)/\eta$ , where  $0 < \eta = c_{\text{net}} - 9$ .*

*Proof:* The proof will be by induction. Let  $m_1, \dots, m_t$  be the indices of the iterations of the while loop in which **net** gets called. For the base case, observe that in order for **net** to get called, we must have  $\eta \ell_{m_1} < c_{\text{net}} \ell_{m_1} < f(W_{m_1-1})$ . However, since this is the first iteration in which **net** is called it must be that  $f(W_0) = f(W_{m_1-1})$  (since for all previous iterations **delFar** must have been called).

So now suppose that  $\ell_{m_j} \leq f(W_0)/\eta$  for all  $m_j < m_i$ . If a call to **net** is made in iteration  $m_i$  then again  $c_{\text{net}} \ell_{m_i} < f(W_{(m_i)-1}) = f(W_{m_{(i-1)}})$ . Thus, by **Lemma 3.2.6** and induction, we have

$$\ell_{m_i} < \frac{f(W_{m_{(i-1)}})}{c_{\text{net}}} \leq \frac{f(W_0) + 9\ell_{m_{(i-1)}}}{c_{\text{net}}} \leq \frac{f(W_0) + 9f(W_0)/\eta}{c_{\text{net}}} = \frac{1 + 9/\eta}{c_{\text{net}}} f(W_0) = \frac{f(W_0)}{\eta},$$

if  $\eta = \frac{c_{\text{net}}}{1 + 9/\eta}$ . This in turn is equivalent to  $\eta + 9 = c_{\text{net}}$ , which is true by definition.  $\blacksquare$

Setting  $c_{\text{net}} = 37$ , results in  $\eta = 28$ , and by **Lemma 3.2.7**, for all  $i$  that correspond to a net iteration,  $\ell_i \leq f(W_0)/28$ . By **Lemma 3.2.6**, for any net iteration  $i$ , we have

$$\left| f(W_i) - f(W_0) \right| \leq 9\ell_i \leq f(W_0)/3.$$

In particular, we conclude that  $|f(W_i) - f(W_0)| \leq f(W_0)/3$  for any iteration  $i$ . We thus get the following.

**Corollary 3.2.8.** *For  $c_{\text{net}} \geq 37$ , and any  $i$ , we have:*

- (A)  $(2/3)f(W_0) \leq f(W_i) \leq (4/3)f(W_0)$ .
- (B) If  $f(W_i) \in [x, y]$  then  $f(W_0) \in [(3/4)x, (3/2)y] \subseteq [x/2, 2y]$ .
- (C) If  $f(W_0) > 0$  then  $f(W_i) > 0$ .

**Lemma 3.2.9.** *For  $c_{\text{net}} \geq 37$ , given an instance  $(W, \Gamma)$  of a  $\varphi$ -NDP, **ndpAlg** $(W, \Gamma)$  returns an interval  $[x', y']$  containing  $f(W, \Gamma)$ , where  $\text{sprd}([x', y']) \leq 4 \max(\varphi, c_{\text{net}})$ .*

*Proof:* Consider the iteration of the while loop at which **ndpAlg** terminates, see **Figure 3.1**<sub>p11</sub>. If **Line 3** or **Line 4** get executed at this iteration, then the interval  $[x, y]$  was computed by the  $\varphi$ -decider, and has spread  $\leq \varphi$ . As such, by **Corollary 3.2.8**, the returned interval  $[x', y'] = [x/2, 2y]$  contains the optimal value, and its spread is  $\leq 4\varphi$ .

A similar argumentation holds if **Line 5** gets executed. Indeed, the returned interval contains the desired value, and its spread is  $4c_{\text{net}}$ .  $\blacksquare$

### 3.3. The result

In **Definition 3.1.3** required that the decision procedure runs in  $O(n)$  time. However, since our analysis also applies for decision procedures with slower running times, we first state the main result in these more general settings.

**Lemma 3.3.1.** *Given an instance of a  $\varphi$ -NDP defined by a set of  $n$  points in  $\mathbb{R}^d$ , one can get a  $\varphi$ -approximation to its optimal value, in  $O(T(n))$  expected time, where  $T(n)$  is the running time of the **decider** procedure, and  $\varphi \geq 3/2$ .*

*If  $\varphi = 1 + \varepsilon$ , for some  $\varepsilon \in (0, 1)$ , (i.e., **decider** is  $(1 + \varepsilon)$ -decider), then one can compute a  $(1 + \varepsilon)$ -approximation to the optimal value, and the expected running time is  $O(T(n) \log \varepsilon^{-1})$ .*

*Proof:* Let  $(W, \Gamma)$  be the given  $\varphi$ -NDP instance, and let **decider** be the corresponding  $\varphi$ -decider. By [Lemma 3.2.4](#) and [Lemma 3.2.9](#), in expected  $O(n)$  time, one can get a bounded spread interval  $[\gamma, c\gamma]$ , for some constant  $c \geq 1$  such that  $c = O(\varphi)$ , such that  $f = f(W, \Gamma) \in [\gamma, c\gamma]$ . If  $c \leq \varphi$  then we are done. Otherwise, perform a binary search over this interval.

Specifically, for  $i = 0, 1, \dots, m = \lceil \log_\varphi c \rceil$ , let  $\gamma_i = \varphi^i \gamma$  and let  $\gamma_{m+1} = c\gamma$ . Now perform a binary search over the  $\gamma_i$ 's using **decider**. If any of the calls returns an interval  $[x, \varphi x]$  that contains the optimal value, then  $f \leq \varphi x \leq \varphi f$  and so  $\varphi x$  can be returned as the desired  $\varphi$ -approximation. Otherwise, if  $f < \gamma_i$  or  $f > \gamma_i$  the binary search moves left or right, respectively. In the end, the search ends up with an interval  $(\gamma_i, \gamma_{i+1}) = (\gamma_i, \varphi \gamma_i)$  which must contain the optimal value, and again, its spread is  $\varphi$ , and it thus provide the required approximation.

The running time is dominated by the calls to the decider procedure (we are assuming here that  $T(n) = \Omega(n)$ ). Clearly, the number of calls to the decider performed by this algorithm is  $O(\log m) = O(\log \log_\varphi 4\varphi) = O(1)$ , if  $\varphi \geq 3/2$ . Otherwise, if  $\varphi = 1 + \varepsilon$ , then

$$O(\log m) = O(\log \log_\varphi 4\varphi) = O\left(\ln\left(1 + \frac{\ln 4}{\ln(1 + \varepsilon)}\right)\right) = O\left(\log\left(1 + \frac{1}{\varepsilon}\right)\right) = O\left(\log \frac{1}{\varepsilon}\right),$$

since  $\ln(1 + \varepsilon) \geq \ln e^{\varepsilon/2} = \varepsilon/2$ , for  $\varepsilon \in (0, 1/2)$ , as can be easily verified. ■

The following is a restatement of the main result in the settings of [Definition 3.1.3](#), where it was assume assumed  $T(n) = O(n)$ , or  $T(n) = O(n/\varepsilon^d)$  when  $\varphi = 1 + \varepsilon$ . In the latter case, one can remove the  $\log(1/\varepsilon)$  factor from the running time.

**Theorem 3.3.2.** *Given an instance of a  $\varphi$ -NDP defined by a set of  $n$  points in  $\mathbb{R}^d$ , one can get  $\varphi$ -approximate the optimal value, in expected  $O(n)$  time, assuming  $\varphi \geq 3/2$ .*

*For the case  $\varphi = (1 + \varepsilon) \in (1, 2)$ , given an  $(1 + \varepsilon)$ -decider, with running time  $O(n/\varepsilon^c)$ , one can  $(1 + \varepsilon)$ -approximate, in  $O(n/\varepsilon^c)$  expected time, the given  $(1 + \varepsilon)$ -NDP, where  $c \geq 1$  is some constant.*

*Proof:* [Lemma 3.3.1](#) readily implies the first half, as we are now assuming  $T(n) = O(n)$ . As for the second half, [Lemma 3.3.1](#) implies that one can get a  $(1 + \varepsilon)$ -approximation in expected  $O(n \log(1/\varepsilon)/\varepsilon^c)$  time. However, by using the same procedure as in the proof of [Lemma 3.3.1](#) but with exponentially decreasing values for  $\varepsilon$  in the binary search, a factor of  $\log(1/\varepsilon)$  can be avoided. This is a standard idea and was also used by Aronov and Har-Peled [[AH08](#)].

Let **decider** $_\varepsilon$  be our  $(1 + \varepsilon)$ -decider. If  $\varepsilon$  is set to any constant (say  $1/2$ ), then by [Lemma 3.2.4](#) and [Lemma 3.2.9](#), in expected linear time, one can get a bounded spread interval  $[\gamma, c\gamma]$ , for some  $c = O(\varphi)$ , such that  $f = f(W, \Gamma) \in \mathcal{I}_0 = [\gamma, c\gamma]$ .

Set  $c = 1 + \varepsilon_0$  and  $i = 1$ . The algorithm now proceeds in rounds, doing the following in the  $i$ th round:

(A) Assume that in the beginning of the  $i$ th iteration, we know that

$$f \in \mathcal{I}_{i-1} = \left[ x_{i-1}, y_{i-1} \right], \quad \text{where} \quad y_{i-1} = (1 + \varepsilon_{i-1})x_{i-1}.$$

If  $\varepsilon_{i-1} \leq \varepsilon$ , then we found the desired approximation, and the algorithm stops.

- (B) Set  $\varepsilon_i = \sqrt{1 + \varepsilon_{i-1}} - 1$  and  $m_i = (1 + \varepsilon_i)x_{i-1}$ .
- (C)  $R_i \leftarrow \mathbf{decider}_{\varepsilon_i}(m_i, \mathbf{W}, \Gamma)$ , see [Definition 3.1.2<sub>p9</sub>](#).
- (D) There are three possibilities:
  - (i) If  $R_i = "f < m_i"$ , then set  $\mathcal{I}_i \leftarrow [x_{i-1}, m_i]$ .
  - (ii) If  $R_i = "f > m_i"$ , then set  $\mathcal{I}_i \leftarrow [m_i, y_{i-1}]$ .
  - (iii) If  $R_i = "f \in \mathcal{I} = [z, (1 + \varepsilon_i)z]"$ , then set  $\mathcal{I}_i \leftarrow \mathcal{I}$ .
- (E)  $i \leftarrow i + 1$ .

In each round, the algorithm computes an interval of spread  $y_i/x_i = 1 + \varepsilon_i$  that contains  $f$ , and

$$\varepsilon_i = \sqrt{1 + \varepsilon_{i-1}} - 1 = \frac{1 + \varepsilon_{i-1} - 1}{\sqrt{1 + \varepsilon_{i-1}} + 1} \leq \frac{\varepsilon_{i-1}}{2}.$$

Since the main bottleneck in each iteration is calling  $\mathbf{decider}_{\varepsilon_i}$ , which runs in  $O(n/\varepsilon_i^c)$  time, the total running time is bounded by,

$$\sum_{i=1} O(n/\varepsilon_i^c) = O(n/\varepsilon^c),$$

since the sum behaves like a geometric series and the last term is  $O(n/\varepsilon^c)$ . ■

### 3.3.1. A deterministic algorithm for the bounded spread case

The random sampling in the algorithm of [Theorem 3.3.2](#) can be removed if the spread of  $\mathbf{P}$  is polynomially bounded. Indeed, then one can snap the points to a grid, so that the points have integer coordinates with polynomially bounded values. A compressed quadtree and a  $O(1)$ -WSPD of the snapped point set can be computed in linear time [[Har11](#), [Cha08](#)]. Given such a WSPD, one can compute for each point its (approximate) nearest neighbor, and compute the median distance in this list in linear time. Using this value in each iteration of  $\mathbf{ndpAlg}$  (instead of picking a random point and computing its nearest neighbor distance) results in an algorithm with linear running time. Note that the resulting algorithm is deterministic.

## 4. Applications

We now show that [Theorem 3.3.2](#) can be applied to a wide array of problems. In order to apply it, one needs to show that the given problem meets the requirements of an **NDP**. Often the inputs to the problems considered are unweighted point sets. As such, when we say such a problem is an **NDP**, it is actually defined for unit weighted point sets. We also continue making the simplifying assumption that the optimum value is strictly positive (see [Remark 3.2.1](#)).

### 4.1. $k$ -center clustering

Since computing a  $k$ -center clustering is an **NDP** problem ([Lemma 3.1.8](#)), plugging this into [Theorem 3.3.2](#), immediately yields a constant factor approximation to  $k$ -center in linear time. It is easy to convert such an approximation to a 2-approximation using a grid, see Har-Peled [[Har04a](#), Lemma 6.5]. Thus we get the following.

**Theorem 4.1.1.** *Given a set  $\mathbf{P}$  of  $n$  points in  $\mathbb{R}^d$ , and a parameter  $k$ ,  $1 \leq k \leq n$ , one can compute a 2-approximation to the optimal  $k$ -center clustering of  $\mathbf{P}$  in (expected) linear time.*



A result similar to [Theorem 4.1.1](#) was already known for the case  $k = O(n^{1/3}/\log n)$  [[Har04a](#)], and the above removes this restriction. In addition, the new algorithm and its analysis are both simpler than the previous algorithm.

## 4.2. The $k$ th smallest distance

**Claim 4.2.1.** *Let  $S$  and  $S'$  be subsets of  $\mathbb{R}$  of size  $n$ , such that  $S'$  is obtained by taking each value in  $S$  and increasing or decreasing it by less than  $\Delta$ . Let  $v$  and  $v'$  be the  $k$ th smallest values in  $S$  and  $S'$ , respectively. Then  $|v - v'| \leq \Delta$ .*

*Proof:* Suppose for contradiction that  $|v - v'| > \Delta$ . If  $v' - v > \Delta$  then  $S'$  has at least  $n - k + 1$  values strictly larger than  $v + \Delta$  which implies  $S$  has at least  $n - k + 1$  values strictly larger than  $v$ . Similarly if  $v - v' > \Delta$  then  $S'$  has at least  $k$  values strictly smaller than  $v - \Delta$  which implies  $S$  has at least  $k$  values strictly smaller than  $v$ .  $\blacksquare$

**Lemma 4.2.2.** *Let  $W$  be a weighted point set in  $\mathbb{R}^d$ , and let  $k > 0$  be an integer parameter. Let  $\binom{W}{2}$  denote the multi-set of pairwise distances determined by  $W$ .<sup>§</sup> Given an instance  $(W, k)$  the  **$k$ thDistance** problem asks you to output the  $k$ th smallest distance in  $\binom{W}{2}$ . Given such an instance, one can  $(1 + \varepsilon)$ -approximate the  $k$ th smallest distance in (expected) time  $O(n/\varepsilon^d)$ .*

*Proof:* Let  $f(W, k)$  be the function that returns the  $k$ th smallest distance. We prove that this function is an  $(1 + \varepsilon)$ -NDP (see [Definition 3.1.3<sub>p9</sub>](#)).

**Decider:** Given  $r$  and  $\varepsilon$ , build a grid where every cell has diameter  $\delta = \varepsilon r/8$ , and store the points of  $W$  in this grid. Now, for any non-empty grid cell  $\square$ , let  $\omega(\square)$  be the total weight of points in  $\square \cap W$ , and register this weight with all the grid cells in distance at most  $r$  from it, i.e. all cells in  $N_{\leq r}(\square)$  (where  $N_{\leq r}(\square)$  is defined analogously for grid cells as  $N_{\leq r}(p)$  was for points in [Definition 2.1.2](#)). Let  $\omega_N(\square)$  denote the total weight registered with a cell  $\square$  (which includes its own weight). Any point in  $\square$  determines  $\omega_N(\square) - 1$  distances to other points in  $N_{\leq r}(\square)$ . Therefore the total number of distances between points in  $\square$  and points in  $N_{\leq r}(\square)$  is  $\omega(\square)(\omega_N(\square) - 1)$ . Summing this over all cells (and dividing by 2 to avoid double counting) gives the quantity

$$S = \sum_{\square, \omega(\square) \neq 0} \frac{\omega(\square)}{2} (\omega_N(\square) - 1).$$

Note, that the count  $S$  includes all distances which are  $\leq r$ , and no distances  $> r + 2\delta$ . So let the desired  $k$ th distance be denoted by  $\ell_k$ . Thus, if  $k \leq S$  then  $\ell_k \leq r + 2\delta = (1 + \varepsilon/4)r$ . Similarly, if  $k > S$  then  $\ell_k > r$ . Therefore, to build a  $(1 + \varepsilon)$ -decider for distance  $r$ , we run the above counting procedure on  $r_1 = r/(1 + \varepsilon/3)$  and  $r_2 = r$ , and let  $S_1$  and  $S_2$  be the two counts computed, respectively. There are three cases:

- (I) If  $k \leq S_1$  then  $\ell_k \leq (1 + \varepsilon/4)r/(1 + \varepsilon/3) < r$ , and return this result.
- (II) If  $k \leq S_2$  then  $\ell_k \leq (1 + \varepsilon/4)r_2 = (1 + \varepsilon/4)r$ , then  $\ell_k \in [r/(1 + \varepsilon/3), (1 + \varepsilon/4)r]$ , and this interval has spread  $(1 + \varepsilon/4)(1 + \varepsilon/3) < 1 + \varepsilon$ . The decider returns that this interval contains the desired value.
- (III) The only remaining possibility is that  $\ell_k > r_2$ , and the decider returns this.

---

<sup>§</sup>In particular a point of weight  $m$  is viewed as  $m$  unit weight points when determining the values of  $\binom{W}{2}$ . For simplicity we assume that the  $k$ th distance in  $W$  is not zero (i.e. it is determined by two distinct points).

The running time of this decision procedure is  $O(n/\varepsilon^d)$ .

**Lipschitz:** Since the distance between any pair of points changes by at most  $2\Delta$  in a  $\Delta$ -drift, the Lipschitz condition holds by [Claim 4.2.1](#).

**Prune:** By assumption, the  $k$ th smallest distance is determined by two distinct weighted points  $\mathbf{p}$  and  $\mathbf{q}$ . Clearly these points are not in  $W^{\geq r}$  since  $d(\mathbf{p}, \mathbf{q}) = f(W, k) < r$ . So consider any point  $\mathbf{s} \in W^{\geq r}$ . Since removing  $\mathbf{s}$  does not remove the distance  $d(\mathbf{p}, \mathbf{q})$  from the set of remaining distances, all that is needed is to show how to update  $k$ . Clearly,  $\mathbf{s}$  contributes  $\omega(\mathbf{s}) \times \omega(W \setminus \{\mathbf{s}\})$  distances, all of value  $\geq r$ , to  $\binom{W}{2}$ , and  $\binom{\omega(\mathbf{s})}{2}$  pairwise distances of value zero. Thus, after removing  $\mathbf{s}$  from  $W$ , the new context is  $k - \binom{\omega(\mathbf{s})}{2}$ . ■

The algorithm of [Lemma 4.2.2](#) also works (with minor modifications) if we are interested in the  $k$ th distance between two sets of points (i.e., the bipartite version).

**Corollary 4.2.3.** *Given two sets  $P$  and  $Q$  of points in  $\mathbb{R}^d$ , of total size  $n$ , and parameters  $k$  and  $\varepsilon > 0$ , one can  $(1 + \varepsilon)$ -approximate, in expected  $O(n/\varepsilon^d)$  time, the following:*

- (A) *The  $k$ th smallest bichromatic distance. Formally, this is the smallest  $k$ th number in the multiset  $X = \{d(\mathbf{p}, \mathbf{q}) \mid \mathbf{p} \in P, \mathbf{q} \in W\}$ .*
- (B) *The closest bichromatic pair between  $P$  and  $Q$ .*

### 4.3. The $k$ th smallest $m$ -nearest neighbor distance

For a set  $P = \{\mathbf{p}_1, \dots, \mathbf{p}_n\} \subseteq \mathbb{R}^d$  of  $n$  points, and a point  $\mathbf{p} \in P$ , its  *$m$ th nearest neighbor* in  $P$  is the  $m$ th closest point to  $\mathbf{p}$  in  $P \setminus \{\mathbf{p}\}$ . In particular, let  $d_m(\mathbf{p}, P)$  denote this distance. Here, consider the multiset of these distances defined for each point of  $P$ ; that is,  $X = \{d_m(\mathbf{p}_1, P), \dots, d_m(\mathbf{p}_n, P)\}$ . Interestingly, for any  $m$ , we can approximate the  $k$ th smallest number in this set in linear time.

In order to use **ndpAlg**, one needs to generalize the problem to the case of weighted point sets. Specifically, a point  $\mathbf{p}$  of weight  $\omega(\mathbf{p})$  is treated as  $\omega(\mathbf{p})$  distinct points at the same location. As such, the set  $X$  from the above is actually a multiset containing  $\omega(\mathbf{p})$  copies of the value  $d_m(\mathbf{p}, P)$ . In particular, if  $\omega(\mathbf{p}) > m$  then  $d_m(\mathbf{p}, P) = 0$ .

**Theorem 4.3.1.** *Let  $W$  be a set of  $n$  weighted points in  $\mathbb{R}^d$ , and  $m, k, \varepsilon$  parameters. Then one can  $(1 + \varepsilon)$ -approximate, in expected  $O(n/\varepsilon^d)$  time, the  $k$ th smallest  $m$ -nearest neighbor distance in  $W$ . Formally, the algorithm  $(1 + \varepsilon)$ -approximates the  $k$ th smallest number in the multiset  $X = \{d_m(\mathbf{p}_1, W), \dots, d_m(\mathbf{p}_n, W)\}$*

*Proof:* We establish the required properties, see [Definition 3.1.3<sub>p9</sub>](#).

**Decider:** Let  $f(W, k, m)$  denote the desired quantity. We first need a  $(1 + \varepsilon)$ -decider for this problem.

To this end, given  $r, k, m, \varepsilon$  and  $W$  as input, create a grid with cells of diameter  $\varepsilon r/4$ , and mark for each point of  $W$  all the grid cells in distance at most  $r$  from it. Each non-empty grid cell has a count of the total weight of the points in distance at most  $r$  from it. Thus, each point of  $W$  can compute the total weight of all points which are approximately in distance at most  $r$  from it. Now, a point  $\mathbf{p} \in W$ , can decide in constant time if (approximately)  $d_m(\mathbf{p}, W) \leq r$ . If the number of points which declare  $d_m(\mathbf{p}, W) \leq r$  (where a point  $\mathbf{p}$  is counted  $\omega(\mathbf{p})$  times) is greater than  $k$  then the distance  $r$  is too large, and if it is smaller than  $k$  then  $r$  is too small. Being slightly more careful about the details (as was done in the proof of [Lemma 4.2.2](#)) one can verify this leads to a  $(1 + \varepsilon)$ -decider for this problem.

**Lipschitz:** Clearly the Lipschitz property holds in this case, as  $\Delta$ -drift only changes inter-point distances by at most an additive term of  $2\Delta$ .

**Prune:** We need to show that  $k$  can be updated properly for a weighted point  $\mathbf{p}$  whose nearest neighbor is further away than  $f(W, k, m)$ . If  $\omega(\mathbf{p}) > m$  then we are throwing away  $\omega(\mathbf{p})$  points, all with  $d_m(\mathbf{p}, W) = 0$ , and so  $k$  should be update to  $k - \omega(\mathbf{p})$  when  $\mathbf{p}$  is removed. Similarly, if  $\omega(\mathbf{p}) \leq m$  then all the points, that  $\mathbf{p}$  corresponds to, have  $d_m(\mathbf{p}, W)$  larger than the threshold, and  $k$  does not have to be updated.

Plugging this into [Theorem 3.3.2](#) implies the desired result. ■

**Remark 4.3.2.** [Theorem 4.3.1](#) can be easily extended to work in the bichromatic case. That is, there are two point sets  $P$  and  $Q$ , and we are interested in the  $m$ th nearest neighbor of a point  $\mathbf{p} \in P$  in the set  $Q$ . It is easy to verify that the same time bounds of [Theorem 4.3.1](#) hold in this case.

In particular, setting  $k = |P|$  and  $m = 1$  in the bichromatic case, the computed distance will be the minimum radius of the balls needed to be placed around the points of  $Q$  to cover all the points of  $P$  (or vice versa).

### 4.3.1. Exact nearest neighbor distances, and furthest nearest neighbor

Using [Theorem 4.3.1](#) with  $m = 1$  and  $\varepsilon = 1$ , results in a 2-approximation, in  $O(n)$  time, to the  ***$k$ th nearest neighbor distance*** of  $P$ . We now show how this can be converted into an exact solution.

So we have a quantity  $r$  that is larger than the  $k$ th nearest neighbor distance, but at most twice larger. We build a grid with a cell diameter being  $r/4$ . Clearly, any point whose nearest neighbor distance is at least the  $k$ th nearest neighbor distance must be the only point in its grid cell. For each such point  $\mathbf{p}$ , we compute its distance to all the points stored in its neighborhood  $N_{\leq r}(\mathbf{p})$ , by scanning the list of points associated with these cells. This computes for these “lonely” points their exact nearest neighbor distance. The  $k$ th smallest nearest neighbor distance will then be the  $(n - k)$ th largest distance computed. Clearly, every cell’s list of points get scanned a constant number of times, so overall the running time is linear. We summarize the result.

For a point set  $P$ , the ***furthest nearest neighbor distance***, is the maximum distance of a point of  $P$ , from the remaining points. Formally, it is  $\max_{\mathbf{p} \in P} d(\mathbf{p}, P)$ .


**Theorem 4.3.3.** *Let  $P$  be a set of  $n$  points in  $\mathbb{R}^d$ . For an integer  $k$ ,  $1 \leq k \leq n$ , one can compute exactly, in expected linear time, the  $k$ th nearest neighbor distance in  $P$ . In particular, setting  $k = n$ , one can compute the furthest nearest neighbor distance exactly, in expected linear time.*

## 4.4. The spanning forest partitions, and the $k$ th longest MST edge

The net computation provides a natural way to partition the data into clusters, and leads to a fast clustering algorithm. One alternative partition scheme is based on distance connectivity.

**Definition 4.4.1.** For a set of points  $P$  and a number  $r > 0$ , let  $\mathcal{C}_{\leq r}(P)$  be the  ***$r$ -connectivity clustering*** of  $P$ .

Specifically, it is a partition of  $P$  into connected components of the MST of  $P$  after all edges strictly longer than  $r$  are removed from it (alternatively, these are the connected components of the intersection graph where we replace every point of  $P$  by a disk of radius  $r/2$  centered at that point), see figure on the right. As such, a set in such a partition of  $P$  is referred to as a **connected component**.



Consider two partitions  $\mathcal{P}, \mathcal{Q}$  of  $P$ . The partition  $\mathcal{P}$  is a refinement of  $\mathcal{Q}$ , denoted by  $\mathcal{P} \sqsubseteq \mathcal{Q}$ , if for any set  $X \in \mathcal{P}$ , there exists a set  $Y \in \mathcal{Q}$  such that  $X \subseteq Y$ .

**Lemma 4.4.2.** *Given a set  $P \subseteq \mathbb{R}^d$ , and parameters  $r$  and  $\varepsilon$ , one can compute, in  $O(n/\varepsilon^d)$  time, a partition  $\mathcal{P}$ , such that  $\mathcal{C}_{\leq r}(P) \sqsubseteq \mathcal{P} \sqsubseteq \mathcal{C}_{\leq (1+\varepsilon)r}(P)$ .*

*Proof:* Build a grid with every cell having diameter  $\varepsilon r/4$ . For every point in  $P$  mark all the cells in distance  $r/2$  from it. That takes  $O(n/\varepsilon^d)$  time. Next, for every marked cell create a node  $v$ , and let  $V$  be the resulting set of nodes. Next, create a bipartite graph  $G = (V \cup P, E)$  connecting every node of  $V$ , to all the points of  $P$  marking its corresponding grid cell. Now, compute the connected components in  $G$ , and for each connected component extract the points of  $P$  that belong to it. Clearly, the resulting partition  $\mathcal{P}$  of  $P$ , is such that any two points in distance  $\leq r$  are in the same set (since for such points there must be a grid cell in distance  $\leq r/2$  from both). Similarly, if the distance between two subsets  $X, Y \subseteq P$  is at least  $(1 + \varepsilon)r$ , then they are in different connected components of  $\mathcal{P}$ . In particular, in order for  $X$  and  $Y$  to be in the same component, there must exist points  $x \in X$  and  $y \in Y$  which marked the same grid cell. However, such a grid cell would need to be in distance at most  $r/2$  from both  $x$  and  $y$ , implying  $d(x, y) \leq r/2 + r/2 + \varepsilon r/4 < (1 + \varepsilon)r$ , a contradiction. Thus, this is the desired partition. Clearly, this takes  $O(n/\varepsilon^d)$  time overall.  $\blacksquare$

**Theorem 4.4.3.** *Given a set of  $n$  points  $P$  in  $\mathbb{R}^d$ , and parameters  $k$  and  $\varepsilon$ , one can output, in expected  $O(n/\varepsilon^d)$  time, a  $(1 + \varepsilon)$ -approximation to the  $k$ th longest (or  $k$ th shortest) edge in the MST of  $P$ .*

*Proof:* Consider the cost function  $\ell_{MST}(P, k)$  which returns the  $k$ th longest edge of the MST of  $P$ . It is easy to verify that  $\ell_{MST}(P, k)$  is the minimum  $r$  such that  $\mathcal{C}_{\leq r}(P)$  has  $k$  connected components. For the sake of simplicity of exposition, we assume all the pairwise distances of points in  $P$  are distinct. We need to fill-in/verify the properties of [Definition 3.1.3<sub>p9</sub>](#):

**Decider:** The  $(1 + \varepsilon)$ -decider for  $\ell_{MST}(P, k)$  uses the algorithm of [Lemma 4.4.2](#). Specifically, for a specified  $r > 0$ , compute a partition  $\mathcal{P}$ , such that  $\mathcal{C}_{\leq r}(P) \sqsubseteq \mathcal{P} \sqsubseteq \mathcal{C}_{\leq (1+\varepsilon)r}(P)$ . Now, if  $\mathcal{P}$  has  $\leq k$  connected components, then  $\ell_{MST}(P, k) \leq (1 + \varepsilon)r$ . Similarly, if  $\mathcal{P}$  has  $> k$  connected components, then  $\ell_{MST}(P, k) > r$ . By calling on two values, and using slightly smaller  $\varepsilon$ , it is straightforward to get the desired  $(1 + \varepsilon)$ -decider, as was done in [Lemma 4.2.2<sub>p17</sub>](#).

**Lipschitz:** Let  $P'$  be some  $\Delta$ -drift of  $P$ , and let  $G$  (resp.  $G'$ ) denote the complete graph on  $P$  (resp.  $P'$ ). For an edge  $e \in E(G)$  (or in  $E(G')$ ), let  $\ell(e)$  denote the length of  $e$ . For the sake of simplicity of exposition, we consider  $P'$  and  $P$  to be multi-sets of the same cardinality (i.e.,  $n = |P| = |P'| > k$ ). (Note this implies some of the edges may have length 0.) For a point  $p \in P$ , let  $p'$  denote its image in  $P'$  after the  $\Delta$ -drift. Let  $e_1, \dots, e_{n-1}$  be the edges in the MST of  $G$  sorted by increasing order by their length. Let  $e'_1, \dots, e'_{n-1}$  denote the corresponding edges in  $G'$ , i.e. for  $e = uv$ , the corresponding edge is  $e' = u'v'$ .

Observe that the subgraph of  $G'$  having  $e'_1, \dots, e'_i$  as edges, is a forest with  $n - i$  connected components. As such, we have that  $\ell_{MST}(P', n - i) \leq \max_{j=1}^i \ell(e'_j)$ , since by the correctness of the

Kruskal's algorithm, the first  $i$  edges of the MST form such a forest with the maximum weight edge being minimum (among all such forests). As such, we have that

$$\begin{aligned}\ell_{MST}(\mathbf{P}', n - i) &\leq \max_{j=1}^i \ell(e'_j) \leq \max_{j=1}^i (\ell(e_j) + 2\Delta) = 2\Delta + \max_{j=1}^i \ell(e_j) \\ &= 2\Delta + \ell_{MST}(\mathbf{P}, n - i).\end{aligned}$$

A symmetric argument implies that  $\ell_{MST}(\mathbf{P}, n - i) \leq 2\Delta + \ell_{MST}(\mathbf{P}', n - i)$ . Setting  $i = n - k$ , we have that  $|\ell_{MST}(\mathbf{P}, k) - \ell_{MST}(\mathbf{P}', k)| \leq 2\Delta$ , as desired.

**Prune:** Consider a point  $\mathbf{p} \in \mathbf{W}$ , such that  $\mathbf{d}(\mathbf{p}, \mathbf{W}) \geq r$ , and  $r > \ell_{MST}(\mathbf{W}, k)$ . This implies that all the MST edges adjacent to  $\mathbf{p}$  are of distance  $\geq r$ . Assume there are  $j$  such edges, and observe that removing these edges creates  $j + 1$  connected components of the MST, such that each connected component is of distance  $\geq r$  from each other (indeed, otherwise, one of the removed edges should not be in the MST, as can be easily verified). Thus, if we delete  $\mathbf{p}$ , and recompute the MST, we will get  $j$  new edges that replace the deleted edges, all of these edges are of length  $\geq r$ . That is, deleting  $\mathbf{p}$ , and decreasing  $k$  by one, ensures that the target value remains the same in the pruned instance.

This establish that  $\ell_{MST}(\mathbf{P}, k)$  is an  $(1 + \varepsilon)$ -NDP, and by plugging this into [Theorem 3.3.2](#), we get the desired result.  $\blacksquare$

**Remark 4.4.4.** Note that in the above, as in other parts of the paper, we treat a point of weight  $\omega$  as  $\omega$  copies of the same point. This naturally leads to edges of length zero in the MST, which are being handled correctly by the algorithm.

## 4.5. Computing the minimum cluster

### 4.5.1. Sketchable families

**Definition 4.5.1 (Upward Closed Set System).** Let  $\mathbf{P}$  be a finite ground set of elements, and let  $\mathcal{F}$  be a family of subsets of  $\mathbf{P}$ . Then  $(\mathbf{P}, \mathcal{F})$  is an *upward closed set system* if for any  $X \in \mathcal{F}$  and any  $Y \subseteq \mathbf{P}$ , such that  $X \subseteq Y$ , we have that  $Y \in \mathcal{F}$ . Such a set system is a *sketchable family*, if for any set  $S \subseteq \mathbf{P}$  there exists a constant size *sketch*  $\text{sk}(S)$  such that:

- (A) For any  $S, T \subseteq \mathbf{P}$  that are disjoint,  $\text{sk}(S \cup T)$  can be computed from  $\text{sk}(S)$  and  $\text{sk}(T)$  in  $O(1)$  time. We assume the sketch of a singleton can be computed in  $O(1)$  time, and as such the sketch of a set  $S \subseteq \mathbf{P}$  can be computed in  $O(|S|)$ .
- (B) There is a membership oracle for the set system based on the sketch. That is, there is a procedure **orac** such that given the sketch of a subset  $\text{sk}(S)$ , **orac** returns whether  $S \in \mathcal{F}$  or not, in  $O(1)$  time.

An example for such a sketchable family, is the set system  $(\mathbf{P}, \mathcal{F})$ , where  $S \subseteq \mathbf{P}$  is in  $\mathcal{F}$  if  $|S| \geq 10$ . Here the sketch of the set is simply the number of elements in the set, and combining two sketches  $\text{sk}(S)$  and  $\text{sk}(T)$  is adding the numbers to get  $\text{sk}(S \cup T)$  (for  $S \cap T = \emptyset$ ).

We will be interested in two natural problems induced by such a family: (i) smallest cluster – find the smallest set in the family with certain properties, and (ii) min-max clustering – find disjoint sets in the family such that they cover the original set, and the maximum price of these sets is minimized.

**Remark 4.5.2.** In the following, we consider sketchable families  $(\mathbf{P}, \mathcal{F})$ , where  $\mathbf{P}$  is a set of Euclidean points. In this case,  $\mathbf{P} = \{\mathbf{p}_1, \dots, \mathbf{p}_n\}$  can be viewed as a set of elements, where each  $\mathbf{p}_i$  is additionally

given a realization at some location in  $\mathbb{R}^d$ . Specifically, changing the locations of the realizations does not change the subsets in  $\mathcal{F}$ . This is required to ensure the Lipschitz property holds. Specifically, sketches remain the same after a  $\Delta$ -drift. Furthermore, we do not have to maintain the sketches under deletions (of the underlying points), as pruned points corresponds to complete sketches that are being thrown away, see [Remark 4.5.8](#).

**Remark 4.5.3.** Note that it is not necessary for the sketches to have  $O(1)$  size or the oracle to run in  $O(1)$  time, however, assuming otherwise will affect the running time of **ndpAlg** if it is used to solve a problem involving a sketchable family.

**Example 4.5.4.** Consider associating a positive  $k$ -dimensional vector  $\psi_{\mathbf{p}}$  with each point  $\mathbf{p} \in \mathbf{P}$  (a vector is *positive* if it is non-zero, and all its coordinates are non-negative). A *positive linear inequality* is an inequality of the form  $\sum_i \alpha_i x_i \geq c$ , where the coefficients  $\alpha_1, \dots, \alpha_k$  are all non-negative. For such a linear inequality, consider the set system  $(\mathbf{P}, \mathcal{F})$ , where a set  $\mathbf{Q}$  is in  $\mathcal{F}$  if the linear inequality holds for the vector  $\sum_{\mathbf{p} \in \mathbf{Q}} \psi_{\mathbf{p}}$ . Clearly, this family is a sketchable family, the sketch being the sum of the vectors associated with the points of the set (here  $k$  is assumed to be a constant).

It is easy to verify that sketchable families are closed under finite intersection. Specifically, given a collection of  $m$  such positive inequalities, the family of sets such that their sketch vector complies with all these inequalities is a sketchable family (of course, checking if a set, given its sketch, is in the family naively would take  $O(mk)$  time).

As a concrete application, consider the scenario where every element in  $\mathbf{P}$  has  $k = 4$  attributes. One might be interested in subsets,  $\mathbf{Q}$ , such that the sum of the first two attributes of the vector  $\sum_{\mathbf{p} \in \mathbf{Q}} \psi_{\mathbf{p}}$  is at least 1, and the sum of last two attributes is at least 2.

Of course, in general, a membership oracle for the attributes space that has the property that if  $\psi$  is valid then  $\psi + \psi'$  is also valid, for any positive  $\psi'$ , would define a sketchable family. As a concrete example, consider the non-linear (and not convex) condition that the sum of at least two attributes is larger than 1. Clearly, this defines a sketchable family.

Clearly the above definition of sketchable family is very general and widely applicable. Next, we show how **ndpAlg** can be used to approximate certain objectives over sketchable families.

## 4.5.2. Min cluster

We now consider the problem of minimizing the cluster radius of a subset of a given point set subject to inclusion in a sketchable family. Specifically, we consider the case when the cluster is defined by a ball of radius  $r$  or when the cluster is defined by a connected component of  $\mathcal{C}_{\leq r}(\mathbf{P})$  for a radius  $r$ . Note that as a ball or component grows, both the set of points the cluster contains and the inclusion of this set of points in the sketchable  $\mathcal{F}$ , are monotone properties. This correspondence is what allows us to apply our general framework.

**Theorem 4.5.5.** *For a set of  $n$  points  $\mathbf{P} \subseteq \mathbb{R}^d$ , and a sketchable family  $(\mathbf{P}, \mathcal{F})$ , one can  $(1 + \varepsilon)$ -approximate, in expected  $O(n/\varepsilon^d)$  time, the radius of the smallest ball,  $\mathbf{b}$ , such that  $\mathbf{b} \cap \mathbf{P} \in \mathcal{F}$ .*

*Proof:* Let  $f(\mathbf{P})$  be the diameter of the smallest ball  $\mathbf{b}$  in  $\mathbb{R}^d$  such that  $\mathbf{P} \cap \mathbf{b} \in \mathcal{F}$ . We claim that  $f(\cdot)$  is an **NDP**, see [Definition 3.1.3<sub>p9</sub>](#). The Lipschitz and Prune properties readily hold for  $f(\cdot)$ . Specifically,  $f(\mathbf{P})$  was chosen to be the diameter rather than the radius so that the Prune property holds. Note that below, for the Decider property, we refer to approximating the radius, which gives us the same approximation to the diameter.

For the Decider property, given  $r$  and  $\varepsilon > 0$ , construct a grid with cells of diameter  $\varepsilon r/4$ , and register each point of  $\mathbf{P}$  in all the grid cells in distance at most  $r$  from it. If there is a ball,  $\mathbf{b}$ , of radius  $r$  such that  $\mathbf{b} \cap \mathbf{P} \in \mathcal{F}$  then the set of points registered with the grid cell containing the center of this ball will be a superset of  $\mathbf{b} \cap \mathbf{P}$  and hence the set is in  $\mathcal{F}$ , by the upward closed property of sketchable families. Moreover, the set registered at this grid cell requires a ball of radius at most  $r + 2\varepsilon r/4$  (centered at any point in this grid cell) to cover it. Thus, the decision procedure simply checks the set of points associated with each grid cell to see whether or not it is in  $\mathcal{F}$ . The definition of sketchable families implies this can be done in linear time in the total sizes of these sets (i.e.,  $O(n/\varepsilon^d)$ ). Furthermore, if there is no ball of radius  $(1 + \varepsilon)r$  whose point set is in  $\mathcal{F}$  then the decision procedure would fail to find such a cell. Thus, this is a  $(1 + \varepsilon)$ -decision procedure, and its running time is  $O(n/\varepsilon^d)$ .

Plugging this into the algorithm of [Theorem 3.3.2](#) implies the result. ■

**Remark 4.5.6.** In the algorithm of the above theorem, as the algorithm progresses, every weighted point that is computed also has an associated sketch from the original points that it corresponds to. One can view this collection of sketches as either an attribute in addition to the weights, or as being passed down via the context.

The following is a sample of what the above theorem implies.

**Corollary 4.5.7.** *We can  $(1 + \varepsilon)$ -approximate, in  $O(n/\varepsilon^d)$  time, the following problems for a set of  $n$  points in  $\mathbb{R}^d$ :*

- (A) *The smallest ball containing  $k$  points of  $\mathbf{P}$ .*
- (B) *The points of  $\mathbf{P}$  are weighted, and there is an additional threshold  $\alpha$ . Compute the smallest ball containing points with total weight at least  $\alpha$ .*
- (C) *If the points of  $\mathbf{P}$  are colored by  $k$  colors, the smallest ball containing points of  $\mathbf{P}$ , such that they are colored by at least  $t$  different colors. (Thus, one can find the smallest non-monochromatic ball [ $t = 2$ ], and the smallest ball having all colors [ $t = k$ ].) The running time is  $O(nk/\varepsilon^d)$ , as the sketch here is a  $k$ -dimensional vector.*

**Remark 4.5.8.** A careful inspection of our algorithm reveals that we only merge sketches or throw them away. Specifically, a pruning stage of our algorithm corresponds to throwing away some of the sketches, and netting corresponds to merging together some of the sketches. See also [Remark 4.5.2](#).

### 4.5.3. Min cluster in the spanning forest

Note that  $\mathcal{C}_{\leq r}(\mathbf{P})$  is a monotone partition as  $r$  increases (i.e., components merge, but never split apart). It is natural to ask what is the minimum  $r$ , for which there is a connected component in  $\mathcal{C}_{\leq r}(\mathbf{P})$  that is in a sketchable family.

**Theorem 4.5.9.** *For a set of  $n$  points  $\mathbf{P} \subseteq \mathbb{R}^d$ , and a sketchable family  $(\mathbf{P}, \mathcal{F})$ , one can  $(1 + \varepsilon)$ -approximate, in expected  $O(n/\varepsilon^d)$  time, the minimum  $r$ , such that there is a connected component in  $\mathcal{C}_{\leq r}(\mathbf{P})$  that is in  $\mathcal{F}$ .*

*Proof:* The target function  $f(\mathbf{P})$  is the smallest  $r$  such that  $\mathcal{C}_{\leq r}(\mathbf{P})$  contains a set that is in  $\mathcal{F}$ . Let  $r_{opt} = f(\mathbf{P})$ . The algorithm first checks if any of the input points by themselves have the desired property, if so then  $r_{opt} = 0$  and the algorithm returns 0. (For the case that  $\mathbf{P}$  is a multiset, see [Remark 3.2.2<sub>p12</sub>](#).) As such, we can assume that  $f(\mathbf{P}) > 0$ , and furthermore, by [Corollary 3.2.8 \(C\)](#),  $f(\cdot)$  is always non-zero during the algorithm execution. Again, we need to verify the properties of [Definition 3.1.3<sub>p9</sub>](#) for  $f(\cdot)$ :

**Decider:** One can  $(1 + \varepsilon/4)$ -approximate the connected components of  $\mathcal{C}_{\leq r/(1+\varepsilon/3)}(\mathbf{P})$ , using the algorithm of [Lemma 4.4.2](#), and for each approximate connected component, use their sketch to decide if any of them is in  $\mathcal{F}$ . If so, return that  $r$  is larger than the optimal value. Otherwise, perform a  $(1 + \varepsilon/3)$ -approximate computation of  $\mathcal{C}_{\leq r}(\mathbf{P})$ , and if one of its clusters is in  $\mathcal{F}$ , then return that the desired value is in the interval  $[r/(1 - \varepsilon/3), r(1 + \varepsilon/3)]$ . Otherwise, return that  $r$  is smaller than the optimal value. Clearly, this is  $(1 + \varepsilon)$ -decider that works in  $O(n/\varepsilon^d)$  time.

**Lipschitz:** Let  $r' = f(\mathbf{P})$ . By the definition of  $f(\cdot)$ , there exists  $k$ , such that  $r' = \ell_{MST}(\mathbf{P}, k)$ . Now, arguing as in the proof of [Theorem 4.4.3<sub>p20</sub>](#), implies that  $f(\cdot)$  has the desired Lipschitz property.

**Prune:** Consider a point  $\mathbf{p} \in \mathbf{W}$ , such that  $d(\mathbf{p}, \mathbf{W}) \geq r$ , and  $r > f(\mathbf{W}) > 0$ . Now,  $\mathbf{p}$  by itself can not have the desired property, as this would imply that  $f(\mathbf{W}) = 0$ . This implies that in  $\mathcal{C}_{\leq r_{opt}}(\mathbf{W})$  the point  $\mathbf{p}$  is by itself (and not in any cluster that is in  $\mathcal{F}$ ). As such, throwing  $\mathbf{p}$  away does not change the target function. Formally,  $f(\mathbf{W}) = f(\mathbf{W} \setminus \{\mathbf{p}\})$ .

Plugging this into [Theorem 3.3.2](#) implies the result. ■

One natural application for [Theorem 4.5.9](#), is for ad hoc wireless networks. Here, we have a set  $\mathbf{P}$  of  $n$  nodes and their locations (say in the plane), and each node can broadcast in a certain radius  $r$  (the larger the  $r$  the higher the energy required, so naturally we would like to minimize it). It is natural now to ask for the minimum  $r$  such that one of the connected components in the resulting ad hoc network has some desired property. For example, in  $O(n/\varepsilon^d)$  time, we can  $(1 + \varepsilon)$ -approximate the smallest  $r$  such that:

- (A) One of the connected components of  $\mathcal{C}_{\leq r}(\mathbf{P})$  contains half the points of  $\mathbf{P}$ , or more generally if the points are weighted, that one of connected component contains points of total weight at least  $\alpha$ , for a prespecified  $\alpha$ .
- (B) If the points are colored, the desired connected component contains all the colors (for example, each color represent some fraction of the data, and the cluster can recover the data if all the pieces are available), or at least two colors, or more generally a different requirement on each color.

## 4.6. Clustering for monotone properties

**Definition 4.6.1 (Min-Max Clustering).** We are given a sketchable family  $(\mathbf{P}, \mathcal{F})$ , and a cost function  $g : 2^{\mathbf{P}} \rightarrow \mathbb{R}^+$ . We are interested in finding disjoint sets  $S_1, \dots, S_m \in \mathcal{F}$ , such that (i)  $\bigcup_i S_i = \mathbf{P}$ , and (ii)  $\max_i g(S_i)$  is minimized. We will refer to the partition realizing the minimum as the *optimal clustering* of  $\mathbf{P}$ .

**Theorem 4.6.2.** *Let  $\mathbf{P}$  be a set of points in  $\mathbb{R}^d$ , and let  $(\mathbf{P}, \mathcal{F})$  be a sketchable family. For a set  $\mathbf{W} \in \mathcal{F}$ , let  $r_{\min}(\mathbf{W})$  be the radius of the smallest ball centered at a point of and enclosing  $\mathbf{W}$ . One can  $(4 + \varepsilon)$ -approximate, in expected  $O(n/\varepsilon^d)$  time, the min-max clustering under  $r_{\min}$  of  $\mathbf{P}$ .*

*That is, one can cover  $\mathbf{P}$  by a set of balls, and assign each point of  $\mathbf{P}$  to one of these balls, such that the set of points assigned to each ball is in  $\mathcal{F}$ , and the maximum radius of any of these balls is a  $(4 + \varepsilon)$ -approximation to the minimum radius used by any such cover.*

*Proof:* Let  $\mathcal{P}_{opt}$  be the optimal partition with radius  $r_{opt}$ , and consider an  $r$ -net  $\mathcal{N}$  for  $r \geq 4r_{opt}$ , computed using [Corollary 2.2.2](#). Consider a point  $\mathbf{p} \in \mathcal{N}$ , and let  $\mathbf{P}_{\mathcal{N}}[\mathbf{p}]$  be the set of points of  $\mathbf{P}$  assigned to  $\mathbf{p}$  by the nearest net-point assignment.

Next, consider the cluster  $\mathbf{W} \in \mathcal{P}_{opt}$  that contains it. Clearly,  $\text{diam}(\mathbf{W}) \leq 2r_{opt}$ , and the distance of  $\mathbf{p}$  from all other net points in  $\mathcal{N}$  is at least  $4r_{opt}$ . It follows that  $\mathbf{W} \subseteq \mathbf{P}_{\mathcal{N}}[\mathbf{p}]$ , and since  $\mathbf{W} \in \mathcal{F}$ , it follows that  $\mathbf{P}_{\mathcal{N}}[\mathbf{p}] \in \mathcal{F}$ .



A 4-decider for this problem works by computing the  $4r$ -net  $\mathcal{N}$ , and for each  $\mathbf{p} \in \mathcal{N}$ , checking the sketchable property for the set  $\mathbf{P}_{\mathcal{N}}[\mathbf{p}]$ . It is easy to verify that the properties of [Definition 3.1.3<sub>p9</sub>](#) hold in this case. In particular, throwing a far away isolated point corresponds to a cluster that already fulfill the monotone property, and it is too far away to be relevant. Namely, computing  $r_{opt}$  is an **NDP** and so plugging this into [Theorem 3.3.2](#) implies the result.  $\blacksquare$

#### 4.6.1. Lower bounded center clustering

If the required sketchable property is that every cluster contains at least  $k$  points, then [Theorem 4.6.2](#) approximates the *lower bounded center* problem. That is, one has to cover the points by balls, such that every cluster (i.e., points assigned to a ball) contains at least  $k$  points. The price of this clustering is the radius of the largest ball used. A 2-approximation to this problem is known via the usage of flow [[APF<sup>+</sup>10](#)] but the running time is super quadratic. Recently, the authors showed a similar result to [Theorem 4.6.2](#) with running time (roughly)  $O(n \log n)$  [[ERH12](#)]. This paper also shows that this problem cannot be approximated to better than (roughly) 1.8 even for points in the plane.

**Corollary 4.6.3.** *Let  $\mathbf{P}$  be a set of points in  $\mathbb{R}^d$ , and let  $k$  and  $\varepsilon > 0$  be parameters. One can  $(4 + \varepsilon)$ -approximate the lower bounded center clustering in  $O(n/\varepsilon^d)$  time.*

#### 4.6.2. Other clustering problems

One can plug-in any sketchable family into [Theorem 4.6.2](#). For example, if the points have  $k$  colors, we can ask for the min-max radius clustering, such that every cluster contains (i) all colors, (ii) at least two different colors, or (iii) a different requirement on each color, etc.

As another concrete example, consider that we have  $n$  customers in the plane, and each customer is interested in  $k$  different services (i.e., there is a  $k$ -dimensional vector associated with each customer specifying his/her demand). There are  $t$  types of service centers that can be established, but each such center type requires a minimum level of demand in each of these  $k$  categories (i.e., each type is specified by a minimum demand  $k$ -dimensional vector, and a set of customers can be the user base for such a service center if the sum of their demands vector is larger than this specification). The problem is to partition the points into clusters (of minimum maximum radius), such that for every cluster there is a valid service center assigned to it. Clearly, this falls into the framework of [Theorem 4.6.2](#), and can be  $(4 + \varepsilon)$ -approximated in  $O(nkt/\varepsilon^d)$  time.

#### 4.6.3. Clustering into spanning forests

One can get a similar result to [Theorem 4.6.2](#) for the connectivity version of clustering of  $\mathbf{P}$ . Formally, a set of points  $W \subseteq \mathbf{P}$  is  *$r$ -valid* if  $W$  is contained in some set of  $\mathcal{C}_{\leq r}(\mathbf{P})$ . Given a sketchable family  $(\mathbf{P}, \mathcal{F})$ , a partition  $\mathcal{P}$  of  $\mathbf{P}$  is an  *$r$ -connected clustering* if all the sets in  $\mathcal{P}$  are in  $\mathcal{F}$ , and are  $r$ -valid.

**Theorem 4.6.4.** *Let  $\mathbf{P}$  be a set of points in  $\mathbb{R}^d$ , and let  $(\mathbf{P}, \mathcal{F})$  be a sketchable family. One can  $(1 + \varepsilon)$ -approximate  $r_{opt}$ , in expected  $O(n/\varepsilon^d)$  time, where  $r_{opt}$  is the minimum value such that there is a  $r_{opt}$ -connected clustering of  $\mathbf{P}$ .*

*Proof:* Let  $f(\cdot)$  be the target function in this case. We verify the properties of [Definition 3.1.3<sub>p9</sub>](#):

**Decider:** Given  $r$ , we use [Lemma 4.4.2](#) to compute a partition  $\mathcal{P}$  of  $\mathbf{P}$  such that  $\mathcal{C}_{\leq r}(\mathbf{P}) \sqsubseteq \mathcal{P} \subseteq \mathcal{C}_{\leq (1+\varepsilon)r}(\mathbf{P})$ . If the sketchable property holds for each cluster of  $\mathcal{P}$ , then return that  $r$  is too large, otherwise return that it is too small. As for the quality of this decider, observe that if the optimal

partition has a cluster  $W$  that uses points from two different clusters of  $\mathcal{P}$ , then  $W$  is not  $r$ -valid, as otherwise these two points would be in the same cluster of  $\mathcal{P}$  (namely,  $r_{opt} > r$ ).

**Lipschitz:** Follows easily by arguing as in [Theorem 4.4.3<sub>p20</sub>](#).

**Prune:** If an isolated point exists, then it can be thrown away because the cluster of original points it corresponds to, is a valid cluster that can be used in the final clustering, and it does not interact with any other clusters.

Plugging this into [Theorem 3.3.2](#) now implies the result. ■

A nice application of [Theorem 4.6.4](#) is for ad hoc networks. Again, we have a set  $P$  of  $n$  wireless clients, and some of them are base stations; that is, they are connected to the outside world. We would like to find the minimum  $r$ , such that each connected component of  $\mathcal{C}_{\leq r}(P)$  contains a base station.

## 4.7. Smallest non-zero distance

Given a point set  $P$ , it was already shown in [Section 4.3.1](#) that, with a small amount of post processing, [ndpAlg](#) can be used to compute the closest pair distance exactly in expected linear time. If one allows  $P$  to contain duplicate points (i.e.  $P$  is a multiset) then a couple modifications to [ndpAlg](#) must be made. First, modify the algorithm so that for the selected point  $p$  it finds the closest point to  $p$  that has non-zero distance from it. Second, modify [delFar](#) (see [Lemma 2.3.1<sub>p9</sub>](#)) so that a singleton cell means all points in the cell have the same location, rather than meaning the cell contains a single point. This can easily be checked for all cells in overall linear time. With these two modifications, the above algorithm works verbatim, and we get the following.

**Lemma 4.7.1.** *Let  $P$  be a multiset of weighted points in  $\mathbb{R}^d$ . Then one can solve the **SmallestNonZeroDist** problem exactly for  $P$ , in expected linear time. In particular, if  $P$  contains no duplicates then this corresponds to computing the closest pair distance.*

Interestingly, the algorithm of [Lemma 4.7.1](#) is a prune-and-search algorithm, as the net stage never gets executed. Observe, that it is not hard to extend the algorithm of Golin *et al.* [[GRSS95](#)] to solve this variant, and the result of [Lemma 4.7.1](#) is included in the paper only for the sake of completeness. In particular, the resulting (simplified) algorithm is similar to the algorithm of Khuller and Matias [[KM95](#)].

## 5. Linear Time with high probability

In [Section 3](#), we presented an algorithm ([ndpAlg](#), see [Figure 3.1<sub>p11</sub>](#)) that provides a constant factor approximation to a wide variety of problems in expected linear time. Here we present an extension of that algorithm that runs in linear time with high probability.

Our algorithm works by computing a quantity that is not directly related to the problem at hand (but in linear time!), and then relating this quantity to the desired quantity, enabling us to compute it in linear time. Specifically, the resulting algorithm is somewhat counterintuitive and bizarre. As such, to make the algorithm description more accessible to the reader, we introduce the basic ideas in [Section 5.1](#). We describe the algorithm in detail in [Section 5.2](#), and analyze it in [Section 5.3](#).

### 5.1. Preliminaries

**Notation** See [Table 5.1](#) for a summary of the notation used in this part of the paper. As it is sufficient for our purposes, for simplicity (in this section)  $P$  will always be an unweighted points set, i.e.  $n = \omega(P) = |P|$ . The algorithm can be easily extended to handle the weighted case.

$d(\mathbf{p}, \mathbf{P})$	$\min_{\mathbf{q} \in \mathbf{P} \setminus \{\mathbf{p}\}} d(\mathbf{p}, \mathbf{q})$	Distance from $\mathbf{p}$ to its nearest neighbor in $\mathbf{P} \setminus \{\mathbf{p}\}$ .
$d_i(\mathbf{p}, \mathbf{P})$		Distance of $\mathbf{p}$ to its $i$ th nearest neighbor in $\mathbf{P} \setminus \{\mathbf{p}\}$ .
$\mathbf{D}(\mathbf{X}, \mathbf{P})$	$\{d(\mathbf{x}, \mathbf{P}) \mid \mathbf{x} \in \mathbf{X}\}$	Multi-set of distances between points in $\mathbf{X}$ and nearest neighbors in $\mathbf{P}$ .
$\mathbf{D}_i(\mathbf{X}, \mathbf{P})$	$\{d_i(\mathbf{x}, \mathbf{P}) \mid \mathbf{x} \in \mathbf{X}\}$	Multi-set of distances between point in $\mathbf{X}$ and their $i$ th nearest neighbors in $\mathbf{P}$ .
$d^\delta(\mathbf{X}, \mathbf{P})$		Value of rank $\lfloor \delta  \mathbf{X}  \rfloor$ in $\mathbf{D}(\mathbf{X}, \mathbf{P})$ , for some $\delta \in [0, 1]$ .
$d_i^\delta(\mathbf{X}, \mathbf{P})$		Value of rank $\lfloor \delta  \mathbf{X}  \rfloor$ in $\mathbf{D}_i(\mathbf{X}, \mathbf{P})$ , for some $\delta \in [0, 1]$ .

Table 5.1: Notation used. If  $\mathbf{X} = \mathbf{P}$ , the second argument is omitted. Thus,  $d_i^\delta(\mathbf{P}) = d_i^\delta(\mathbf{P}, \mathbf{P})$ .

### 5.1.1. Basic idea

**Sample in the middle** Observe that if one could sample a nearest neighbor distance that lies in the interval  $[d^\delta(\mathbf{P}), d^{1-\delta}(\mathbf{P})]$  for some fixed constant  $\delta \in (0, 1)$ , then we are done. Indeed, for each iteration of **ndpAlg** for which the sampled distance was near the middle, it is guaranteed that the algorithm either terminates or removes a constant fraction of the points. As such it is sufficient to present an algorithm which returns a value from this range with probability  $\geq 1 - 1/n^c$ , for some sufficiently large constant  $c$ , where  $n = |\mathbf{P}|$ . Such a value near the median is a *middle* nearest neighbor distance.

For the rest of this section we therefore focus on the problem of computing such an approximate middle nearest neighbor distance, in linear time, with high probability. To turn **ndpAlg** into an algorithm which runs in linear time with high probability, we replace **Line 1** and **Line 2** in **Figure 3.1<sub>p11</sub>** with this new subroutine.

**Sampling for the median** The problem with the random sampling from  $\mathbf{D}(\mathbf{P})$  done in **ndpAlg** is that though, with good probability, the sampled distance lies near the middle, it does not do so with high probability. This is the same problem one is faced when performing the standard basic randomized median selection. For median selection one known solution is rather than randomly sampling a pivot, to instead take a sufficiently large random sample, and take the median of the random sample as the pivot. This pivot is good with high probability [MR95], and it is this result we wish to mimic.

**The challenge** Given the above discussion, the natural solution to finding a good middle value from  $\mathbf{D}(\mathbf{P})$ , would be to sample multiple values from  $\mathbf{D}(\mathbf{P})$ , and then take the median. By the Chernoff inequality, if one samples at least a logarithmic number of values from a set of  $n$  values ( $\mathbf{D}(\mathbf{P})$  in our case) then, with high probability, the median of the sample will lie near the median of the entire set. (See **Lemma 5.3.1<sub>p30</sub>** for a formal statement and proof of this observation.) The problem with this approach is that sampling  $O(\log n)$  values from  $\mathbf{D}(\mathbf{P})$  takes  $\Theta(n \log n)$  time (at least naively). Specifically, while the sampling is easy, computing the associated values requires  $\Omega(n)$  time per sampled point (note this is not a problem for median selection as there we are sampling values directly rather than points from which values are determined).

Therefore, we instead not only take a sample from which to select our median value, but also a sample from which we will determine what these values are. Namely, for each point in a  $\Theta(\log n)$  sized random sample, compute its nearest neighbor in a second  $\Theta(n/\log n)$  sized sample, and then take the median of the computed values, and let this quantity be  $\ell$ . Though this approach takes  $\Theta(n)$  time, we need to somehow relate  $\ell$  to the desired middle nearest neighbor distance.

```

midNN(P):
   $\ell \leftarrow \text{estLogDist}(P)$ .

  res  $\leftarrow$  deciderM(P,  $\ell$ , 3/4)
  if res = " $d^{3/4}(P) \in [x, y]$ " then return x //  $y/x = O(1)$ 
1: if res = " $\ell < d^{3/4}(P)$ " then return  $\ell$  //Remark 5.3.3

  // Must be that res = " $\ell > d^{3/4}(P)$ ".
   $l \leftarrow \ell / 64 \lceil \ln^2 n \rceil$ 

  reslog  $\leftarrow$  deciderM(P,  $l$ , 1/2)
  if reslog = " $d^{1/2}(P) \in [x, y]$ " then return x //  $y/x = O(1)$ 
  if reslog = " $l < d^{1/2}(P)$ " then
2: return lowSpread(P,  $l/4$ ,  $4\ell$ ) //Section 5.2.3

  // Must be that reslog = " $l > d^{1/2}(P)$ "
3: return smallComp( $\ell$ , P) //Section 5.2.4

```

Figure 5.1: Algorithm for computing a constant factor approximation to a value in the interval  $[d^\delta(P), d^{1-\delta}(P)]$ , for some fixed constant  $\delta \in (0, 1)$ .

**Intuition** First note that for most points in the  $\Theta(\log n)$  sample, their nearest neighbor in  $P$  will not make it into the  $\Theta(n/\log n)$  sample from which we compute nearest neighbor distances. This implies  $\ell$  is an over estimate of middle nearest neighbor distances in  $P$ . The high level intuition is that how well  $\ell$  approximates middle distances tells us something about the distribution of  $P$ . There are two cases.

In the first case, most of the points of  $P$  are in small clusters<sup>¶</sup>. This case can be detected since here  $\ell$  will be way off, since for any point in the  $\Theta(\log n)$  sample, a  $\Theta(n/\log n)$  sample is not likely to have a point from the same small cluster. The advantage is that if points are in small clusters, then one only needs to look within a small cluster to compute a nearest neighbor distance (and this is done using connectivity clustering).

In the second case, rather than being clustered,  $P$  is roughly evenly distributed. Hence, intuitively, for any  $k$ ,  $d_k(p, P) \leq kd(p, P)$ . Note that for any point in the  $\Theta(\log n)$  sample, with high probability we are expected to see one of its  $\log^2 n$  nearest neighbors in a  $\Theta(n/\log n)$  sample, and so  $\ell \leq d_{\log^2 n}^{1/2}(P)$ . Combining these observations, roughly speaking  $\ell \leq (\log^2 n)d^{1/2}(P)$ . In other words, we know the median lies in a  $\log^2 n$  spread interval,  $[\ell/\log^2 n, \ell]$ . This case is then resolved using nets and grids (the details here are non-trivial).

## 5.2. The Algorithm

We now present the algorithm in full detail. Specifically, let **midNN** denote the new algorithm for computing a value in  $\mathcal{I} = [d^\delta(P), d^{1-\delta}(P)]$ , for some fixed constant  $\delta \in (0, 1)$ . In Section 4.3.1 it was shown how to convert a constant factor approximation to such a value into an exact value. Hence, computing a constant factor approximation to a value in  $\mathcal{I}$  in linear time, with high probability, suffices for our purposes.

<sup>¶</sup>Think of constant size clusters here, although in the algorithm they may be up to  $\log^2 n$  sized. Also, note the clusters are non-trivial as otherwise we are really in the second case.

**estLogDist**(P):  
 $n = |\mathbf{P}|$   
 $\mathbf{X} \leftarrow$  random sample of P of size  $\lceil c_5 \ln n \rceil$ .  
 $\mathbf{S} \leftarrow$  random sample of P of size  $\lceil c_6 n / \ln n \rceil$   
 Compute  $\mathbf{D}(\mathbf{X}, \mathbf{S})$   
**return** the element of rank  $|\mathbf{X}|/4$  in  $\mathbf{D}(\mathbf{X}, \mathbf{S})$ .

Figure 5.2: Roughly estimating  $d_{O(\log n)}^{1/4}(\mathbf{P})$ . Here  $c_5$  and  $c_6$  are sufficiently large constants.

The full algorithm is shown in Figure 5.1. Before analyzing the algorithm, we first define several subroutines that it employs.

### 5.2.1. **deciderM**: Counting distances

Given a point set P and a distance  $r$ , the decision procedure **deciderM**(P,  $r$ ,  $\alpha$ ) decides if the  $t$ th nearest neighbor distance in  $\mathbf{D}(\mathbf{P})$  is smaller than  $r$ , larger than  $r$ , or roughly equal to it, where  $t = \lfloor \alpha |\mathbf{P}| \rfloor$ . This decider is implicitly described in Theorem 4.3.1, but a simpler decider suffices, which we describe here in detail for the sake of clarity. To this end, throw the points of P into a grid with cell diameter  $r$ . Now, for every point  $\mathbf{p} \in \mathbf{P}$ , let  $w_r(\mathbf{p})$  be the number of the points in  $\mathbf{N}_{\leq r}(\mathbf{p})$ , not including  $\mathbf{p}$  itself. Clearly  $w_r(\mathbf{p})$  can be computed in constant time (given we have first recorded the number of points in each non-empty grid cell). Now, we compute the count

$$T = \sum_{\mathbf{p} \in \mathbf{P}} \frac{w_r(\mathbf{p})}{2},$$

in  $O(|\mathbf{P}|)$  time. Now, if  $T < t$  the procedure returns that  $r$  is smaller than the  $t$ th nearest-neighbor distance. Similarly, if  $T \geq t$ , then the  $t$ th nearest-neighbor distance must be at most  $3r$ . In this case, by also computing the count for the distance  $r/3$ , one will be able to either return a constant spread interval that contains the desired distance, or that this distance is strictly smaller than  $r$ . Thus yielding the required decider.

### 5.2.2. **estLogDist**: Sampling a distance

Given a set P of  $n$  points in  $\mathbb{R}^d$ , take two random samples from P, a sample X of size  $\Theta(\log n)$  and a sample S of size  $\Theta(n/\log n)$ , and compute  $\mathbf{d}^{1/4}(\mathbf{X}, \mathbf{S})$ . This can be done in  $O(n)$  time, by scanning all of S for each point in X to compute the set of distances  $\mathbf{D}(\mathbf{X}, \mathbf{S})$ . Next, compute value of rank  $|\mathbf{X}|/4$  in this set. We denote this algorithm by **estLogDist**(P), and it is shown in Figure 5.2.

### 5.2.3. **lowSpread**: Logarithmic Spread

The subroutine **lowSpread** used by **midNN** (i.e. Line 2 in Figure 5.1<sub>p28</sub>) handles the case when the desired middle nearest neighbor distance lies in a poly-logarithmic spread interval. The analysis of **lowSpread** can be found in Section 5.3.3. In the following, P is a set of unweighted points such that  $[\mathbf{d}^{1/2}(\mathbf{P}), \mathbf{d}^{3/4}(\mathbf{P})] \subseteq [r, R]$ , for some values  $0 < r \leq R$ , where  $n = |\mathbf{P}|$ . (Note, in the following, when sampling from a set S we treat it as a set of  $\omega(\mathbf{S})$  points.)

**Algorithm 5.2.1** (**lowSpread**(P,  $r$ ,  $R$ )).

- (A) Compute  $\mathbf{S} = \mathbf{net}(r/8, \mathbf{P})$ .
- (B) Throw the points of  $\mathbf{S}$  into a grid of sidelength  $R$ .
- (C)  $\mathbf{U} \leftarrow$  random sample from  $\mathbf{S}$  of  $\Theta(\log(n))$  points.
- (D) Let  $\mathbf{Z}$  be an approximation to the set  $\mathbf{D}(\mathbf{U}, \mathbf{S})$ .  
Specifically, for each point  $\mathbf{p} \in \mathbf{U}$  look within  $\mathbf{N}_{\leq 2R}(\mathbf{p})$  for its nearest neighbor, see [Definition 2.1.2<sub>p7</sub>](#). If such a neighbor is found record this distance (which might be zero if  $\mathbf{p}$  is of weight  $> 1$ ) and if no neighbor is found set this value to  $\infty$ .
- (E) Let  $\mathbf{p} \in \mathbf{P}$  be the point that corresponds to the number of rank  $(5/8)|\mathbf{Z}|$  in  $\mathbf{Z}$ , and let  $\tau$  be this number.
- (F) Return  $\tau$  ( $\tau$  is a constant approximation to  $\mathbf{d}(\mathbf{p}, \mathbf{P})$ ).

#### 5.2.4. **smallComp**: Many small components

The subroutine **smallComp**, used by **midNN** (i.e. [Line 3](#) in [Figure 5.1](#)), handles the case when the desired middle nearest neighbor distance is considerably smaller than the value sampled by **estLogDist**. This corresponds, intuitively, to the situation where the input is mainly clustered in tiny clusters. The analysis of **smallComp** is delegated to [Section 5.3.4](#).

**Algorithm 5.2.2** (**smallComp**( $\ell, \mathbf{P}$ )). *The algorithm works as follows:*

- (A) Compute a partition  $\mathcal{C}$ , such that  $\mathcal{C}_{\leq \rho}(\mathbf{P}) \sqsubseteq \mathcal{C} \sqsubseteq \mathcal{C}_{\leq 2\rho}(\mathbf{P})$  in linear time, using the algorithm of [Lemma 4.4.2<sub>p20</sub>](#), where  $\rho = \ell/8m$ , where  $m = \lceil \ln^2 n \rceil$  and  $n = |\mathbf{P}|$ .
- (B) Identify the components of  $\mathcal{C}$  which are neither singletons nor contain more than  $c_3 \ln^2 n$  points, where  $c_3$  is a sufficiently large constant.
- (C) Collect all the points in these components into a set of points  $\mathbf{X}$  (this set is not empty).
- (D) Take a random sample  $\mathbf{U}$  of  $c_4 \log n$  points from  $\mathbf{X}$ , where  $c_4$  is a sufficiently large constant.
- (E) For each point in  $\mathbf{U}$ , compute its nearest neighbor in  $\mathbf{P}$  by scanning the non-trivial connected component in  $\mathcal{C}$  that contains it. Let  $\mathbf{Z}$  be this set of numbers.
- (F) Return the median value in  $\mathbf{Z}$ .

### 5.3. Algorithm Analysis

#### 5.3.1. Why median of a sample is a good estimate for the median

The following lemma is by now standard, and we include its proof for the sake of completeness.

**Lemma 5.3.1.** *Let  $\mathbf{X}$  be a set of  $n$  real numbers, constants  $\delta \in (0, 1/2)$ ,  $\alpha \in (0, 1)$ , and a parameter  $t > 0$ . Let  $\mathbf{U}$  be a random sample of size  $t$  picked (with repetition and uniformly) from  $\mathbf{X}$ . Then, the element of rank  $\alpha t$  in  $\mathbf{U}$  has rank in the interval  $\left[ (1 - \delta)\alpha n, (1 + \delta)\alpha n \right]$  in the original set  $\mathbf{X}$ , with probability  $\geq 1 - 2 \exp(-\delta^2 \alpha t / 8)$ . In particular, for  $t \geq (8c \ln n) / (\delta^2 \alpha)$ , this holds with probability  $\geq 1 - 2/n^c$ .*

*Proof:* Let  $Y_L$  be the number of values in  $\mathbf{U}$  of rank  $\leq (1 - \delta)\alpha n$  in  $\mathbf{X}$ . Similarly, let  $Y_R$  be the number of values in  $\mathbf{U}$  of rank  $\leq (1 + \delta)\alpha n$  in  $\mathbf{X}$ . Clearly  $\mu_L = \mathbf{E}[Y_L] = (1 - \delta)\alpha t$  and  $\mu_R = \mathbf{E}[Y_R] = (1 + \delta)\alpha t$ .

Observe that if  $Y_L \leq \alpha t \leq Y_R$  then the element of rank  $\alpha t$  in  $\mathbf{U}$  has rank in  $\mathbf{X}$  in the interval

$\left[(1 - \delta)\alpha n, (1 + \delta)\alpha n\right]$ , as desired. By the Chernoff inequality [MR95], we have that

$$\begin{aligned} \Pr\left[Y_L > \alpha t\right] &= \Pr\left[Y_L > \frac{\alpha t}{\mu_L}\mu_L\right] = \Pr\left[Y_L > \frac{\alpha t}{(1 - \delta)\alpha t}\mu_L\right] \leq \Pr\left[Y_L > (1 + \delta)\mu_L\right] \\ &\leq \exp\left(-\frac{\delta^2}{4}\mu_L\right) \leq \exp\left(-\frac{\delta^2}{8}\alpha t\right), \end{aligned}$$

as  $1/(1 - \delta) > 1 + \delta$ . Arguing as above, the other bad event probability is

$$\begin{aligned} \Pr\left[Y_R < \alpha t\right] &\leq \Pr\left[Y_R < \frac{\alpha t}{\mu_R}\mu_R\right] = \Pr\left[Y_R < \frac{\alpha t}{(1 + \delta)\alpha t}\mu_R\right] \\ &\leq \Pr\left[Y_R < \left(1 - \frac{\delta}{2}\right)\mu_R\right] \leq \exp\left(-\frac{\delta^2}{8}\mu_R\right) \leq \exp\left(-\frac{\delta^2}{8}\alpha t\right). \end{aligned}$$

■


### 5.3.2. **estLogDist**: Analysis

**Lemma 5.3.2.** *Let  $\mathbf{P}$  be a set of  $n$  points in  $\mathbb{R}^d$ . Let  $\ell$  be the distance returned by executing **estLogDist**( $\mathbf{P}$ ), see [Figure 5.2<sub>p29</sub>](#). With high probability, we have the following:*

- (A) For all points  $\mathbf{p} \in \mathbf{P}$ , the set  $|\mathbf{P} \cap \mathbf{b}(\mathbf{p}, r_{\mathbf{p}})|$  has at most  $m = \lceil \ln^2 n \rceil$  points, where (i)  $r_{\mathbf{p}} = \mathbf{d}(\mathbf{p}, \mathbf{S})$ , (ii)  $\mathbf{S}$  is the random sample used by **estLogDist**, and (iii)  $\mathbf{b}(\mathbf{p}, r_{\mathbf{p}})$  is the ball of radius  $r_{\mathbf{p}}$  centered at  $\mathbf{p}$ .
- (B) Let  $\mathcal{P}$  be a clustering of  $\mathbf{P}$  such that  $\mathcal{C}_{\leq \rho}(\mathbf{P}) \subseteq \mathcal{P} \subseteq \mathcal{C}_{\leq 2\rho}(\mathbf{P})$ , as computed by the algorithm of [Lemma 4.4.2<sub>p20</sub>](#), for  $\rho = \ell/8m$ . Then, at least  $(5/8)n$  of the points of  $\mathbf{P}$  are in clusters of  $\mathcal{P}$  that contain at most  $m$  points.
- (C)  $\ell \geq \mathbf{d}^{1/8}(\mathbf{P})$ .

*Proof:* (A) Fix a point  $\mathbf{p} \in \mathbf{P}$ , and consider a ball  $\mathbf{b}$  centered at  $\mathbf{p}$  whose interior contains exactly  $\lceil \ln^2 n \rceil$  points of  $\mathbf{P}$ . The probability that a random point from  $\mathbf{P}$  hits  $\mathbf{b}$  is  $\alpha = \ln^2 n/n$ . As such, the probability that all the points of  $\mathbf{S}$  miss it is at most  $(1 - \alpha)^{c_5 n / \ln n} \leq e^{-c_5 \ln n} \leq 1/n^{c_5}$ , implying the claim, if  $c_5$  is sufficiently large.

(B) By [Lemma 5.3.1](#), for  $c_6$  sufficiently large, we have that  $\ell$  is in the range  $\left[\mathbf{d}^{1/8}(\mathbf{P}, \mathbf{S}), \mathbf{d}^{3/8}(\mathbf{P}, \mathbf{S})\right]$ , and this holds with high probability. That is, for at least  $(5/8)n$  points of  $\mathbf{P}$  have  $\mathbf{d}(\mathbf{p}, \mathbf{S}) \geq \ell$ , and let  $\mathbf{Q}$  be this set of points. For every point  $\mathbf{p} \in \mathbf{Q}$ , we have that the ball  $\mathbf{b} = \mathbf{b}(\mathbf{p}, \ell) \subseteq \mathbf{b}(\mathbf{p}, r_{\mathbf{p}})$  contains at most  $m$  points of  $\mathbf{P}$ , by (A). Partitioning  $\mathbf{b}$  into equal radius rings of width  $\ell/(m + 2)$ , one of them must be empty of points of  $\mathbf{P}$ . In particular, no cluster of  $\mathcal{C}_{\leq 2\rho}(\mathbf{P})$  can contain points on both sides of this ring, see figure on the right. As such, at least  $(5/8)n$  of the points of  $\mathbf{P}$  are in clusters of  $\mathcal{P}$  that contain at most  $m$  points.



(C) Follows immediately from (B), as  $\ell \geq \mathbf{d}^{1/8}(\mathbf{P}, \mathbf{S}) \geq \mathbf{d}^{1/8}(\mathbf{P})$ . ■

**Remark 5.3.3.** *Line 1 of **midNN**( $\mathbf{P}$ ) returns  $\ell$  as the approximate middle nearest neighbor distance if **deciderM** returned  $\ell < \mathbf{d}^{3/4}(\mathbf{P})$ . [Lemma 5.3.2](#) (C) then implies that  $\mathbf{d}^{1/8}(\mathbf{P}) \leq \ell < \mathbf{d}^{3/4}(\mathbf{P})$ , with high probability, and so  $\ell$  is a valid approximation. In the unlikely event that  $\mathbf{d}^{1/8}(\mathbf{P}) > \ell$ , one simply runs **midNN**( $\mathbf{P}$ ) again. This event can be determined with the call **deciderM**( $\mathbf{P}, \ell, 1/8$ ).*

### 5.3.3. lowSpread: Analysis

During an execution of **midNN**( $P$ ), **lowSpread**( $P, r, R$ ) gets called (see [Algorithm 5.2.1](#)) only if the interval  $[r, R]$  contains the values  $d^{1/2}(P), \dots, d^{3/4}(P)$ .

**Lemma 5.3.4.** *Let  $P$  be a point set in  $\mathbb{R}^d$  of size  $n$ . Given  $r, R \in \mathbb{R}$ , such that (i)  $R/r = O(\log^2 n)$ , and (ii)  $d^{1/2}(P), \dots, d^{3/4}(P) \in [r, R]$ , then **lowSpread**( $P, r, R$ ) runs in  $O(n)$  time, and returns a value which is a  $O(1)$ -approximation to some nearest neighbor distance  $d(p, P) \in [d^{1/2}(P), d^{3/4}(P)]$ , with high probability, for some point  $p \in P$ .*

*Proof:* Set  $\Delta = r/8$ . Using the notation from **lowSpread**, see [Algorithm 5.2.1<sub>p29</sub>](#), let  $S = \mathbf{net}(\Delta, P)$  and  $U$  the  $\Theta(\log(n))$  size sample from  $S$ . By [Lemma 5.3.1](#), with high probability  $d^{5/8}(U, S) \in [d^{1/2}(S), d^{3/4}(S)]$ , and assume that this is the case. Since  $S$  is an  $\Delta$ -net of  $P$ , by [Claim 4.2.1](#), this changes distances by at most  $2\Delta$ , which implies that

$$\tau = d^{5/8}(U, S) \in [d^{1/2}(P) - 2\Delta, d^{3/4}(P) + 2\Delta] \subseteq \left[\frac{1}{2}d^{1/2}(P), 2d^{3/4}(P)\right] \subseteq \left[\frac{r}{2}, 2R\right].$$

by assumption.

In particular, this implies that all the relevant distances from the computation of this middle value are in the range  $[r/2, 2R]$ . As such, when computing the  $(5/8)$ -middle element, it is allowable to not compute precisely values outside this interval, as done by the algorithm. This implies that **lowSpread** indeed returns the value  $\tau$ , which is the desired approximation.

As for the running time, observe that  $S$  is an  $\Delta$ -net, and an open ball with this radius around any point must be empty of any other net point. In particular, for any point  $p \in S$ , the maximum number of other points that can be in  $N_{\leq 2R}(p)$  is  $O((R/r)^d) = O(\log^{2d} n)$ . Therefore, for a point  $p \in U$ , computing  $d(p, S)$  takes polylogarithmic time (using linear time preprocessing and a grid). Thus, overall, computing the set of distances  $D(U, S)$  takes polylogarithmic time, since  $|U| = \Theta(\log n)$ . Hence, the overall running time is  $O(n + \text{polylog } n) = O(n)$ .  $\blacksquare$

### 5.3.4. smallComp: Small Components

**Lemma 5.3.5.** *The call to **smallComp**( $\ell, P$ ) by **midNN** computes a point  $p \in P$ , such that  $d(p, P) \in [d^{1/32}(P), d^{31/32}(P)]$ . The procedure succeeds and its running time is  $O(n)$ , with high probability, where  $n = |P|$ .*

*Proof:* Computing the components of  $\mathcal{C}$  takes linear time, and given  $\mathcal{C}$ , computing  $X$  takes linear time. For each point in the sample  $U$ , it takes  $O(\log^2 n)$  time to compute its nearest neighbor since one only has to scan the point's  $O(\log^2 n)$  sized connected component in  $\mathcal{C}$ . Since  $|U| = O(\log n)$ , computing the nearest neighbor distances and taking the median takes polylogarithmic time overall.

Now, **smallComp**( $\ell, P$ ) is called by **midNN** only if the following conditions holds:

- (C1)  $\ell/64 \lceil \ln^2 n \rceil > d^{1/2}(P)$ . Namely, at least half of the points of  $P$  have their nearest neighbor in their own cluster in  $\mathcal{C}_{\leq \rho}(P)$ , where  $\rho = \ell / (8 \lceil \ln^2 n \rceil)$ . Let  $Z_1$  be the set of these points.
- (C2) Let  $\mathcal{P}$  be the partition of  $P$  computed by **smallComp**. We have that  $\mathcal{C}_{\leq \rho}(P) \sqsubseteq \mathcal{P} \sqsubseteq \mathcal{C}_{\leq 2\rho}(P)$ . By [Lemma 5.3.2](#) (B), at least  $(5/8)n$  of the points of  $P$  are in clusters of  $\mathcal{P}$  that contain at most  $m$  points (note, that some of these points might be a cluster containing only a single point). Let  $Z_2$  be the set of these points.



As such,  $|Z_1 \cap Z_2| \geq n/8$ . Namely, at least  $n/8$  of the points of  $P$ , are in small clusters and these clusters have strictly more than one point. Furthermore,  $Z_1 \cap Z_2 \subseteq X$ , and we conclude that  $|X| \geq n/8$ . Now, by [Lemma 5.3.1](#), with high probability, the returned value is in the range  $d^{1/4}(X), \dots, d^{3/4}(X)$ , which implies that the returned value is in the range  $d^{1/32}(P), \dots, d^{31/32}(P)$ , as desired. ■

## 5.4. The result

The analysis in the previous section implies the following.

**Lemma 5.4.1.** *Given a set  $P \subseteq \mathbb{R}^d$  of  $n$  points, **midNN**( $P$ ) computes, a value  $x$ , such that there is a value of  $d^{1/32}(P), \dots, d^{31/32}(P)$  in the interval  $[x/c, cx]$ , where  $c > 1$  is a constant. The running time bound holds with high probability.*

As shown in [Section 4.3.1](#), given a constant factor approximation to some nearest neighbor distance, we can compute the corresponding nearest neighbor distance exactly in linear time using a grid computation. Let **midNNExact** denote the algorithm that performs this additional linear time step after running **midNN**. As argued above, replacing [Line 1](#) and [Line 2](#) in [Figure 3.1<sub>p11</sub>](#) with “ $\ell_i = \mathbf{midNNExact}(P_{i-1})$ ”, turns **ndpAlg** into an algorithm that runs in linear time with high probability.

We thus have the following analogue of [Theorem 3.3.2](#).

**Theorem 5.4.2.** *For a constant  $\varphi > 3/2$ , given an instance of a  $\varphi$ -NDP defined by a set of  $n$  points in  $\mathbb{R}^d$ , with a  $\varphi$ -decider with (deterministic or high-probability) running time  $O(n)$ , one can get a  $\varphi$ -approximation to its optimal value, in  $O(n)$  time, with high probability.*

*Similarly, for  $\varepsilon \in (0, 1/2)$ , given an instance of a  $(1 + \varepsilon)$ -NDP defined by a set of  $n$  points in  $\mathbb{R}^d$ , with a  $(1 + \varepsilon)$ -decider having (deterministic, or high probability) running time  $O(n/\varepsilon^c)$ , then one can  $(1 + \varepsilon)$ -approximate the optimal value for this NDP, in  $O(n/\varepsilon^c)$  time, with high probability.*

**Corollary 5.4.3.** *All the applications of [Section 4](#) can be modified so that their running time bound holds not only in expectation, but also with high probability.*

**Remark 5.4.4.** Note, that the above high probability guarantee is in the size of the input. As such, in later iterations of the algorithm, when the input size of the subproblem is no longer polynomial in  $n$  (say, the subproblem size is  $O(\text{polylog } n)$ ), the high probability guarantee is no longer meaningful.

Fortunately, this is a minor technicality – once the subproblem size drops below  $n/\log n$ , the remaining running time of **ndpAlg** is  $O(n)$ , with high probability. This follows by a simple direct argument – by the Chernoff inequality, after  $\Theta(\log n)$  iterations, with high probability at least  $\log_{16/15} n$  iterations are successful, see the proof of [Lemma 3.2.4](#) for details. After this number of successful iterations, the algorithm terminates. This implies that the remaining running time is  $O(n)$ , with high probability, once the subproblem size is  $O(n/\log n)$ .

As such, our new high-probability algorithm is needed only in the initial “heavy” iterations of the **ndpAlg** algorithm, until the subproblem size drops below  $O(n/\log n)$ .

## 6. Conclusions

There is further research to be done in investigating the technique presented here. For example, looking into the implementation of this new algorithm in both standard settings and distributed settings (e.g.,

MapReduce). Additionally, since now one can do approximate distance selection in linear time, maybe now one can get a speed up for other algorithms that do (not necessarily point based) distance selection.

Our framework provides a new way of looking at distance based optimization problems, in particular through the lens of nets. We know how to compute nets efficiently for doubling metrics and it seems one can compute approximate nets in near linear time for planar graphs. For example, it seems the new technique implies that approximate  $k$ -center clustering in planar graphs can be done in near linear time. This provides fertile ground for future research.

**Open question: When is linear time not possible** Another interesting open problem arising out of the high probability linear time algorithm is the following: Given sets  $\mathbf{P}$  and  $\mathbf{X}$  in  $\mathbb{R}^d$ , of size  $n$  and  $O(\log n)$ , respectively, how long does it take to approximate the values in the set  $\mathcal{D}(\mathbf{X}, \mathbf{P})$  (i.e., for each point of  $\mathbf{X}$  we approximate its nearest-neighbor in  $\mathbf{P}$ ). This problem can be easily solved in  $O(n \log n)$  time. The open question is whether it is possible solve this problem in linear time.

The algorithm in this paper implies that we can compute exactly a value (of a specified rank) in this set in  $O(n)$  time (even if  $|\mathbf{X}|$  is of size  $n$ ). We currently have an algorithm with running time  $O(n \log |\mathbf{X}|) = O(n \log \log n)$  for approximating  $\mathcal{D}(\mathbf{X}, \mathbf{P})$ .

More generally, this is a proxy to the meta-question of when linear time algorithms do not exist for distance problems induced by points.

## Acknowledgments

The authors thank Sergio Cabello for useful discussions. In particular, the deterministic algorithm of [Section 3.3.1](#) was suggested by him. The authors also thank the anonymous referees for their insightful and detailed comments.

## References

- [AH08] B. Aronov and S. Har-Peled. On approximating the depth and related problems. *SIAM J. Comput.*, 38(3):899–921, 2008.
- [AHV04] P. K. Agarwal, S. Har-Peled, and K. R. Varadarajan. Approximating extent measures of points. *J. Assoc. Comput. Mach.*, 51(4):606–635, 2004.
- [AHV05] P. K. Agarwal, S. Har-Peled, and K. Varadarajan. Geometric approximation via coresets. In J. E. Goodman, J. Pach, and E. Welzl, editors, *Combinatorial and Computational Geometry*, Math. Sci. Research Inst. Pub. Cambridge, New York, NY, USA, 2005.
- [APF<sup>+</sup>10] G. Aggarwal, R. Panigrahy, T. Feder, D. Thomas, K. Kenthapadi, S. Khuller, and A. Zhu. Achieving anonymity via clustering. *ACM Trans. Algo.*, 6(3):49:1–49:19, 2010.
- [AST94] P. K. Agarwal, M. Sharir, and S. Toledo. Applications of parametric searching in geometric optimization. *J. Algorithms*, 17:292–318, 1994.
- [BS02] S. Bespamyatnikh and M. Segal. Fast algorithms for approximating distances. *Algorithmica*, 33(2):263–269, 2002.
- [Cha08] T. M. Chan. Well-separated pair decomposition in linear time? *Inform. Process. Lett.*, 107(5):138–141, 2008.

- [CK95] P. B. Callahan and S. R. Kosaraju. A decomposition of multidimensional point sets with applications to  $k$ -nearest-neighbors and  $n$ -body potential fields. *J. Assoc. Comput. Mach.*, 42:67–90, 1995.
- [Cla83] K. L. Clarkson. Fast algorithms for the all nearest neighbors problem. In *Proc. 24th Annu. IEEE Sympos. Found. Comput. Sci. (FOCS)*, pages 226–232, Washington, DC, USA, 1983. IEEE Computer Society.
- [Cla88] K. L. Clarkson. Applications of random sampling in computational geometry, II. In *Proc. 4th Annu. Sympos. Comput. Geom. (SoCG)*, pages 1–11, New York, NY, USA, 1988. ACM.
- [CS89] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.
- [DHW12] A. Driemel, S. Har-Peled, and C. Wenk. Approximating the Fréchet distance for realistic curves in near linear time. *Discrete Comput. Geom.*, 48:94–127, 2012.
- [EET93] H. ElGindy, H. Everett, and G. Toussaint. Slicing an ear using prune-and-search. *Pattern Recogn. Lett.*, 14(9):719–722, 1993.
- [ERH12] A. Ene, B. Raichel, and S. Har-Peled. Fast clustering with lower bounds: No customer too far, no shop too small. In submission. <http://sarielhp.org/papers/12/lbc/>, 2012.
- [Eri95] J. Erickson. On the relative complexities of some geometric problems. In *Proc. 7th Canad. Conf. Comput. Geom. (CCCG)*, pages 85–90, Ottawa, Canada, 1995. Carleton University.
- [FJ84] G. N. Frederickson and D. B. Johnson. Generalized selection and ranking: Sorted matrices. *SIAM J. Comput.*, 13:14–30, 1984.
- [Gon85] T. Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoret. Comput. Sci.*, 38:293–306, 1985.
- [GRSS95] M. Golin, R. Raman, C. Schwarz, and M. Smid. Simple randomized algorithms for closest pair problems. *Nordic J. Comput.*, 2:3–27, 1995.
- [Har01] S. Har-Peled. Clustering motion. In *Proc. 42nd Annu. IEEE Sympos. Found. Comput. Sci. (FOCS)*, pages 84–93, Washington, DC, USA, 2001. IEEE Computer Society.
- [Har04a] S. Har-Peled. Clustering motion. *Discrete Comput. Geom.*, 31(4):545–565, 2004.
- [Har04b] S. Har-Peled. No coresets, no cry. In *Proc. 24th Conf. Found. Soft. Tech. Theoret. Comput. Sci. (FSTTCS)*, pages 324–335, Berlin, Heidelberg, 2004. Springer-Verlag.
- [Har11] S. Har-Peled. *Geometric Approximation Algorithms*, volume 173 of *Mathematical Surveys and Monographs*. Amer. Math. Soc., Boston, MA, USA, 2011.
- [HK07] S. Har-Peled and A. Kushal. Smaller coresets for  $k$ -median and  $k$ -means clustering. *Discrete Comput. Geom.*, 37(1):3–19, 2007.
- [HM04] S. Har-Peled and S. Mazumdar. Coresets for  $k$ -means and  $k$ -median clustering and their applications. In *Proc. 36th Annu. ACM Sympos. Theory Comput. (STOC)*, pages 291–300, New York, NY, USA, 2004. ACM.

- [HM05] S. Har-Peled and S. Mazumdar. Fast algorithms for computing the smallest  $k$ -enclosing disc. *Algorithmica*, 41(3):147–157, 2005.
- [HM06] S. Har-Peled and M. Mendel. Fast construction of nets in low dimensional metrics, and their applications. *SIAM J. Comput.*, 35(5):1148–1184, 2006.
- [HR11] S. Har-Peled and B. Raichel. The Fréchet distance revisited and extended. In *Proc. 27th Annu. Sympos. Comput. Geom. (SoCG)*, pages 448–457, New York, NY, USA, 2011. ACM. <http://sarielhp.org/papers/10/frechet3d/>.
- [HR13] S. Har-Peled and B. Raichel. Net and prune: A linear time algorithm for euclidean distance problems. In *Proc. 45th Annu. ACM Sympos. Theory Comput. (STOC)*, pages 605–614, New York, NY, USA, 2013. ACM.
- [HR14] S. Har-Peled and B. Raichel. Net and prune: A linear time algorithm for euclidean distance problems. *CoRR*, abs/1409.7425, 2014.
- [KL04] R. Krauthgamer and J. R. Lee. Navigating nets: simple algorithms for proximity search. In *Proc. 15th ACM-SIAM Sympos. Discrete Algs. (SODA)*, pages 798–807, Philadelphia, PA, USA, 2004. Society for Industrial and Applied Mathematics.
- [KM95] S. Khuller and Y. Matias. A simple randomized sieve algorithm for the closest-pair problem. *Inform. Comput.*, 118:34–37, 1995.
- [LMS94] C.-Y. Lo, J. Matoušek, and W. L. Steiger. Algorithms for ham-sandwich cuts. *Discrete Comput. Geom.*, 11:433–452, 1994.
- [Meg83] N. Megiddo. Applying parallel computation algorithms in the design of serial algorithms. *J. Assoc. Comput. Mach.*, 30(4):852–865, 1983.
- [Meg84] N. Megiddo. Linear programming in linear time when the dimension is fixed. *J. Assoc. Comput. Mach.*, 31:114–127, 1984.
- [MR95] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, Cambridge, UK, 1995.
- [MSW96] J. Matoušek, M. Sharir, and E. Welzl. A subexponential bound for linear programming. *Algorithmica*, 16:498–516, 1996.
- [Mul94] K. Mulmuley. *Computational Geometry: An Introduction Through Randomized Algorithms*. Prentice Hall, Englewood Cliffs, NJ, 1994.
- [Rab76] M. O. Rabin. Probabilistic algorithms. In J. F. Traub, editor, *Algorithms and Complexity: New Directions and Recent Results*, pages 21–39. Academic Press, Orlando, FL, USA, 1976.
- [Sal97] J. Salowe. Parametric search. In J. E. Goodman and J. O’Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 37, pages 683–698. CRC Press LLC, Boca Raton, FL, 1997.
- [Sch79] A. Schönhage. On the power of random access machines. In *Proc. 6th Internat. Colloq. Automata Lang. Prog.*, volume 71 of *Lecture Notes Comput. Sci.*, pages 520–529, London, UK, 1979. Springer-Verlag.

- [Smi00] M. Smid. Closest-point problems in computational geometry. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 877–935. Elsevier, Amsterdam, The Netherlands, 2000.
- [SW92] M. Sharir and E. Welzl. A combinatorial bound for linear programming and related problems. In *Proc. 9th Sympos. Theoret. Aspects Comput. Sci.*, volume 577 of *Lect. Notes in Comp. Sci.*, pages 569–579, London, UK, 1992. Springer-Verlag.
- [vOV04] R. van Oostrum and R. C. Veltkamp. Parametric search made practical. *Comput. Geom. Theory Appl.*, 28(2-3):75–88, 2004.