

On Locality-Sensitive Orderings and Their Applications

Timothy M. Chan

Department of Computer Science, University of Illinois at Urbana-Champaign, USA
tmc@illinois.edu

Sariel Har-Peled¹

Department of Computer Science, University of Illinois at Urbana-Champaign, USA
sariel@illinois.edu

Mitchell Jones

Department of Computer Science, University of Illinois at Urbana-Champaign, USA
mfjones2@illinois.edu

Abstract

For any constant d and parameter $\varepsilon > 0$, we show the existence of (roughly) $1/\varepsilon^d$ orderings on the unit cube $[0, 1]^d$, such that any two points $p, q \in [0, 1]^d$ that are close together under the Euclidean metric are “close together” in one of these linear orderings in the following sense: the only points that could lie between p and q in the ordering are points with Euclidean distance at most $\varepsilon\|p - q\|$ from p or q . These orderings are extensions of the \mathcal{Z} -order, and they can be efficiently computed.

Functionally, the orderings can be thought of as a replacement to quadtrees and related structures (like well-separated pair decompositions). We use such orderings to obtain surprisingly simple algorithms for a number of basic problems in low-dimensional computational geometry, including (i) dynamic approximate bichromatic closest pair, (ii) dynamic spanners, (iii) dynamic approximate minimum spanning trees, (iv) static and dynamic fault-tolerant spanners, and (v) approximate nearest neighbor search.

2012 ACM Subject Classification Theory of computation \rightarrow Computational geometry

Keywords and phrases Approximation algorithms, Data structures, Computational geometry

Digital Object Identifier 10.4230/LIPIcs.ITCS.2019.21

Acknowledgements The authors thank the anonymous reviewers for their helpful suggestions.

1 Introduction


In this paper, we describe a technique that leads to new simple algorithms for a number of fundamental proximity problems in low-dimensional Euclidean spaces.

Notation

The O notation hides constants that depends (usually exponentially) on d . Throughout, we assume (without loss of generality) that ε is a power of 2; that is $\varepsilon = 2^{-E}$ for some positive integer E .

¹ Supported in part by NSF AF award CCF-1421231.



Z-order

Consider a point set $P \subseteq [0, 1)^2$, its quadtree, and a DFS traversal of this quadtree. One can order the points of P according to this traversal, resulting in an ordering \prec of the underlying set $[0, 1)^2$. The relation \prec is the ordering along a space filling mapping, known as the **Z-order**. Specifically, there is a bijection z from the unit interval $[0, 1)$ to the unit square $[0, 1)^2$ such that the ordering along the resulting “curve” is the Z-order. The Z-order mapping z is not continuous. Nevertheless, the Z-order mapping has the advantage of being easy to define. Indeed, given a real number $\alpha \in [0, 1)$, with the binary expansion $\alpha = 0.x_1x_2x_3\dots$ (i.e., $\alpha = \sum_{i=1}^{\infty} x_i2^{-i}$), the Z-order mapping of α is the point $z(\alpha) = (0.x_2x_4x_6\dots, 0.x_1x_3x_5\dots)$. Computing the Z-order or its inverse is quite easy, if one is allowed bitwise-logical operations – in particular, the ability to compute compressed quadtrees efficiently is possible only if such operations are available [19]. The approach extends to higher constant dimensions. The idea of using the Z-order can be traced back to the work of Gargantini [16], and it is widely used in databases and seems to improve performance in practice [24]. Once comparison by Z-order is available, building a compressed quadtree is no more than storing the points according to the Z-order, and this yields simple data structures for various problems. For example, Laio *et al.* [27] and Chan [7, 8, 9] applied the Z-order to obtain simple efficient algorithms for approximate nearest neighbor search and related problems.

Shifting

The Z-order (and quadtrees) does not preserve distance. That is, two points that are far away might be mapped to two close-together points, and vice versa. This problem is even apparent when using a grid, where points that are close together get separated into different grid cells. One way to get around this problem is to shift the grid (deterministically or randomly) [21]. The same approach works for quadtrees – one can shift the quadtree constructed for a point set several times such that for any pair of points in the quadtree, there will be a shift where the two points are in a cell of diameter that is $O(1)$ times their distance. Improving an earlier work by Bern [3], Chan [6] showed that $2\lceil d/2 \rceil + 1$ deterministic shifts are enough in d dimensions (a proof is reproduced in Appendix A.2). A somewhat similar shifting scheme was also suggested by Feige and Krauthgamer [14]. Random shifting of quadtrees underlines, for example, the approximation algorithm by Arora for Euclidean TSP [2].

By combining Z-order with shifting, both Chan [7] and Laio *et al.* [27] observed an extremely simple data structure for $O(1)$ -approximate nearest neighbor search in constant dimensions: just store the points in Z-order for each of the $2\lceil d/2 \rceil + 1$ shifts; given a query point q , find the successor and predecessor of q in the Z-order by binary search for each of the shifts, and return the closest point found. The data structure can be easily made dynamic to support insertions and deletions of points, and can also be adapted to find $O(1)$ -approximate bichromatic closest pairs.

For approximate nearest neighbor (ANN) search, the $O(1)$ approximation factor can be reduced to $1 + \varepsilon$ for any fixed $\varepsilon > 0$, though the query algorithm becomes more involved [7] and unfortunately cannot be adapted to compute $(1 + \varepsilon)$ -approximate bichromatic closest pairs. (In the monochromatic case, however, the approach can be adapted to find *exact* closest pairs, by considering $O(1)$ successors and predecessors of each point [7].)

For other proximity-related problems such as spanners and approximate minimum spanning trees (MST), this approach does not seem to work as well: for example, the static algorithms in [9], which use the \mathcal{Z} -order, still require explicit constructions of compressed quadtrees and are not easily dynamizable.

Main new technique: Locality Sensitive Orderings

For any given $\varepsilon > 0$, we show that there is a family of $O((1/\varepsilon^d) \log(1/\varepsilon))$ orderings of $[0, 1]^d$ with the following property: For any $p, q \in [0, 1]^d$, there is an ordering in the family such that all points lying between p and q in this ordering are within distance at most $\varepsilon \|p - q\|$ from p or q . The order between two points can be determined efficiently using some bitwise-logical operations. See Theorem 10. We refer to these as **locality-sensitive orderings**. They generalize the previous construction of $2\lceil d/2 \rceil + 1$ shifted copies of the \mathcal{Z} -order, which guarantees the stated property only for a large specific constant ε (dependent on d). The property ensures, for example, that a $(1 + \varepsilon)$ -approximate nearest neighbor of a point q can be found among the immediate predecessors and successors of q in these orderings.

Applications

Locality-sensitive orderings immediately lead to simple algorithms for a number of problems, as listed below. Many of these results are significant simplification of previous work; some of the results are new.

- (a) **Approximate bichromatic closest pair.** Theorem 12 presents a data structure that maintains a $(1 + \varepsilon)$ -approximate closest bichromatic pair for two sets of points in \mathbb{R}^d , with an update time of $O(\log n)$, for any fixed $\varepsilon > 0$, ignoring factors that depend on ε (roughly $1/\varepsilon^d$). Previously, a general technique of Eppstein [12] can be applied in conjunction with a dynamic data structure for ANN, but the amortized update time increases by two $\log n$ factors.
- (b) **Dynamic spanners.** For a parameter $t \geq 1$ and a set of points P in \mathbb{R}^d , a graph $G = (P, E)$ is a t -spanner for P if for all $p, q \in P$, there is a p - q path in G of length at most $t \|p - q\|$. Static algorithms for spanners have been extensively studied in computational geometry. The dynamic problem appears tougher, and has also received much attention (see Table 1.1). We obtain a very simple data structure for maintaining a dynamic $(1 + \varepsilon)$ -spanners in Euclidean space with an update (insertion and deletion) time of $O(\log n)$ and having $O(n)$ edges in total, for any fixed $\varepsilon > 0$, ignoring factors that depend on ε (roughly $1/\varepsilon^d$). See Theorem 14. Although Gottlieb and Roditty [17] have previously obtained the same update time $(1/\varepsilon^{O(d)}) \log n$, their method requires much more intricate details. (Note that Gottlieb and Roditty's method more generally applies to spaces with bounded doubling dimension, but no simpler methods have been reported in the Euclidean setting.)

21:4 On Locality-Sensitive Orderings and Their Applications

■ **Table 1.1** Previous work on dynamic $(1 + \varepsilon)$ -spanners in \mathbb{R}^d . Dependencies on ε (of the form $1/\varepsilon^{O(d)}$) are omitted in the O bounds.

reference	insertion time	deletion time
Roditty [30]	$O(\log n)$	$O(n^{1/3} \log^{O(1)} n)$
Gottlieb and Roditty [18]	$O(\log^2 n)$	$O(\log^3 n)$
Gottlieb and Roditty [17]	$O(\log n)$	$O(\log n)$

■ **Table 1.2** Previous work on static k -vertex-fault-tolerant $(1 + \varepsilon)$ -spanners in \mathbb{R}^d . Dependencies on ε (of the form $1/\varepsilon^{O(d)}$) are omitted in the O bounds.

reference	# edges	degree	running time
Levcopoulos <i>et al.</i> [26]	$2^{O(k)} n$	$2^{O(k)}$	$O(n \log n + 2^{O(k)} n)$
	$O(k^2 n)$	unbounded	$O(n \log n + k^2 n)$
	$O(kn \log n)$	unbounded	$O(kn \log n)$
Lukovszki [28, 29]	$O(kn)$	$O(k^2)$	$O(n \log^{d-1} n + kn \log \log n)$
Czumaj and Zhao [11]	$O(kn)$	$O(k)$	$O(kn \log^d n + k^2 n \log k)$
H. Chan <i>et al.</i> [5]	$O(k^2 n)$	$O(k^2)$	$O(n \log n + k^2 n)$
Kapoor and Li [25]/Solomon [31]	$O(kn)$	$O(k)$	$O(n \log n + kn)$

- (c) **Dynamic approximate minimum spanning trees.** As is well-known [4, 19], a $(1 + \varepsilon)$ -approximate Euclidean MST of a point set P can be computed from the MST of a $(1 + \varepsilon)$ -spanner of P . In our dynamic spanner (and also Gottlieb and Roditty’s method [17]), each insertion/deletion of a point causes $O(1)$ edge updates to the graph. Immediately, we thus obtain a dynamic data structure for maintaining a $(1 + \varepsilon)$ -approximate Euclidean MST, with update time (ignoring dependencies on ε) equal to that for the dynamic graph MST problem, which is currently $O(\log^4 n / \log \log n)$ with amortization [22].
- (d) **Static and dynamic vertex-fault-tolerant spanners.** For parameters $k, t \geq 1$ and a set of points P in \mathbb{R}^d , a k -vertex-fault-tolerant t -spanner is a graph G which is a t -spanner and for any $P' \subseteq P$ of size at most k , the graph $G \setminus P'$ remains a t -spanner for $P \setminus P'$. Fault-tolerant spanners have been extensively studied (see Table 1.2). Locality-sensitive orderings lead to a very simple construction for k -vertex-fault-tolerant $(1 + \varepsilon)$ -spanners, with $O(kn)$ edges, maximum degree $O(k)$, and $O(n \log n + kn)$ running time (ignoring dependencies on ε). See Theorem 16. Although this result was known before, all previous constructions (including suboptimal ones), from Levcopoulos *et al.* [26] to Solomon’s work [31], as listed in Table 1.2, require intricate details. It is remarkable how effortlessly we achieve optimal $O(k)$ degree, compared to the previous methods. (Note, however, that some of the more recent previous constructions more generally apply to spaces with bounded doubling dimension, and some also achieve good bounds on other parameters such as the total weight and the hop-diameter.) Our algorithm can be easily made dynamic, with $O(\log n + k)$ update time. No previous results on dynamic fault-tolerant spanners were known.
- (e) **Approximate nearest neighbors.** Locality-sensitive orderings lead to a simple dynamic data structure for $(1 + \varepsilon)$ -approximate nearest neighbor search with $O(\log n)$ time per update/query (ignoring dependencies on ε). While this result is not new [7], we emphasize that the query algorithm is the simplest so far – it is just a binary search in the orderings maintained.

Computational models

The model of computation we have assumed is a unit-cost real RAM, supporting standard arithmetic operations and comparisons (but no floor function), augmented with bitwise-logical operations (bitwise-exclusive-or and bitwise-and), which are commonly available in programming languages (and in reality are cheaper than some arithmetic operations like multiplication).

If we assume that input coordinates are integers bounded by U and instead work in the word RAM model with $(\log U)$ -bit words ($U \geq n$), then our approach can actually yield *sublogarithmic* query/update time. For example, we can achieve $O(\log \log U)$ expected time for dynamic approximate bichromatic closest pair, dynamic spanners, and dynamic ANN, by replacing binary search with van Emde Boas trees [32]. Sublogarithmic algorithms were known before for dynamic ANN [7], but ours is the first sublogarithmic result for dynamic $(1 + \varepsilon)$ -spanners. Our results also answer the open problem of dynamic $(1 + \varepsilon)$ -approximate bichromatic closest pair in sublogarithmic time, originally posed by Chan and Skrepetos [10].

2 Locality-sensitive orderings

2.1 Grids and orderings

► **Definition.** For a set X , consider a total order (or ordering) \prec on the elements of X . Two elements $x, y \in X$ are **adjacent** if there is no element $z \in X$, such that $x \prec z \prec y$ or $y \prec z \prec x$.

Given two elements $x, y \in X$, such that $x \prec y$, the **interval** $[x, y)$ is the set $[x, y) = \{x\} \cup \{z \in X \mid x \prec z \prec y\}$.

The following is well known, and goes back to a work by Walecki in the 19th century [1]. We include a proof in Appendix A.1 for the sake of completeness.

► **Lemma 1.** *For n elements $\{0, \dots, n-1\}$, there is a set \mathfrak{D} of $\lceil n/2 \rceil$ orderings of the elements, such that, for all $i, j \in \{0, \dots, n-1\}$, there exist an ordering $\sigma \in \mathfrak{D}$ in which i and j are adjacent.*

► **Definition 2.** Consider an axis-parallel cube $\mathcal{C} \subseteq \mathbb{R}^d$ with side length ℓ . Partitioning it uniformly into a $t \times t \times \dots \times t$ grid \mathbf{G} creates the **t -grid** of \mathcal{C} . The grid \mathbf{G} is a set of t^d identically sized cubes with side length ℓ/t .

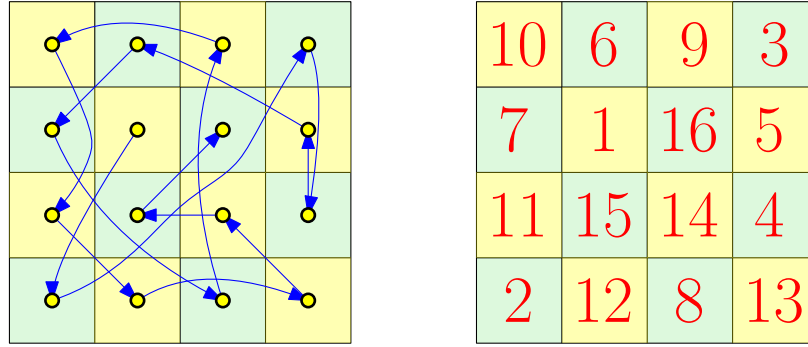
For a cube $\square \subseteq \mathbb{R}^d$, its **diameter** is $\text{diam}(\square) = \text{sidelength}(\square)\sqrt{d}$.

By Lemma 1 we obtain the following result.

► **Corollary 3.** *For a t -grid \mathbf{G} of an axis-parallel cube $\mathcal{C} \subseteq \mathbb{R}^d$, there is a set $\mathfrak{D}(t, d)$ of $O(t^d)$ orderings, such that for any $\square_1, \square_2 \in \mathbf{G}$, there exists an order $\sigma \in \mathfrak{D}(t, d)$ where \square_1 and \square_2 are adjacent in σ .*

2.2 ε -Quadtrees

► **Definition 4.** An **ε -quadtrees** \mathcal{T}_ε is a quadtree-like structure, built on a cube with side length ℓ , where each cell is being partitioned into a $(1/\varepsilon)$ -grid. The construction then continues recursively into each grid cell of interest. As such, a node in this tree has up to $1/\varepsilon^d$ children, and a node at level $i \geq 0$ has an associated cube of side length $\ell\varepsilon^i$. We call a $1/2$ -quadtrees a regular quadtree.



■ **Figure 2.1** One ordering of a set of cells.

► **Lemma 5.** *Let $E > 0$ be an integer number, $\varepsilon = 2^{-E}$, and \mathcal{T} be a regular quadtree over $[0, 2)^d$. Then there are ε -quadtrees $\mathcal{T}_\varepsilon^1, \dots, \mathcal{T}_\varepsilon^E$, such that the collection of cells at each level in \mathcal{T} is contained in exactly one of these ε -quadtrees.*

Proof. For $i = 0, \dots, E - 1$, construct the ε -quadtree $\mathcal{T}_\varepsilon^i$ using the cube $[0, 2^{E-i})^d \supseteq [0, 2)^d$ as the root. Now for $j \in \{0, \dots, E - 1\}$, observe that the collection of cells at levels $j, j + E, j + 2E, \dots$, of \mathcal{T} will also be in the quadtree $\mathcal{T}_\varepsilon^j$. Indeed, any node at level $j + \ell E$ in \mathcal{T} corresponds to a cell of side length $2^{-(j+\ell E)}$. Now in the $(\ell + 1)$ th level of quadtree $\mathcal{T}_\varepsilon^j$, this same node will have side length $\varepsilon^{\ell+1} 2^{E-j} = 2^{-(j+\ell E)}$. ◀

Consider an ε -quadtree \mathcal{T}_ε . Every node has up to $1/\varepsilon^d$ children. Consider any ordering σ of $\{1, \dots, 1/\varepsilon^d\}$, and consider a DFS of \mathcal{T}_ε that always visits the children of a node in the order specified by σ . This induces an ordering on the points in the cube which is the root of \mathcal{T}_ε . Indeed, for any two points, imagine storing them in an ε -quadtree – this implies that the two points are each stored in their own leaf node, which contains no other point of interest. Now, independently of what other points are stored in the quadtree, this DFS traversal would visit these two points in the same order. This can be viewed as a space filling curve (which is not continuous) which maps a cube to an interval. This is a generalization of the \mathcal{Z} -order. In particular, given a point set stored in \mathcal{T}_ε , and σ , one can order the points according to this DFS traversal, resulting in 1-dimensional ordering of the points. We denote the resulting ordering by $(\mathcal{T}_\varepsilon, \sigma)$.

► **Definition 6.** Let Π be the set of all orderings of $[0, 2)^d$, induced by picking one of the $\lg(1/\varepsilon)$ trees of Lemma 5, together with an ordering $\sigma \in \mathcal{D}(1/\varepsilon, d)$, as defined by Lemma 1. Each ordering in Π is called an **ε -ordering**.

Any two points that are lucky enough to lie in a cell of the quadtree that has diameter close to their distance, are going to be adjacent in one of the ε -orderings. Indeed, consider two points $p, q \in [0, 1)^d$, a parameter $\varepsilon > 0$, such that p, q are both contained in a cell \square of the regular quadtree \mathcal{T} with $\|p - q\| > \varepsilon \text{diam}(\square)$. Then, there is an ε -quadtree \mathcal{T}_ε that has \square as a node, and let \square_p and \square_q be the two children of \square in \mathcal{T}_ε , containing p and q respectively. Furthermore, there is an ordering $\sigma \in \mathcal{D}(1/\varepsilon, d)$, such that \square_p and \square_q are adjacent. As such, the cube \square_p (resp., \square_q) corresponds to an interval $[x, x')$ (resp., $[x', x'')$) in the ordering $(\mathcal{T}_\varepsilon, \sigma)$, and these two intervals are adjacent.

Now consider the case when there are two points close together, but no appropriately sized quadtree cell contains both p and q . In other words, two points that are close together might get separated by nodes that are much bigger in the quadtree. This can be resolved using shifting. We need the following result of Chan [6, Lemma 3.3] – a proof is provided in Appendix A.2.

► **Lemma 7.** Consider any two points $p, q \in [0, 1]^d$, and let \mathcal{T} be the infinite quadtree of $[0, 2]^d$. For $D = 2 \lceil d/2 \rceil$ and $i = 0, \dots, D$, let $v_i = (i/(D+1), \dots, i/(D+1))$. Then there exists an $i \in \{0, \dots, D\}$, such that $p + v_i$ and $q + v_i$ are contained in a cell of \mathcal{T} with side length $\leq 2(D+1) \|p - q\|$.

2.3 Comparing two points according to an ε -ordering

We now show how to efficiently compare two points in P according to a given ε -ordering σ with a shift v_i . The shift can be added to the two points directly, and as such only need to worry about how to compare two points according to σ .

► **Observation 8.** Let \oplus denote the bitwise-exclusive-or operator. Define $\text{msb}(a) := -\lfloor \lg a \rfloor$ to be the index of the most significant bit in the binary expansion of $a \in (0, 2]$. Given $a, b \in [0, 2)$, one can compare the msb of two numbers using the following:

- (a) $\text{msb}(a) > \text{msb}(b)$ if and only if $a < b$ and $a < a \oplus b$.
- (b) If $a \oplus b \leq a \wedge b$ then $\text{msb}(a) = \text{msb}(b)$, where \wedge is the bitwise-and operator.

► **Lemma 9.** Let $p = (p_1, \dots, p_d)$ and $q = (q_1, \dots, q_d)$ be two distinct points in $P \subseteq [0, 2]^d$ and $\sigma \in \Pi$ be an ε -ordering over the cells of some ε -quadtree \mathcal{T}_ε storing P . Then one can determine if $p \prec_\sigma q$ using $O(d \log(1/\varepsilon))$ bitwise-logical operations.

Proof. Recall ε is a power of two and $E = \lg(1/\varepsilon)$. In order to compare p and q , for $i = 1, \dots, d$, compute $a_i = p_i \oplus q_i$. Find the index i' such that $\text{msb}(a_{i'}) \leq \text{msb}(a_i)$ for all i , using $O(d)$ comparisons. Given $p_{i'}$ and $q_{i'}$, our first goal is to determine the place in which $p_{i'}$ and $q_{i'}$ first differ in their binary representation. Note that because ε is a power of two, each bit in the base $1/\varepsilon$ expansion of $p_{i'}$ corresponds to a block of E bits in the binary expansion of $p_{i'}$.

First we find the position within the block of size E where $p_{i'}$ and $q_{i'}$ differ in their binary expansion. For $j = 1, \dots, E$, let $b_j = 2^{E-j}/(2^E - 1) \in (0, 1)$ be the number whose binary expansion has a 1 in positions $j, j + E, j + 2E, \dots$, and 0 everywhere else. For $j = 1, \dots, E$, compute $b_j \wedge a_{i'}$ and check if $\text{msb}(b_j) = \text{msb}(b_j \wedge a_{i'})$. Once we find a j such that $\text{msb}(b_j) = \text{msb}(b_j \wedge a_{i'})$, stop. We know that that $p_{i'}$ and $q_{i'}$ first differ in the j th position inside some block.

It remains to extract the E bits from each block in p_1, \dots, p_d . For $i = 1, \dots, d$, let $B_i \in \{0, 1\}^E$ be the bits inside the block associated with p_i . For $k = 1, \dots, E$, set $B_{i,k} = \mathbb{1}[\text{msb}(2^{j-k} a_{i'}) = \text{msb}((2^{j-k} a_{i'}) \wedge p_i)]$ (where $\mathbb{1}[\cdot]$ is the indicator function). By repeating a similar process for all q_1, \dots, q_d , we obtain the coordinates of the cells in which p and q differ. We can then consult σ to determine whether or not $p \prec_\sigma q$.

This implies that p and q can be compared using $O(d \log(1/\varepsilon))$ operations by Observation 8. ◀

Remark

In the word RAM model for integer input, the extra $\log(1/\varepsilon)$ factor in the above time bound can be eliminated: msb can be explicitly computed in $O(1)$ time by a complicated algorithm of Fredman and Willard [15]; this allows us to directly jump to the right block of each coordinate and extract the relevant bits.

2.4 The result

► **Theorem 10.** *For a parameter $\varepsilon \in (0, 1)$, there is a set Π^+ of $O((1/\varepsilon^d) \log(1/\varepsilon))$ orderings of $[0, 1]^d$, such that for any two points $p, q \in [0, 1]^d$ there is an ordering $\sigma \in \Pi^+$ defined over $[0, 1]^d$, such that for any point u with $p \prec_\sigma u \prec_\sigma q$ it holds that either $\|p - u\| \leq \varepsilon \|p - q\|$ or $\|q - u\| \leq \varepsilon \|p - q\|$.*

Furthermore, given such a ordering σ , and two points p, q , one can compute their ordering, according to σ , using $O(d \log(1/\varepsilon))$ arithmetic and bitwise-logical operations.

Proof. Let Π^+ be the set of all ordering defined by picking an ordering from Π , as defined by Definition 6 using the parameter ε , together with a shift from Lemma 7.

Consider any two points $p, q \in [0, 1]^d$. By Lemma 7 there is a shift for which the two points fall into a quadtree cell \square with side length at most $2(D+1)\|p - q\|$. Next, there is an ε -quadtree \mathcal{T}_ε that contains \square , and the two children that correspond to two cells \square_p and \square_q with side length at most $2(D+1)\varepsilon\|p - q\|$, which readily implies that the diameter of these cells is at most $2(D+1)\sqrt{d}\varepsilon\|p - q\|$. Furthermore, there is an ε -ordering in Π such that all the points of \square_p are adjacent to all the points \square_q in this ordering. This implies the desired claim, after adjusting ε by a factor of $2(D+1)\sqrt{d}$ (and rounding to a power of 2). ◀

From now on, we refer to the set of orderings Π^+ in the above Theorem as locality-sensitive orderings. We remark that by the readjustment of ε in the final step of the proof, the number of locality-sensitive orderings when including the factors involving d is $O(d^{3d/2}(1/\varepsilon^d) \log(1/\varepsilon))$.

2.4.1 Discussion

Connection to locality-sensitive hashing

Let P be a set of n points in $\{0, 1\}^d$. Consider the decision version of the $(1 + \varepsilon)$ -approximate nearest neighbor problem. Specifically, for a pre-specified radius r and any given query point q , we would like to efficiently decide whether or not there exists a point $p \in P$ such that $\|q - p\|_1 \leq (1 + \varepsilon)r$ or conclude that all points in P are at least distance r from q . The locality-sensitive hashing (LSH) technique [23] can be used to create a data structure supporting these types of decision queries in time $O(dn^{1/(1+\varepsilon)} \log n)$ time (which is correct with high probability) and using total space $O(dn^{1+1/(1+\varepsilon)} \log n)$. Similar results also hold in Euclidean space.

At a high level, LSH works as follows. Start by choosing $k := k(\varepsilon, r, n)$ indices in $[d]$ at random (with replacement). Let R denote the resulting multiset of coordinates. For each point $p \in P$, let p_R be the projection p onto these coordinates of R . We can group the points of P into buckets, where each bucket contains points with the same projection. Given a query point q , we check if any of the points in the same bucket as q_R is at distance at most $(1 + \varepsilon)r$ from q . This construction can also be repeated a sufficient number of times in order to guarantee success with high probability.

The idea of bucketing can also be viewed as an implicit ordering on the randomly projected point set by ordering points lexicographically according to the k coordinates. In this sense, the query algorithm can be viewed as locating q within each of the orderings, and comparing q to similar points nearby in each ordering. From this perspective, every locality-sensitive ordering can be viewed as a LSH scheme. Indeed, for a given query point q , the approximate nearest neighbor to q can be found by inspecting the elements adjacent to q in each of the locality-sensitive orderings and returning the closest point to q found (see Theorem 17).

Of course, the main difference between the two schemes is that for every fixed ε , the number of “orderings” in a LSH scheme is polynomial in both d and n . While for locality-sensitive orderings, the number of orderings remain exponential in d . This trade-off is required, as locality-sensitive orderings guarantee a much stronger property than that of a LSH scheme.

Extension of locality-sensitive orderings to other norms in Euclidean space

The result of Theorem 10 also holds for any L_p norm. The key change that is needed is in the proof of Lemma 7: For any two points $s, t \in [0, 1]^d$, there exists a shift v such that $s + v$ and $t + v$ are contained in a quadtree cell of side length at most $2(D + 1) \|s - t\|_p$. This extension follows easily from the proof of the Lemma, see Appendix A.2. Theorem 10 then follows by adjusting ε by a factor of $2(D + 1)d^{1/p}$ in the last step, implying that the number of orderings will be $O(d^{d(1+1/p)}(1/\varepsilon^d) \log(1/\varepsilon))$.

Extension of locality-sensitive orderings for doubling metrics

An abstraction of low-dimensional Euclidean space, is a metric space with (low) doubling dimension. Formally, a metric space (\mathcal{M}, d) has *doubling dimension* λ if any ball of \mathcal{M} of radius r can be covered by at most 2^λ balls of half the radius (i.e., $r/2$). It is known that \mathbb{R}^d has doubling dimension $O(d)$ [33]. We point out that locality-sensitive orderings still exist in this case, but they are less constructive in nature, since one needs to be provided with all the points of interest in advance.

For a point set $P \subseteq \mathcal{M}$, the analogue of a quadtree for a metric space is a net tree [20]. Being somewhat imprecise, a net tree can be constructed as follows: The root node corresponds to the point set $P \subseteq \mathcal{M}$. Compute a randomized partition of P of diameter $1/2$ (assume P has diameter one), and for each cluster in the partition, create an associated node and hang it on the root. The tree is computed recursively in this manner, at each level i computing a random partition of diameter 2^{-i} . The leaves of the tree are points of P .

As with quadtrees, it is possible during this randomized construction for two nearby points to be placed in different clusters and be separated further down the tree. If $\ell = d(p, q)$ for two points $p, q \in P$, then the probability that p and q lie in different clusters of diameter $r = 2^{-i}$ in the randomized partition is at most $O((\ell/r) \log n)$ [13]. In particular, for $r \approx 1/(\ell \log n)$, the probability p and q are separated is at most a constant. If we want even this property to hold with high probability for all pairs of points, one needs to construct $O(\log n)$ (randomly constructed) net trees of P . (This corresponds to randomly shifting a quadtree $O(\log n)$ times in the Euclidean setting.)

Given such a net tree T , each node has $I = 2^{O(\lambda)}$ children. We can arbitrarily and explicitly number the children of each node by a distinct label from $[I]$. One can define an ordering of such a tree as we did in the Euclidean case, except that the gap (in diameter) between a node and its children is $O(\varepsilon/\log n)$ instead of ε . Repeating our scheme in the Euclidean case, this implies that one would expect to require $(\varepsilon^{-1} \log n)^{O(\lambda)}$ orderings of P .

This requires having all the points of P in advance, which is a strong assumption for a dynamic data structure (like the applications we show next). For example, Gottlieb and Roditty [17] show how to maintain dynamic spanners in a doubling metric, but only assuming that after a point has been deleted from P , the distance between the deleted point and a point currently in P can still be computed in constant time.

3 Applications

3.1 Bichromatic closest pair

Given an ordering $\sigma \in \Pi^+$, and two finite sets of points R, B in \mathbb{R}^d , let $\mathcal{Z} = \mathcal{Z}(\sigma, R, B)$ be the set of all pairs of points in $R \times B$ that are adjacent in the ordering of $R \cup B$ according to σ . Observe that inserting or deleting a single point from these two sets changes the content of \mathcal{Z} by a constant number of pairs. Furthermore, a point participates in at most two pairs.

► **Lemma 11.** *Let R and B be two sets of points in $[0, 1]^d$, and let $\varepsilon \in (0, 1)$ be a parameter. Let $\sigma \in \Pi^+$ be a locality-sensitive ordering (see Theorem 10). Then, one can maintain the set $\mathcal{Z} = \mathcal{Z}(\sigma, R, B)$ under insertions and deletions to R and B . In addition, one can maintain the closest pair in \mathcal{Z} (under the Euclidean metric). Each update takes $O(d \log(n) \log(1/\varepsilon))$ time, where n is the total size of R and B during the update operation.*

Proof. Maintain two balanced binary search trees T_R and T_B storing the points in R and B , respectively, according to the order σ . Insertion, deletion, predecessor query and successor query can be implemented in $O(d \log(1/\varepsilon) \log n)$ time (since any query requires $O(\log n)$ comparisons costing each $O(d \log(1/\varepsilon))$ time by Lemma 9). We also maintain a min-heap of the pairs in \mathcal{Z} sorted according to the Euclidean distance. The minimum is the desired closest pair. Notice that a single point can participate in at most two pairs in \mathcal{Z} .

We now explain how to handle updates. Given a newly inserted point r (say a red point that belongs to R), we compute its (potential) pairs it participates in, by computing its successor r' in R , and its successor b' in B . If $r \prec_\sigma b' \prec_\sigma r'$ then the new pair rb' should be added to \mathcal{Z} . The pair before r in the ordering that might use r is computed in a similar fashion. In addition, we recompute the predecessor and successor of r in R , and we recompute the pairs they might participate in (deleting potentially old pairs that are no longer valid).

Deletion is handled in a similar fashion – all points included in pairs with the deleted point recompute their pairs. In addition, the successor and predecessor (of the same color) need to recompute their pairs. This all requires a constant number of queries in the two trees, and thus takes the running time as stated. ◀

► **Theorem 12.** *Let R and B be two sets of points in $[0, 1]^d$, and let $\varepsilon \in (0, 1)$ be a parameter. Then one can maintain a $(1 + \varepsilon)$ -approximation to the bichromatic closest pair in $R \times B$ under updates (i.e., insertions and deletions) in $O(\log(n) \log^2(1/\varepsilon)/\varepsilon^d)$ time per operation, where n is the total number of points in the two sets. The data structure uses $O(n \log(1/\varepsilon)/\varepsilon^d)$ space, and at all times maintains a pair of points $r \in R$, $b \in B$, such that $\|r - b\| \leq (1 + \varepsilon)d(R, B)$, where $d(R, B) = \min_{r \in R, b \in B} \|r - b\|$.*

Proof. We maintain the data structure of Lemma 11 for all the locality-sensitive orderings of Theorem 10. We might as well maintain all the good pairs for these data structures together in one global min-heap. We claim that the minimum length pair in this heap is the desired approximation.

To see that, consider the bichromatic closest pair $r \in R$ and $b \in B$. By Theorem 10 there is a locality-sensitive ordering σ , such that the interval I in the ordering between r and b contains points that are in distance at most $\ell = \varepsilon \|r - b\|$ from either r or b . In particular, let P_r (resp., P_b) be all the points in I in distance at most ℓ from r (resp., b). Observe that $P_r \subseteq R$, as otherwise, there would be a bichromatic pair in P_r , and since the diameter of this set is at most ℓ , this would imply that (r, b) is not the closest bichromatic pair – a contradiction. Similarly, $P_b \subseteq B$. As such, there must be two points $b' \in B$ and $r' \in R$, that are consecutive in σ , and this is one of the pairs considered by the algorithm (as it is stored

in the min-heap). In particular, by the triangle inequality, we have

$$\|r' - b'\| \leq \|r' - r\| + \|r - b\| + \|b - b'\| \leq 2\ell + \|r - b\| \leq (1 + 2\varepsilon) \|r - b\|.$$

The theorem follows after adjusting ε by a factor of 2. \blacktriangleleft

Remark

In the word RAM model, for integer input in $\{1, \dots, U\}^d$, the update time can be improved to $O((\log \log U) \log^2(1/\varepsilon)/\varepsilon^d)$ expected, by using van Emde Boas trees [32] in place of the binary search trees (and the min-heaps as well). With standard word operations, we may not be able to explicitly map each point to an integer in one dimension following each locality-sensitive ordering, but we can still simulate van Emde Boas trees on the input as if the mapping has been applied. Each recursive call in the van Emde Boas recursion focuses on a specific block of bits of each input coordinate value (after shifting); we can extract these blocks, and perform the needed hashing operations on the concatenation of these blocks over the d coordinates of each point.

3.2 Dynamic spanners

► **Definition 13.** For a set of n points P in \mathbb{R}^d and a parameter $t \geq 1$, a t -**spanner** of P is an undirected graph $G = (P, E)$ such that for all $p, q \in P$,

$$\|p - q\| \leq d_G(p, q) \leq t\|p - q\|,$$

where $d_G(p, q)$ is the length of the shortest path from p to q in G using the edge set E .

Using a small modification of the results in the previous section, we easily obtain a dynamic $(1 + \varepsilon)$ -spanner. Note that there is nothing special about how the data structure in Theorem 12 deals with the bichromatic point set. If the point set is monochromatic, modifying the data structure in Lemma 11 to account for the closest monochromatic pair of points leads to a data structure with the same bounds and maintains the $(1 + \varepsilon)$ -approximate closest pair.

The construction of the spanner is very simple: Given P and $\varepsilon \in (0, 1)$, maintain orderings of the points specified by Π^+ (see Theorem 10). For each $\sigma \in \Pi^+$, let E_σ be the edge set consisting of edges connecting two consecutive points according to σ , with weight equal to their Euclidean distance. Thus $|E_\sigma| = n - 1$. Our spanner $G = (P, E)$ then consists of the edge set $E = \bigcup_{\sigma \in \Pi^+} E_\sigma$.

► **Theorem 14.** *Let P be a set of n points in $[0, 1]^d$ and $\varepsilon \in (0, 1)$. One can compute a $(1 + \varepsilon)$ -spanner G of P with $O(n \log(1/\varepsilon)/\varepsilon^d)$ edges, where every vertex has degree $O(\log(1/\varepsilon)/\varepsilon^d)$. Furthermore, insertions and deletions can be performed in $O(\log(n) \log^2(1/\varepsilon)/\varepsilon^d)$ time, with at most $O(\log(1/\varepsilon)/\varepsilon^d)$ edges being deleted or inserted into the spanner.*

Proof. The construction is described above. The same analysis as in the proof of Theorem 12 implies the number of edges in G and the update time.

It remains to prove that G is a spanner. By Theorem 10, for any pair of points $s, t \in P$, there is a locality-sensitive ordering $\sigma \in \Pi^+$, such that the σ -interval $[s, t)$ contains only points that are in distance at most $\varepsilon \|s - t\|$ from either s or t . In particular, there must be two points in $s', t' \in P$ that are adjacent in σ , such that one of them, say s' (resp., t') is in distance at most $\varepsilon \|s - t\|$ from s (resp., t). As such, the edge $s't'$ exists in the graph being maintained.

This property is already enough to imply that this graph is a $(1 + c\varepsilon)$ -spanner for a sufficiently large constant c – this follows by an induction on the distances between the points (specifically, in the above, we apply the induction hypothesis on the pairs s, s' and t, t'). We omit the easy but somewhat tedious argument – see [4] or [19, Theorem 3.12] for details. The theorem follows after adjusting ε by a factor of c . ◀

3.2.1 Static and dynamic vertex-fault-tolerant spanners

► **Definition 15.** For a set of n points P in \mathbb{R}^d and a parameter $t \geq 1$, a k -**vertex-fault-tolerant t -spanner** of P , denoted by (k, t) -VF T S, is a graph $G = (P, E)$ such that

- (i) G is a t -spanner (see Definition 13), and
- (ii) For any $P' \subseteq P$ of size at most k , the graph $G \setminus P'$ is a t -spanner for $P \setminus P'$.

A $(k, 1 + \varepsilon)$ -VF T S can be obtained by modifying the construction of the $(1 + \varepsilon)$ -spanner in Section 3.2. Construct a set of locality-sensitive orderings Π^+ . For each $\sigma \in \Pi^+$ and each $p \in P$, connect p to its $k + 1$ successors and $k + 1$ predecessors according to σ with edge weights equal to the Euclidean distances. Thus each ordering maintains $O(nk)$ edges and there are $O(|\Pi^+|kn) = O(kn \log(1/\varepsilon)/\varepsilon^d)$ edges overall. We now prove that this graph G is in fact a $(k, 1 + \varepsilon)$ -VF T S.

► **Theorem 16.** *Let P be a set of n points in $[0, 1]^d$ and $\varepsilon \in (0, 1)$. One can compute a k -vertex-fault-tolerant $(1 + \varepsilon)$ -spanner G for P in time $O((n \log n \log(1/\varepsilon) + kn) \log(1/\varepsilon)/\varepsilon^d)$. The number of edges is $O(kn \log(1/\varepsilon)/\varepsilon^d)$ and the maximum degree is bounded by $O(k \log(1/\varepsilon)/\varepsilon^d)$.*

Furthermore, one can maintain the k -vertex-fault-tolerant $(1 + \varepsilon)$ -spanner G under insertions and deletions in $O((\log n \log(1/\varepsilon) + k) \log(1/\varepsilon)/\varepsilon^d)$ time per operation.

Proof. The construction algorithm, number of edges, and maximum degree follows from the discussion above. So, consider deleting a set $P' \subseteq P$ of size at most k from G . Consider an ordering $\sigma \in \Pi^+$ with the points P' removed. By the construction of G , all the pairs of points of $P \setminus P'$ that are (now) adjacent in σ remain connected by an edge in $G \setminus P'$. The argument of Theorem 14 implies that the remaining graph is spanner. We conclude that that $G \setminus P'$ is a $(1 + \varepsilon)$ -spanner for $P \setminus P'$.

As for the time taken to handle insertions and deletions, one simply maintains the orderings of the points using balanced search trees. After an insertion to one of the orderings in $O(\log n \log(1/\varepsilon))$ time, $O(k)$ edges have to be added and deleted. Therefore insertions take $O((\log n \log(1/\varepsilon) + k) |\Pi^+|) = O((\log n \log(1/\varepsilon) + k) \log(1/\varepsilon)/\varepsilon^d)$ time total. Deletions are handled similarly.

The total construction time follows by inserting each of the points into the dynamic data structure. ◀

3.3 Dynamic approximate nearest neighbors

Another application of the same data structure in Theorem 12 is supporting $(1 + \varepsilon)$ -approximate nearest neighbor queries. In this scenario, the data structure must support insertions and deletions of points and the following queries: given a point q , return a point $t \in P$ such that $\|q - t\| \leq (1 + \varepsilon) \min_{p \in P} \|q - p\|$.

► **Theorem 17.** *Let P be a set of n points in $[0, 1]^d$. For a given $\varepsilon > 0$, one can build a data structure using $O(n \log(1/\varepsilon)/\varepsilon^d)$ space, that supports insertion and deletion in time $O(\log(n) \log^2(1/\varepsilon)/\varepsilon^d)$. Furthermore, given a query point $q \in [0, 1]^d$, the data structure returns a $(1 + \varepsilon)$ -approximate nearest neighbor in P in $O(\varepsilon^{-d} \log^2(1/\varepsilon) \log(n))$ time.*

Proof. Maintain the data structure of Lemma 11 for all locality-sensitive orderings of Theorem 10, with one difference: Since the input is monochromatic, the data structures for each ordering store distances between all consecutive pairs. The space and update time bounds easily follow by the same analysis.

Given a query point $q \in [0, 1]^d$, for each of the orderings the algorithm inspects the predecessor and successor to q . The algorithm returns the closest point to q encountered. We claim that the returned point p is the desired approximate nearest neighbor.

Let $p^* \in P$ be the nearest neighbor to q and $\ell = \|q - p^*\|$. By Theorem 10, there is a locality-sensitive ordering $\sigma \in \Pi^+$ such that the σ -interval $I = [p^*, q)$ contains points that are of distance at most $\varepsilon\ell$ from p^* or q (and this interval contains at least one point of P , namely, p^*). Note that no point of P can be at distance less than $\varepsilon\ell$ to q . Thus, the point $p \in P$ adjacent to q in I is of distance at most $\varepsilon\ell$ from p^* . Therefore, for such a point p , we have $\|p - q\| \leq \|p - p^*\| + \|p^* - q\| \leq (1 + \varepsilon)\ell$.

The final query time follows from the time taken for these predecessor and successor queries, as in the proof of Lemma 11. ◀


References

- 1 Brian Alspach. The wonderful Walecki construction. *Bull. Inst. Combin. Appl.*, 52:7–20, 2008.
- 2 Sanjeev Arora. Polynomial time approximation schemes for Euclidean TSP and other geometric problems. *J. Assoc. Comput. Mach.*, 45(5):753–782, September 1998. URL: <http://www.cs.princeton.edu/~arora/pubs/tsp.ps>.
- 3 Marshall W. Bern. Approximate Closest-Point Queries in High Dimensions. *Inform. Process. Lett.*, 45(2):95–99, 1993. doi:10.1016/0020-0190(93)90222-U.
- 4 Paul B. Callahan and S. Rao Kosaraju. Faster Algorithms for Some Geometric Graph Problems in Higher Dimensions. In Vijaya Ramachandran, editor, *Proc. 4th ACM-SIAM Sympos. Discrete Alg. (SODA)*, pages 291–300. ACM/SIAM, 1993. URL: <http://dl.acm.org/citation.cfm?id=313559.313777>.
- 5 T.-H. Hubert Chan, Mingfei Li, Li Ning, and Shay Solomon. New Doubling Spanners: Better and Simpler. *SIAM J. Comput.*, 44(1):37–53, 2015. doi:10.1137/130930984.
- 6 Timothy M. Chan. Approximate nearest neighbor queries revisited. *Discrete Comput. Geom.*, 20(3):359–373, 1998. doi:10.1007/PL00009390.
- 7 Timothy M. Chan. Closest-point problems simplified on the RAM. In *Proc. 13th ACM-SIAM Sympos. Discrete Alg. (SODA)*, pages 472–473. SIAM, 2002. URL: <http://dl.acm.org/citation.cfm?id=545381.545444>.
- 8 Timothy M. Chan. A minimalist’s implementation of an approximate nearest neighbor algorithm in fixed dimensions, 2006. URL: <http://tmc.web.engr.illinois.edu/sss.ps>.
- 9 Timothy M. Chan. Well-separated pair decomposition in linear time? *Inform. Process. Lett.*, 107(5):138–141, 2008. doi:10.1016/j.ipl.2008.02.008.
- 10 Timothy M. Chan and Dimitrios Skrepetos. Dynamic data structures for approximate Hausdorff distance in the word RAM. *Comput. Geom. Theory Appl.*, 60:37–44, 2017. doi:10.1016/j.comgeo.2016.08.002.
- 11 Artur Czumaj and Hairong Zhao. Fault-Tolerant Geometric Spanners. *Discrete Comput. Geom.*, 32(2):207–230, 2004. URL: <http://www.springerlink.com/index/10.1007/s00454-004-1121-7>.
- 12 David Eppstein. Dynamic Euclidean minimum spanning trees and extrema of binary functions. *Discrete Comput. Geom.*, 13:111–122, 1995. URL: <http://www.ics.uci.edu/~eppstein/pubs/Epp-DCG-95.pdf>.

- 13 Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. Sys. Sci.*, 69(3):485–497, 2004. doi:10.1016/j.jcss.2004.04.011.
- 14 Uriel Feige and Robert Krauthgamer. Stereoscopic families of permutations, and their applications. In *Proc. 5th Israel Symp. Theo. Comput. and Systems (ISTCS)*, pages 85–95. IEEE Computer Society, 1997. doi:10.1109/ISTCS.1997.595160.
- 15 Michael L. Fredman and Dan E. Willard. Surpassing the Information Theoretic Bound with Fusion Trees. *J. Comput. Sys. Sci.*, 47(3):424–436, 1993. doi:10.1016/0022-0000(93)90040-4.
- 16 Irene Gargantini. An Effective Way to Represent Quadrees. *Commun. ACM*, 25(12):905–910, 1982. doi:10.1145/358728.358741.
- 17 Lee-Ad Gottlieb and Liam Roditty. An Optimal Dynamic Spanner for Doubling Metric Spaces. In *Proc. 16th Annu. Euro. Sympos. Alg. (ESA)*, pages 478–489, 2008. doi:10.1007/978-3-540-87744-8_40.
- 18 Lee-Ad Gottlieb and Liam Roditty. Improved algorithms for fully dynamic geometric spanners and geometric routing. In *Proc. 19th ACM-SIAM Sympos. Discrete Alg. (SODA)*, pages 591–600, 2008. URL: <http://dl.acm.org/citation.cfm?id=1347082.1347148>.
- 19 Sariel Har-Peled. *Geometric Approximation Algorithms*, volume 173 of *Math. Surveys & Monographs*. Amer. Math. Soc., Boston, MA, USA, 2011. doi:10.1090/surv/173.
- 20 Sariel Har-Peled and Manor Mendel. Fast Construction of Nets in Low Dimensional Metrics, and Their Applications. *SIAM J. Comput.*, 35(5):1148–1184, 2006. doi:10.1137/S0097539704446281.
- 21 Dorit S. Hochbaum and Wolfgang Maass. Approximation Schemes for Covering and Packing Problems in Image Processing and VLSI. *J. Assoc. Comput. Mach.*, 32(1):130–136, 1985. doi:10.1145/2455.214106.
- 22 Jacob Holm, Eva Rotenberg, and Christian Wulff-Nilsen. Faster Fully-Dynamic Minimum Spanning Forest. In Nikhil Bansal and Irene Finocchi, editors, *Proc. 23rd Annu. Euro. Sympos. Alg. (ESA)*, volume 9294 of *Lect. Notes in Comp. Sci.*, pages 742–753. Springer, 2015. doi:10.1007/978-3-662-48350-3_62.
- 23 Piotr Indyk and Rajeev Motwani. Approximate Nearest Neighbors: Towards Removing the Curse of Dimensionality. In *Proc. 30th ACM Sympos. Theory Comput. (STOC)*, pages 604–613, 1998. doi:10.1145/276698.276876.
- 24 Ibrahim Kamel and Christos Faloutsos. On Packing R -trees. In Bharat K. Bhargava, Timothy W. Finin, and Yelena Yesha, editors, *Proc. 2nd Intl. Conf. Info. Knowl. Mang.*, pages 490–499. ACM, 1993. doi:10.1145/170088.170403.
- 25 Sanjiv Kapoor and Xiang-Yang Li. Efficient Construction of Spanners in d -Dimensions. *CoRR*, abs/1303.7217, 2013. URL: <http://arxiv.org/abs/1303.7217>.
- 26 Christos Levcopoulos, Giri Narasimhan, and Michiel H. M. Smid. Efficient Algorithms for Constructing Fault-Tolerant Geometric Spanners. In Jeffrey Scott Vitter, editor, *Proc. 30th ACM Sympos. Theory Comput. (STOC)*, pages 186–195. ACM, 1998. doi:10.1145/276698.276734.
- 27 Swanwa Liao, Mario A. López, and Scott T. Leutenegger. High Dimensional Similarity Search With Space Filling Curves. In *Proc. 17th Int. Conf. on Data Eng. (ICDE)*, pages 615–622, 2001. doi:10.1109/ICDE.2001.914876.
- 28 Tamás Lukovszki. New Results of Fault Tolerant Geometric Spanners. In Frank K. H. A. Dehne, Arvind Gupta, Jörg-Rüdiger Sack, and Roberto Tamassia, editors, *Proc. 6th Workshop Alg. Data Struct. (WADS)*, volume 1663 of *Lect. Notes in Comp. Sci.*, pages 193–204. Springer, 1999. doi:10.1007/3-540-48447-7_20.
- 29 Tamás Lukovszki. *New results on geometric spanners and their applications*. PhD thesis, University of Paderborn, Germany, 1999. URL: <http://d-nb.info/958256713>.

- 30 Liam Roditty. Fully Dynamic Geometric Spanners. *Algorithmica*, 62(3-4):1073–1087, 2012. doi:10.1007/s00453-011-9504-7.
- 31 Shay Solomon. From hierarchical partitions to hierarchical covers: Optimal fault-tolerant spanners for doubling metrics. In David B. Shmoys, editor, *Proc. 46th ACM Sympos. Theory Comput. (STOC)*, pages 363–372. ACM, 2014. doi:10.1145/2591796.2591864.
- 32 Peter van Emde Boas. Preserving Order in a Forest in Less Than Logarithmic Time and Linear Space. *Inf. Process. Lett.*, 6(3):80–82, 1977. doi:10.1016/0020-0190(77)90031-X.
- 33 Jean-Louis Verger-Gaugry. Covering a Ball with Smaller Equal Balls in \mathbb{R}^n . *Discrete Comput. Geom.*, 33(1):143–155, 2005. doi:10.1007/s00454-004-2916-2.

A Proofs



■ **Figure A.1** For n even, a decomposition of K_n into $n/2$ Hamiltonian paths.

A.1 Proof of Lemma 1

► **Lemma 1 (Restated).** For n elements $\{0, \dots, n-1\}$, there is a set \mathfrak{D} of $\lceil n/2 \rceil$ orderings of the elements, such that, for all $i, j \in \{0, \dots, n-1\}$, there exist an ordering $\sigma \in \mathfrak{D}$ in which i and j are adjacent.

Proof. As mentioned earlier this is well known, see [1]. Assume n is even, and consider the clique K_n , with its vertices v_0, \dots, v_{n-1} . The edges of this clique can be covered by $n/2$ Hamiltonian paths that are edge disjoint. Tracing one of these path gives rise to one ordering, and doing this for all paths results with orderings with the desired property, since edge $v_i v_j$ is adjacent in one of these paths.

To get this cover, draw K_n by using the vertices of an n -regular polygon, and draw all the edges of K_n as straight segments. For every edge $v_i v_{i+1}$ of K_n there are exactly $n/2$ parallel edges with this slope (which form a matching). Let M_i denote this matching. Similarly, for the vertex v_i , consider the segment $v_i v_{i+n/2}$ (indices are here modulo n), and the family of segments (i.e., edges) of K_n that are orthogonal to this segment. This family is also a matching M'_i of size $n/2 - 1$. Observe that $\sigma_i = M_i \cup M'_i$ form a Hamiltonian path, as shown in Figure A.1. Since the slopes of the segments in M_i and M'_i are unique, for $i = 0, \dots, n/2 - 1$, it follows that $\sigma_0, \dots, \sigma_{n/2-1}$ are an edge disjoint cover of all the edges of K_n by $n/2$ Hamiltonian paths.

If n is odd, use the above construction for $n + 1$, and delete the redundant symbol from the computed orderings. ◀

A.2 Proof of 7 (shifting)

For two positive real numbers x and y , let

$$x \% y = x - y \lfloor x/y \rfloor.$$

The basic idea behind shifting is that one can pick a set of values that look the “same” in all resolutions.

► **Lemma 18.** *Let $n > 1$ be a positive odd integer, and consider the set $X = \{i/n \mid i = 0, \dots, n-1\}$. Then, for any $\alpha = 2^{-\ell}$, where $\ell \geq 0$ is integer, we have that $X \% \alpha = \{i/n \% \alpha \mid i = 0, \dots, n-1\}$ is equal to the set $\alpha X = \{\alpha i/n \mid i = 0, \dots, n-1\}$.*

Proof. The proof is by induction. For $\ell = 0$ the claim clearly holds. Next, assume the claim holds for some $i \geq 0$, and consider $\ell = i + 1$. Setting $m = (n-1)/2$ and $\Delta = 2^{-i}/n$, we have by induction (and rearrangement) that

$$\begin{aligned} X \% 2^{-i} &= 2^{-i} X = \{0, \Delta, \dots, 2m\Delta\} \\ &= \{0, (m+1)\Delta, \Delta, (m+2)\Delta, 2\Delta, \dots, (m+m)\Delta, m\Delta\}. \end{aligned}$$

Settings $\delta = \Delta/2 = 2^{-i-1}/n$, and observing that $(2m+1)\delta = n\delta = 2^{-i-1}$, we have

$$\begin{aligned} X \% 2^{-i-1} &= (X \% 2^{-i}) \% 2^{-i-1} \\ &= \{0, (m+1)\Delta, \Delta, (m+2)\Delta, 2\Delta, \dots, (m+j)\Delta, j\Delta, \dots, (m+m)\Delta, m\Delta\} \% 2^{-i-1} \\ &= \{0, 2(m+1)\delta, 2\delta, 2(m+2)\delta, 4\delta, \dots, 2(m+j)\delta, 2j\delta, \dots, 2(m+m)\delta, 2m\delta\} \% 2^{-i-1} \\ &= \{0, \delta, 2\delta, 3\delta, 4\delta, \dots, (2j-1)\delta, 2j\delta, \dots, (2m-1)\delta, 2m\delta\}, \end{aligned}$$

since $2(m+j)\delta \% 2^{-i-1} = (2m+1+2j-1)\delta \% 2^{-i-1} = (2j-1)\delta$, for $j = 1, \dots, m$. ◀

► **Lemma 7 (Restated).** *Consider any two points $p, q \in [0, 1]^d$, and let \mathcal{T} be the infinite quadtree of $[0, 2]^d$. For $D = 2 \lceil d/2 \rceil$ and $i = 0, \dots, D$, let $v_i = (i/(D+1), \dots, i/(D+1))$. Then there exists an $i \in \{0, \dots, D\}$, such that $p + v_i$ and $q + v_i$ are contained in a cell of \mathcal{T} with side length $\leq 2(D+1) \|p - q\|$.*

Proof. Without loss of generality, suppose d is even (as such $D = d$). Let $\ell \in \mathbb{N}$, such that for $\alpha = 2^{-\ell}$, we have

$$(d+1) \|p - q\| < \alpha \leq 2(d+1) \|p - q\|.$$

For $\tau \in [0, 1]$, let $\mathbf{G} + \tau$ denote the (infinite) grid with sidelength α shifted by the point (τ, \dots, τ) .

Let $X = \{i/(d+1) \mid i = 0, \dots, d\}$ be the set of shifts considered. Since we are shifting a grid with sidelength α , the shifting is periodical with value α . It is thus sufficient to consider the shifts modulo α .

Let $p = (p_1, \dots, p_d)$ and $q = (q_1, \dots, q_d)$. Assume that $p_1 \leq q_1$. A shift τ is bad, for the first coordinate, if there is an integer i , such that $p_1 \leq \tau + i\alpha \leq q_1$. The set of bad shifts in the interval $[0, \alpha]$ is

$$B_1 = \{(p_1, q_1) + i\alpha \mid i \in \mathbb{Z}\} \cap [0, \alpha].$$

The set B_1 is either an interval of length $|p_1 - q_1| \leq \|p - q\| < \alpha/(d+1)$, or two intervals (of the same total length) adjacent to 0 and α . In either case, B_1 can contain at most one point of $\alpha X = X \% \alpha$, since the distance between any two values of αX is at least $\alpha/(d+1)$, by Lemma 18.

Namely, the first coordinate rules out at most one candidate shift in $X \times \alpha$. Repeating the above argument for all d coordinates, we conclude that there is at least one shift in αX that is good for all coordinates. Let $\beta = \alpha i / (d + 1) \in \alpha X$ this be good shift. Namely, p and q belongs the same cell of $G + \beta$. The final step is to observe that shifting the points by $-\beta$, instead of the grid by distance β has the same effect (and $-\beta \times \alpha \in \alpha X$), and as such, the canonical cell containing both p and q is in the quadtree \mathcal{T} as desired, and the sidelength of this cell is α . ◀