Parallel Nearest Neighbors in Low Dimensions with Batch Updates

Guy E. Blelloch [*]   Magdalen Dobson [†]

Abstract

We present a set of parallel algorithms for computing exact k-nearest neighbors in low dimensions. Many k-nearest neighbor algorithms use either a kd-tree or the Morton ordering of the point set; our algorithms combine these approaches using a data structure we call the *zd-tree*. We show that this combination is both theoretically efficient under common assumptions, and fast in practice. For point sets of size n with bounded expansion constant and bounded ratio, the zd-tree can be built in $O(n)$ work with $O(n^\epsilon)$ span for constant $\epsilon < 1$, and searching for the k-nearest neighbors of a point takes expected $O(k \log k)$ time. We benchmark our k-nearest neighbor algorithms against existing parallel k-nearest neighbor algorithms, showing that our implementations are generally faster than the state of the art as well as achieving 75x speedup on 144 hyperthreads. Furthermore, the zd-tree supports parallel batch-dynamic insertions and deletions; to our knowledge, it is the first k-nearest neighbor data structure to support such updates. On point sets with bounded expansion constant and bounded ratio, a batch-dynamic update of size k requires $O(k \log n/k)$ work with $O(k^\epsilon + \text{polylog}(n))$ span.

1 Introduction

Computing nearest neighbors is one of the most fundamental problems in computer science, with applications in diverse areas ranging from graphics [19, 32, 33, 34] to AI [31] to as far afield as particle physics [36]. Research on nearest neighbors can be roughly divided into two areas: one area focuses on computing approximate nearest neighbors in high dimensions, primarily with clustering as an application. The second focuses on exact (or closer to exact) nearest neighbors in lower dimensions, with tasks such as surface reconstruction [3, 24] as a prominent application. This work focuses on the latter category.

The most common method of computing nearest neighbors in low dimensions is via a kd-tree [9], a tree which keeps the entire bounding box of the point set at its root, and whose children represent progressively smaller enclosed bounding boxes. Kd-trees have many applications in point-based graphics, and have been the data structure of choice for many graphics practitioners [33], even though other methods have better worst-case guarantees. One of the best kd-tree implementations is Arya et al's [8], which has been used widely by researchers [19, 32, 34]. Another commonly used library of kd-trees is that of the Computational Geometry Algorithms Library (CGAL) [40]. Recent work on kd-trees has focused on better theoretical guarantees [35], and with better performance in high dimensions [18].

Another approach for computing nearest neighbors uses space-filling curves known as the Morton ordering, z-ordering, or Lebesgue ordering (henceforth Morton ordering). Recursing by splitting the Morton ordering roughly splits space, making it possible to effectively search for nearest neighbors. Two nearest neighbor algorithms that make use of Morton ordering are Chan's minimalist nearest neighbor algorithm [16], and Connor and Kumar's k-nearest neighbor graph algorithm [22]. Other approaches to computing nearest neighbors include well-separated decompositions [15], and Delaunay triangulation [11].

Some important considerations when choosing a k-nearest neighbors algorithm are how it performs (theoretically as well as practically), does it run efficiently in parallel (since todays machines only have multiple processors), what kind of point sets it handles, and whether it supports dynamic updates (since in many applications point sets change over time [39]). Vaidya [41] and Callahan and Kosaraju [15] give strong bounds for general point sets computing all nearest neighbors in $O(n \log n)$ time using variants of kd-trees. Chan improved this to $O(n)$ time if the the ratio of the largest distance to the smallest is polynomially bounded [17]. However these results are limited to static point sets and have not yet shown to be practical. Connor and Kumar give bounds under the assumption of bounded expansion constants [22] for a practical algorithm they implement. There has also been significant interest in parallel algorithms for the problem. This includes implementations based on MapReduce [2], for GPUs [26], the STANN library [22], and an implementation in CGAL [4]. Although in principle kd-trees should be able to support dynamic updates we know of no libraries that efficiently support them, and few interesting theoretical bounds for the problem in low dimensions. When considering paral-

---
[*]Carnegie Mellon University.
[†]Carnegie Mellon University.

lelism and updates together one should be interested in batches of updates that can be processed in parallel.

In this paper, we present a technique that combines the ideas of kd-trees and Morton ordering to achieve efficient algorithms for k-nearest neighbors in bounded dimension. Some guiding intuition for such a combination is that Morton-based algorithms tend to have quick preprocessing (since only a sort is required) and slower queries; on the other hand, tree-based algorithms can have slower building times but their additional structure leads to faster queries. Thus, combining these approaches may allow us to achieve the advantages of both. In particular we present a k-nearest neighbor algorithm that hybridizes the kd-tree and Morton order approaches by using a kd-tree whose splitting rule is based on the Morton ordering; we call this tree the **zd-tree**. We also present what is to our knowledge the first parallel batch-dynamic update algorithm for a k-nearest neighbor data structure. We prove the following theoretical results in the context of a point sets with bounded expansion constant and bounded ratio, two reasonable and broadly used assumptions when computing nearest neighbors [5, 6, 8, 10, 17, 22, 25, 29, 30, 37].

The first result concerns the work and span required to build the zd-tree:

THEOREM 2.1. *For a point set P of size n with bounded ratio, the zd-tree can be built using $O(n)$ work with $O(n^\epsilon)$ span, and the resulting tree height $O(\log n)$.*

The second result bounds the work for a k-nearest neighbor query on the zd-tree.

THEOREM 2.2. *For a zd-tree representing a point set P of size n with bounded expansion, finding the k-nearest neighbors of a point $p \in P$ requires expected $O(k \log k)$ work.*

These two theorems together imply a linear-work algorithm for finding the k-nearest neighbors among a set of points (i.e. the k-nearest neighbor graph). They also imply that for a point $q \notin P$, finding the nearest neighbors requires $O(\log n + k \log k)$ work.

The third result bounds the work and span for batches of updates.

THEOREM 2.3. *Let T be a pruned zd-tree representing point set P, and let Q be a point set of size k, such that $|P| + |Q| = n$. Then if $P \cup Q$ and Q both have bounded expansion and bounded ratio in the same hypercube X, Q can be inserted into T in $O(k \log(n/k))$ work and $O(k^\epsilon + \mathrm{polylog}(n))$ span.*

In additional to the theoretical contributions, we implement both our nearest neighbor searching algorithm
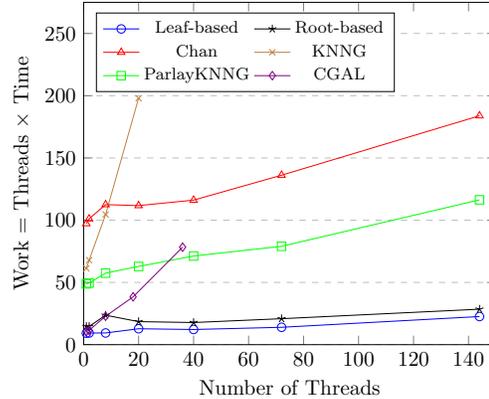
Figure 1: A figure showing the work (threads × time) performed by various nearest neighbor algorithms as the number of threads increases. The k-nearest neighbor graph was computed on 10 million points from a random dataset within a 3D cube. Ideally the line for a particular algorithm would be both low (small total work) and straight (indicating more threads does not change the total work).

and the batch-dynamic updates described above, and we measure our nearest neighbor searching algorithm against a large number of competitors. Our algorithms are optimized for parallelism: in addition to presenting a thread-safe data structure so that queries can be conducted in parallel, we use parallelism when recursively building or updating our kd-tree. A snapshot of our practical results can be found in Figure 1, which compares the work needed to preprocess and query a point set across our implementation and competitors.

Our experimental results show the following:

1. Our k-nearest neighbor algorithms achieve **high parallelism.** Using our basic algorithm to query all nearest neighbors of a 3D dataset with 10 million points achieves 75-fold speedup on a 72-core Dell R930 with 144 hyper-threads.

2. Our algorithms are **fast**. Our algorithm's speed is robust across all the measures which we tested—adversarial datasets, varying k, varying the size of the dataset, and varying the number of threads. In most cases, it beats its competitors by close to an order of magnitude.

3. Our batch-dynamic updates **drastically decrease the cost per insertion**. An insertion of one point into a tree of $5,000,000$ points takes about 10^{-5} seconds, while an insertion of $100,000$ points takes 10^{-7} seconds per pointelement.

1.1 Preliminaries. For the special case where we are given a point set P and wish to calculate the k-nearest neighbors of all points in P, we refer to the result as the **k-nearest neighbor graph** of P. A query of a

point not in P is a **dynamic query**, and a query of a point in the tree is sometimes referred to as as a **non-dynamic query**. Similarly, adding a point to or deleting a point from the tree is referred to as a **dynamic update**, which is **batch-dynamic** if the updates are processed in batches rather than one at a time.

Kd-trees. Many nearest neighbor algorithms use a kd-tree as the data structure to query for nearest neighbors. Given a set of points P in a d-dimensional bounding box, a kd-tree splits the data into two smaller bounding boxes at every level of the tree. Designing a kd-tree requires making a choice of **splitting rule**—that is, how the bounding box will be divided. One common splitting rule is to divide the bounding box of points along its largest dimension; another is to divide the space such that equal numbers of points are on each side. A variant of the kd-tree is the {*quad,oct*}-*tree*, where each internal node of the tree has 2^d equal sized children in d-dimensional space.

Morton ordering. A common tool used in designing k-nearest neighbor algorithms is the Morton ordering. For a set of points whose coordinates are d-vectors of integers $(x_1, x_2, \ldots, x_d)$, the Morton ordering is calculated by taking each integer coordinate in binary form, interleaving the coordinates to create one integer per point, then sorting using the interleave integers. Nearest neighbor algorithms take advantage of the following property of the Morton ordering: given any two points p and q and the rectangle defined by those points as the corners, then all points in the rectangle must fall between p and q in the Morton ordering. This allows pruning regions of the ordering.

Bounded expansion constant. Several previous works for nearest neighbors in metric spaces have assumed a bounded expansion constant [5, 6, 10, 22, 25, 29, 30, 37], which roughly requires that the density of points in the metric space does not change rapidly. In the context of Euclidean space we use the following definition. Given a point p_i and a positive real r, let box(p_i, r) denote the box centered at p_i with radius r (that is, half the side length).

DEFINITION 1. *Given a point set P contained in a bounded Euclidean space X, P has* **expansion constant** γ *if for all $x \in X$ and all positive real r, if $|box(x, r)| = k$ for any $k > 1$ then*

$$|box(x, 2r)| \leq \gamma k.$$

The expansion constant is referred to as **bounded** *if* $\gamma = O(1)$.

Bounded Ratio. Another property that will be needed to prove some of our theorems is that of the point set having bounded ratio. This is a commonly used property in problems such as nearest neighbors and closest pair [8, 17].

DEFINITION 2. *Given a point set P of size n, let $d_{\max}$ denote the maximum distance between any two points in the set, and let $d_{\min}$ denote the minimum distance between any two points in the set. Then P has* **bounded ratio** *if*

$$\frac{d_{\max}}{d_{\min}} = poly(n).$$

Model of computation. Our results for the parallel algorithms are given for the binary-fork-join model [12]. In this model, a process can fork two child processes, which work in parallel and when both complete, the parent process continues. Costs are measured in terms of the work (total number of instructions across all processes) and span or depth (longest dependence path among processes). Any algorithm in the binary forking model with W work and S span can be implemented on a CRCW PRAM with P processors in $O(W/P + S)$ time with high probability [7, 14], so the results here are also valid on the PRAM, maintaining work efficiency.

1.2 Related Work. Arya and Mount's k-nearest neighbor implementation [8] is commonly referred to as the state-of-the-art sequential k-nearest neighbor algorithm for low dimensions. Their implementation uses a kd-tree known as a *balanced box decomposition (BBD) tree*, whose splitting rule attempts to get the best of both commonly used splitting rules—that is, splitting a bounding box into approximately equal areas which also have approximately equal numbers of points. They show that using a BBD tree, an approximate k-nearest neighbor query takes $O(\frac{k}{\epsilon} \log n)$ work. The BBD tree theoretically supports insertions, but their given implementation does not. In [22], Connor and Kumar parallelize Arya and Mount's all-nearest neighbors implementation and show that theirs produces faster results, so we compare against Connor and Kumar's instead of theirs.

One nearest neighbor algorithm which uses the Morton ordering instead of a kd-tree is Chan's "minimalist" nearest neighbors algorithm [16], which has a theoretical guarantee of $O(n \log n)$ expected preprocessing time and $O(\frac{1}{\epsilon} \log n)$ expected time per query for approximate nearest neighbors. The algorithm is notable for both its simple proof and strikingly minimalist implementation, whose sequential version requires fewer than 100 lines of code in C++. Chan's algorithm first randomly shifts the coordinate of each point, then sorts the points using the Morton ordering. The algorithm then uses an implicit

tree, recursively dividing the sorted points and visiting every implicit vertex which is within some radius of the query point. An adversarial case for this algorithm is when some query point q is in the right half of the sorted data and its nearest neighbor p is in the left half, causing the algorithm to search a large number of vertices. The random shift helps avoid this case in expectation.

The k-nearest neighbor implementation that most closely matched ours—in that it is tailored for parallelism and for exact nearest neighbor searching in low dimensions—is that of Connor and Kumar in [22], where STANN stands for Simple Threaded Approximate Nearest Neighbor. Their algorithm makes several improvements on Chan's algorithm, especially for the case of computing the k-nearest neighbor graph. Their main improvement is to search from the leaf of the implicit tree rather than the root, which allows for the possibility of searching only $O(k)$ implicit nodes instead of at least $O(\log n)$ (and this would be a best case scenario where the tree is perfectly balanced). Indeed, they find that if the input point set has bounded expansion constant, their data structure uses $O(n \log n)$ work and their nearest neighbor queries use expected $O(k \log k)$ work. Their algorithm only works for static point sets and as our experiments show is not as fast as ours.

Another well-known tree used for computing nearest neighbors is Callahan and Kosaraju's well-separated pair decomposition [15]. For n points, they can build their tree (similar to a kd-tree) in $O(n \log n)$ work and polylogarithmic span (in parallel). Based on the tree, they can build the decomposition and find nearest neighbors in $O(n)$ work and polylogarithmic span. The approach is only described for the static case. Another approach is to use the Delaunay triangulation of the set of points [11]. Although this seems to work reasonably well in two dimensions, in three and higher dimensions it can be very expensive. Beyond bounded expansion constant, another common geometric assumption used for finding nearest neighbors or the closest pair is a bound on the ratio of the furthest pair and the closest pair in a dataset [8, 17, 20, 21, 23].

2 Algorithm Design and Bounds

Here we describe our algorithms for constructing a data structure for k-nearest neighbors, querying the structure, and batch updating it.

Data structure. The data structure we use for nearest neighbor searching is a kd-tree whose splitting rule uses the Morton ordering; this is what we refer to as the zd-tree. Since the Morton ordering is just the interleaving of the bits of each coordinate, the tree is built by letting the root represent the entire bounding box, and splitting the points into child nodes at level i

based on whether the bit at place i is 0 or 1. In three dimensions, our tree is almost equivalent to an oct-tree in which every three levels of our tree corresponds to one level of the oct-tree; however, the leaves can be at different levels. Each internal node of the tree stores the two opposing corners defining its bounding box, its two children, and its parent. Each leaf node stores its two opposing corners, its parent, and the set of points it contains. We bound the number of points in a leaf by a constant, and a leaf can be empty. Note that every point covered by the root bounding box is included in exactly one leaf node.

Construction. Before the zd-tree can be built, we preprocess the input. Firstly, motivated by Chan [16], and necessary for our bounds (the proof of Theorem 2.2), we select a random shift for each coordinate, and shift all the coordinates by this amount. This shift is kept throughout. We then sort the points by the Morton order. This can use Chan's comparison function, which leads to an $O(n \log n)$ work sort, but as we describe in Section 2.1 can be reduced to a linear time radix sort with span $O(n^\epsilon)$ [28] when assuming a bounded expansion constant. In this case the number of bits needed for the Morton order can bounded by $O(\log n)$.

After shifting and sorting we apply a divide-and-conquer algorithm (Algorithm 1) to build the zd-tree. The algorithm recurses at each level of the tree on the two sides of the cut for the given bit of the Morton ordering. Importantly, finding the cut in the routine `splitUsingBit` only requires a binary search since the points are sorted by Morton order. This implies that when the tree is sufficiently shallow (guaranteed by bounded ratio) the work to build the tree is only linear, and the parallel depth is low. Even if completely imbalanced the work would be $O(n \log n)$.

Downward search algorithm. Our downward search algorithm is detailed in Algorithm 2. The algorithm maintains a current set of k nearest neighbors, which starts empty and is improved over time by inserting closer points. In our pseudocode, we use N to represent the nearest neighbor candidate set. The downward search works as follows: let r be the distance from p to the furthest element in N if N contains at least k elements, or infinity otherwise. Now search vertex v only if the bounding box for v intersects a ball of radius r around p. This is determined by the `withinBox` function. If the node is a leaf, iterate through the points it contains and update the set of nearest neighbors if necessary. If it is not a leaf, recurse on its children, searching first the child whose center is closest to the query point p.

Our **root-based** algorithm simply starts at the root of the zd-tree with an empty N and applies `searchDown`, but we also use `searchDown` in our upward algorithm.

Upward search algorithm. Our upward search algorithm is detailed in Algorithm 3. It always starts at the leaf in the tree containing the point p and works its way up the tree. The idea is that in general, only a small part of the tree needs to be examined. It uses the downward search as a subroutine. As in the downward search, it maintains a priority queue N, initially empty, of the current estimate of the nearest k neighbors, which is improved over time helping to prune further search. The algorithm starts at the leaf by adding any points in the leaf to N. Then, as with the downard version, let r be the distance between p and k-th nearest neighbor in N, or infinity if there are not k neighbors in N yet. Now search the parent of the current node if and only if the ball of radius r around p extends outside the bounding box of the current node. Otherwise we know there are no points not included in the current node that could be closer than those in N. This can use the same `withinBox` as used in the downward algorithm, but with a negative r. When searching the parent we search the parent's other child using the downward algorithm.

Finding the leaf in which a point p belongs, which is needed, depends on whether we are generating a k nearest neighbor graph or using the the structure for dynamic searches for points not in the set. In the first case we know the leaf since each point is in a leaf. Therefore to generate a k-nearest neighbor graph we need just build the tree and then run `searchUp` on each point in each leaf. We refer to this as the **leaf-based** version of our algorithm. In the second case we have to search down the tree from the root to find the location of the leaf. This can use the bits of the Morton ordering to decide left or right. We refer to this as the **bit-based** version, and the downward search from the root as the **root-based** version.

Batch-dynamic updates. The tree data structure naturally lends itself to the possibility of dynamic insertions and deletions. Insertion of a new point q into a zd-tree T is conceptually simple: locate the leaf of T which q should be inserted into; then either add q to the sequence of points contained in the leaf, or if q would cause the number of points in the leaf to exceed some cutoff, split the leaf into two children. This concept can be refined to a parallel batch-dynamic algorithm, which takes a set of points and recurses in parallel down the right and left children of the root. One small subtlety is that to avoid cases where an insertion might require a rebuild of the entire tree, we require a bounding box that all data will be contained in to be specified before building the initial tree.

As with building the tree from scratch, a batch-dynamic insert starts by using the random shift to offset the points and sorting the points to be added based

---

Algorithm 1: buildTree(P, b)

Input: A set of randomly shifted points sorted according to their Morton ordering and an integer b representing the bit we are working on, starting with the highest bit.
Output: The leaf or internal node that contains P's bounding box

1 **if** $b == 0$ or size(P) $<$ sizeCutoff **then**
2   **return** createLeaf(P)
3 **else**
4 $i =$ splitUsingBit(P, b) ;
5   **do in parallel**
6 $L =$ buildTree($P[1 : i], b - 1$) ;
7 $R =$ buildTree($P[i : n], b - 1$) ;
8   **return** createInternalNode(L, R) ;

---

on their Morton ordering. We then apply the recursive algorithm shown in Algorithm 4. Deletions use an almost identical algorithm.

2.1 Theoretical Results. In this section, we give theoretical results on the performance of our algorithms when assuming bounded expansion constant. The results in the rest of the section assume that the point set P has bounded expansion constant $\gamma \geq 2$ as well as bounded ratio, and they assume that the dimension $d = O(1)$. We also require that every coordinate of every point is unique. This is a fair assumption to make in the context of nearest neighbors, since every point set that does not have this property can be transformed into one that does and where each point retains the same nearest neighbors. The presented proofs in this section assume that X is a bounding cube, but the full version of the paper contains the proofs of the same results where X is any convex region. Wherever not otherwise specified, we use B to denote the bounding box of the randomly shifted point set; note that the side lengths of B can be at most twice the side lengths of the smallest bounding box containing X.

Our first theorem concerns the height and build time of a zd-tree on a point set with bounded ratio.

THEOREM 2.1. *For a point set P of size n with bounded ratio, the zd-tree can be built using $O(n)$ work with $O(n^\epsilon)$ span, and the resulting tree height $O(\log n)$.*

Proof. For the tree depth and work bound, we need to show that the longest path in the tree has length $O(\log n)$. The bounding cube of X has side length within a constant factor of $d_{\max}$, and it must be divided until the two points whose distance between them is $d_{\min}$ are in separate cubes. Since $d_{\max}/d_{\min} = \text{poly}(n)$, $d_{\max}$

(a) A graphic illustrating a box situated inside the largest possible quad-tree box as referenced in Lemma 2.2.

(b) A figure illustrating the setup in Lemma 2.3. Here, the empty boxes represent empty splits, the light blue boxes represent unbalanced splits, and the dark blue box represents the box containing p.

(c) An illustration of the three types of cuts referenced in Lemma 2.4: interior in red, exterior in black, and empty in blue.

Figure 2: Figures aiding with the proofs in Section 2.1.

must be halved $O(\log n)$ times to reach $d_{\min}$. Thus the tree has $O(\log n)$ depth.

For the sorting claim, we wish to show that a parallel radix sort can be used. For the radix sort to require $O(n)$ work, we need to guarantee that only $O(\log n)$ bits are required to sort the dataset. This follows directly from the fact that the tree has depth $O(\log n)$. □

Now, we work towards the theorem on the expected time required to query the k-nearest neighbors of a point p. Like the result from Connor and Kumar in [22], the $O(k \log k)$ bound only applies when searching for the nearest neighbors of a point already in the tree; for a dynamic query, finding the leaf to start from requires $O(\log n)$ time. The following proofs assume without loss of generality that a leaf of the zd-tree contains no more than k points. Furthermore, we assume that the tree is a true quad-tree, meaning that each internal node has 2^d children; thus we use the term "quad-tree box" to refer to the box belonging to a tree node. This will slightly simplify the analysis, and it can only be worse than the performance of the actual algorithm.

THEOREM 2.2. *For a zd-tree representing a point set P of size n with bounded expansion, finding the k-nearest neighbors of a point $p \in P$ requires expected $O(k \log k)$ work.*

The proof of the theorem separates the work into two parts: the work of searching through the points in a leaf of the zd-tree, and the work of traversing the zd-tree to get to those leaves. The first lemma concerns the former.

LEMMA 2.1. *When searching for the k-nearest neighbors of a point p, $O(k)$ candidate points will be considered,* resulting in $O(k \log k)$ work to evaluate all the candidate points.

Proof. The multiplicative factor of $\log k$ comes from the fact that a priority queue is used to store nearest neighbors, so an $O(\log k)$ cost is incurred each time the priority queue is updated.

Consider the leaf L of the tree that would contain p. The initial approximation is found by recursing up to L's ancestor until the ancestor has more than k descendants, then adding those descendants to the priority queue. The ancestor's bounding box B must contain $O(k)$ points by the fact of bounded expansion constant, since one of its children contains fewer than k points.

Now, let r be the side length of B. Our algorithm will search a leaf only if the box belonging to that leaf overlaps with box(p, r). Those neighbors are contained within radius r of the quad-tree box containing our initial guess.

The search algorithm evaluates every point at radius r from p as a candidate point. If the radius r of B is expanded twice, all the candidate points will be contained in the resultant box; call this box of candidates Q. If $k' = O(k)$ points are in B, the expansion condition guarantees that at most $\gamma^2 k' = O(k)$ points are in Q. However, the search algorithm does not directly search each point in Q; rather, it searches every leaf whose bounding box overlaps with Q. This means that if a leaf L's bounding box were to fall partially inside Q and partially outside, all the points in L would be counted. Since B is a quad-tree box, this could only occur if the leaf L were to have a bounding box with radius larger than r. Since the radius is larger than r and $d = O(1)$, there can only be a constant number of such leaves. Since

Algorithm 2: searchDown(T, p, N)
needed subroutines:
distance(p, N, k) returns infinity if N has fewer than k points and otherwise returns the distance from p to the furthest point in N; $k = 1$ if not specified.
insert(N, p, k) adds p to the set N keeping only the k closest points to p.
withinBox(T, p, r) returns true if p is within a distance r of the bounding box for T.

Input: A pointer to a tree node T, the query point p, and a current set of up to k nearest neighbors N.
Output: The k-nearest neighbors of p.

1   $r \leftarrow$ distance(p, N, k) ;
2   **if** withinBox(T, p, r) **then**
3   **if** $T =$ Leaf **then**
4       $Q \leftarrow$ set of points contained in T ;
5   **for** $q \in Q$ **do**
6   **if** $q \neq p$ **then**
7   **if** distance(q, p) < distance(p, N, k) **then**
8   insert(N, p, k);
9   **else**
10       $R \leftarrow T$.Right() ;
11       $L \leftarrow T$.Left() ;
12       $\ell \leftarrow$ distance(p, L.center());
13       $r \leftarrow$ distance(p, R.center()) ;
14   **if** $\ell < r$ **then**
15   N ' = searchDown(L, p, N);
16   **return** searchDown(R, p, N');
17   **else**
18   N' = searchDown(R, p, N);
19   **return** searchDown(L, p, N');

Algorithm 3: searchUp(C, p)
withinBox(T, p, r) with negative r returns true if p is in within the bounding box of T and at least r from the boundary.

Input: A leaf C of the kd-tree and a point p within the bounding box of C.

1   $N = \emptyset$
2   $Q \leftarrow$ set of points contained in C ;
3   **for** $q \in Q$ **do**
4   **if** $q \neq p$ **then**
5   **if** $d(q, p) <$ distance(p, N, k) **then**
6   insert(N, p, k);
7   $r \leftarrow$ distance(p, N, k);
8   $P \leftarrow C$.Parent() ;
9   **while** not withinBox($C, p, -r$) and $P \neq \top$ **do**
10   **if** P.Left() $= C$ **then**
11       $N =$ searchDown(P.Right(), p, N);
12   **else**
13       $N =$ searchDown(P.Left(), p, N);
14     $C = P$;
15     $r \leftarrow$ distance(p, N, k);
16     $P \leftarrow C$.Parent() ;
17 **return** N

Algorithm 4: batchInsert(T, P)

Input: A pointer to a node T of the kd-tree and a set of points P contained in its bounding box and sorted according to their Morton order

1   **if** $T =$ Leaf **then**
2   **if** size(P) + size(T) ¡ leafCutoff **then**
3   Insert $p \in P$ into T ;
4   **else**
5   Split T into multiple leaf nodes ;
6   **else**
7     $b = T \rightarrow$ bit ;
8     $i =$ splitUsingBit(P, b) ;
9   **do in parallel**
10   batchInsert($T \rightarrow$ Right, $P[1 : i]$) ;
11   batchInsert($T \rightarrow$ Left, $P[i : n]$) ;

---

each leaf has at most k points, the original bound still holds. □

Now, we move on to considering the expected work required to traverse the zd-tree to access the leaves. The probability calculation given here can also be found in [22].

LEMMA 2.2. *When traversing the zd-tree, the expected number of tree edges traversed to find all the nearest neighbors is $O(k)$.*

Proof. The worst case for the cost of the search algorithm is when the search space described in Lemma 2.1 is only contained in the bounding box containing all of P, as illustrated in Figure 2(a). First, we show that the length of the traversal is bounded by the longest path between two leaves in the search space; all other searches can be charged to the $O(k)$ points that are searched. We will

use the box B to refer to the search space. The largest cut of B divides B into 2^d quad-tree boxes. Each of these boxes must be contained within a quad-tree box Q_i of at most twice the side length of B. Thus, traversing every leaf in Q_i incurs cost at most $O(k)$; since d is constant, the claim follows.

All that remains is compute the expectation of the length of the longest path in the zd-tree. Without loss of generality, assume that the search space is a box with side length 2^h. Then the probability that the search space is contained within a box of side length 2^{h+j} is $\left(\frac{2^j - 1}{2^j}\right)$, since the box must have its upper left corner in

one of $2^j - 1$ grid squares along each dimension. Thus, due to the random shift, the probability that the search space is NOT contained within a box of side length 2^{h+j} is $1 - \left(1 - \frac{1}{2^j}\right)^d$. From the perspective of traversing the zd-tree, this is the event that the path between two leaves in the search space is length j. Thus its expectation can be upper bounded by the following summation, which charges a cost of one for each box the search space is not contained in:

$$\sum_{j=1}^{\infty} \left(1 - \left(1 - \frac{1}{2^j}\right)^d\right) = O(1)$$

and the result follows.   □

Lemmas 2.1 and 2.2 together compose the proof of Theorem 2.2.

Now we move on to batch-dynamic updates. In a weight-balanced tree, the argument for the desired $O(k \log(n/k))$ bound would be as follows: when a batch of points is inserted into the tree, the work required to insert them into $\log k$ levels can be no more than the work that would be required to build them into their own tree. Thus insertion into the first $O(\log k)$ levels of the tree uses $O(k)$ work, and the work bound on an insertion is $O(k(\log n - \log k)) = O(\log(n/k))$.

Hence the goal of Theorem 2.3 is to show that in addition to the traditional notion of balance, the zd-tree also obeys some notion of weight balance—that is, that each split of a point set must produce two halves where each contains a constant fraction of the points. Unfortunately, this is not strictly true, since a set of points with bounded expansion may, for example, have only one element with a 1 at the largest bit if that element is very close to the rest of the elements. However, we will be able to show a slightly weaker notion than weight balance: that enough nodes in the tree are weight-balanced that the same work bound still applies.

One more piece of terminology is needed before the theorem statement. When building the zd-tree by successively splitting the bounding box of the input, a split may have no points in X on one side. During the tree building phase, we face a choice regarding empty cuts: when the algorithm makes an empty cut, we could either fork off two child nodes where one is a leaf containing no points, or we could simply not fork off an empty node, and divide the other node using the next bit. Since it is strictly cheaper, we choose the latter. Thus the following analysis will not deal with empty cuts; in particular, Lemmas 2.4 and 2.5 do not consider empty cuts in their analysis, since empty cuts do not affect the length of paths in the tree.

THEOREM 2.3. *Let T be a pruned zd-tree representing point set P, and let Q be a point set of size k, such that $|P| + |Q| = n$. Then if $P \cup Q$ and Q both have bounded expansion and bounded ratio in the same hypercube X, Q can be inserted into T in $O(k \log(n/k))$ work and $O(k^\epsilon + \text{polylog}(n))$ span.*

When X in its bounding box B is being recursively divided using Algorithm 1, it will be useful to separate the divisions or cuts into several categories. A cut along dimension d divides some sub-cube S of B in two with cutting plane ℓ. The cut is either **empty**, meaning that ℓ does not touch any points in X; or it is **exterior**, meaning it touches the boundary of X; or it is **interior**, meaning that within S, ℓ does not touch any points on the boundary of X. See Figure 2(c) for an illustration.

The first step towards Theorem 2.3 is to show that interior cuts are weight balanced.

LEMMA 2.3. *One out of every d interior cuts must be weight balanced; that is, it must split its bounding box into sets of size αn and $(\alpha - 1)n$ for constant $\alpha \in (0, 1)$.*

Proof. Consider a point $p \in P$. As the zd-tree is built, the point set P is split into smaller pieces along each dimension. Call a split **unbalanced** if it splits P into pieces of size n_1, n_2 such that one of $n_1, n_2 < \frac{1}{2(1+\gamma)}n$. Refer to a split as "involving" p if it splits a box containing p. We will show that after $d - 1$ unbalanced splits involving p, the next split must be balanced.

Assume for contradiction that there are d consecutive unbalanced splits involving p. Let B be the quadtree box containing p after those splits, and let the length of B's longest side be 2^j. By the assumption that d consecutive unbalanced splits have already happened, there must be some side of B where the most recent split on that side was of a region with maximum side length 2^{j+1}; call the hyperrectangle resulting from that split Q. Let $q \in B$ be the unique closest point to Q, as shown in Figure 2(b). Then, consider any $x \in Q$ such that $B_q = \text{box}(x, 2^j)$ contains q and no other point in B, and overlaps only B and Q. Due to its proximity to B, $\text{box}(x, 2^{j+1})$ must completely contain B. By our assumption, B_q contains fewer than $\frac{1}{2(\gamma+1)}n + 1$ points, since it contains one point from B and otherwise only points from Q, which has at most $\frac{1}{2(\gamma+1)n} + 1$ points by our assumption. The box B contains at least $\left(1 - \frac{1}{2(1+\gamma)}\right)^d n$ points, so since $d \geq 2$, the expansion constant is violated and we reach a contradiction.   □

While interior cuts are easily shown to be balanced, the same argument does not hold for a sequence of exterior cuts, since Lemma 2.3 relies on being able to

choose a certain point x as the center of a box, and this point might not be included in X if some of the d unbalanced cuts were exterior. Since an arbitrarily long sequence of nodes formed from exterior cuts might not be weight-balanced, we take a different approach: showing that even if we have to pay the maximum possible cost for each unbalanced path, the number of points on such an unbalanced path is small enough that the overall bound is unchanged.

The first step towards this goal is to quantify, for a given point $p \in P$, how many exterior cuts involving p must be made before the first interior cut involving p.

LEMMA 2.4. *Normalize the length of B to n. Then for every point $p \in P$, let $f(p)$ denote the minimum distance from p to the boundary of X, perpendicular to some side of the bounding box B. Then $O\left(\log \frac{n}{f(p)}\right)$ cuts will be made before the next cut containing p is interior.*

Proof. A cut involving p is guaranteed to be interior when the quad-tree box containing p has radius less than $f(p)$. Since the radius of the quad-tree box is halved every d cuts and side length of B is n, the aforementioned condition is met after $d \cdot \log(n/f(p))$ cuts. □

Lemma 2.4 gives us a way to bound the number of exterior cuts along the path to p. The next step is to bound the *number* of points that can be a given distance or closer to the boundary of one of the faces.

LEMMA 2.5. *Normalize the side of B to n. Let S be a subset of B formed by cutting B parallel to one of its faces at distance r away from the face. Then at most a $\left(\frac{\gamma^2}{\gamma^2+1}\right)^{\log n/r}$ fraction of the total points in B are contained in S.*

Proof. One way of upper bounding the number of points in S is as follows: S can be formed by dividing the bounding box in half along one dimension $\log(n/r)$ times. On each division, at most how many points can be in the resultant rectangle? Consider the first cut which divides the bounding cube in half; the goal is to maximize the number of points in one half without violating the expansion constant. The "sparse" half can be separated into 2^{d-1} hypercubes of radius $n/2$, each of which is expanded twice before the whole space X is encompassed. Note that it is optimal for all the "sparse" sub-cubes to contain the same amount of points since each must be able to expand a box around its center. The following equation solves for the largest fraction f of the total points that can be in any one sub-cube without violating

the expansion constant.

$$\gamma^2 \frac{1-f}{2^{d-1}-1} \geq x + (2^{d-1}-2)\frac{1-f}{2^{d-1}-1}$$
$$\implies f \leq \frac{\gamma^2 - 2^{d-1} + 2}{\gamma^2 + 1}.$$

The following cuts need to be upper bounded in a slightly different way, since they are cutting a hyperrectangle with $d-1$ sides of length n and one side of length s. Thus the area can be decomposed into $2n/s$ boxes of side length $s/2$. Half the boxes will receive the maximum number of points possible; both halves will evenly distribute their points across the boxes. Each "sparse" hypercube of side length $s/2$ must be expanded twice before it contains a cube C of side length s that is completely contained within the region. For that cube C, the same equations as before can be written to bound the number of points in the dense half of C to a $\frac{\gamma^2 - 2^{d-1}_2}{\gamma^2+1}$ fraction of the total points in C. The total number of points in C is upper bounded by $\left(\frac{\gamma^2 - 2^{d-1}_2}{\gamma^2+1}\right)^k$ where k is the total number of cuts that have occurred, assuming that the maximum possible fraction is on one side each time. The overall bound follows. □

Now, we can put all these pieces together to prove the theorem.

Proof. [Proof of Theorem 2.3]
The sorting cost bound follows directly from Theorem 2.1.

For the update bound, we know that over all weight-balanced paths in the tree, the cost of insertion the k points down those paths is $O(k \log(n/k))$. Thus our task is to account for the paths in the tree that are not weight-balanced. In the worst case, for every non-weight-balanced path of length ℓ, we incur an $O(\ell)$ cost for each point that traverses it. We will show that the number of points in Q that are distributed among the unbalanced parts of the tree is small enough that the overall bound is unchanged. Consider one face S of X. For a point p at a given length r away from the boundary of x perpendicular to S, the path traveled from the root to p can encounter $O(\log(n/r))$ unbalanced nodes. Consider all points at a distance r or closer to S. Lemma 2.5 shows that there are at most $k \cdot \left(\frac{\gamma^2}{\gamma^2+1}\right)^{\log(n/r)}$ such points. The cost incurred for each depth is also $\log(n/r)$. Thus, the maximum cost for all the points at depth $f(n)$ or smaller is $k\left(\frac{\gamma^2}{\gamma^2+1}\right)^{\log(n/r)} \log(n/r)$. Renaming $\log(n/r)$ as the variable x and integrating over values of

x from 0 to $n/2$ gives:

$$k \int_0^{n/2} \left(\frac{\gamma^2}{\gamma^2 + 1} \right)^x x \, dx$$
$$< k \int_0^{\infty} \left(\frac{\gamma^2}{\gamma^2 + 1} \right)^x x \, dx$$
$$= O(k).$$

Since the number of faces is constant, the overall bound follows. This shows that even in the worst case where every exterior cut is unbalanced and the maximum number of points are distributed in nodes formed from exterior cuts, the extra work incurred does not change the overall bound. □

We conclude the section with a note on the tradeoff between work and span. The versions of the theorems given in this section have a span that is greater than polylogarithmic due to the radix sort. If the algorithms used a comparison sort, they could run in polylogarithmic span at the cost of needing $O(n \log n)$ work for the sort.

3 Implementation Details

In this section we give more details on the practical implementation of both our algorithms and the other algorithms we use to benchmark our code.

3.1 Our Algorithms. We implemented our algorithms in C++ using the parallel primitives from ParlayLib [13]. Our search implementation closely matches the algorithms shown in Section 2, so here we focus mostly on implementation details and other optimizations.

Numerical details. We work with double-precision floats, which we round to 64-bit integers for building the tree and computing the Morton ordering.

Miscellaneous optimizations. We mention a few optimizations that made significant differences in our runtime. Whenever possible, we used squared distances instead of Euclidian distances in our computations, which made our code about 10% faster. To store the current k-nearest neighbors when traversing the tree we use a vector for small k and the C++ STL priority queue for larger k. The overhead of the vector was significantly less for small k, but the linear instead of logarithmic cost dominates for $k > 40$ or so. A third optimization was to sort the sequence of queries using their Morton ordering so that nearby queries in this order access nearby nodes in the tree, thus reducing cache misses. The savings from reducing cache misses more than compensates for the cost of the sort, in some cases decreasing runtime by a factor of two. This is useful even when querying in

parallel since the parallel scheduler processes chunks of the iteration space on the same core.

3.2 Other Implementations. For the purpose of comparison we use three existing implementations of nearest neighbor search: CGAL [4], STANN [22], and Chan [16]. Here we describe some performance issues with their code, and some modifications we made to improve the performance of their code to ensure a fair comparison. An extended version of this discussion can be found in the full version of this paper.

Chan. Chan's code was fully sequential so we needed to parallelize it. Conceptually this is relatively straightforward since the algorithm just requires using a parallel sort instead of a sequential one, and then running the queries in parallel. Chan's code only searches from the root of his implicit tree. We note that the root-based implementation of our code is significantly faster than Chan's.

STANN. STANN includes both a k-nearest neighbor graph (KNNG) function and a k-nearest neighbor (KNN) function. The first finds the k nearest neighbors among a set of points, and the second supports a function to build a tree and a separate function to query a point for its k nearest neighbors. They supply a parallel version of KNNG, that was parallelized with OpenMP, and only a sequential version of KNN. Their algorithm did not scale well beyond 16 threads, since it left some components sequential. We therefore updated their code to use the parallel primitives and built-in functions from ParlayLib [13]; this drastically improved their performance.

CGAL. CGAL implements a parallel version of their k-nearest neighbor code using the threading building blocks (TBB) [27]. We use their code directly with no modifications. We note that their code does not scale well past 16 or so threads. Furthermore, although the code appears to be thread safe, there seems to be contention when there are many threads, thereby slowing them all down. Due to the particularly bad performance beyond 36 threads (which all are on one chip), we only report numbers up to 36 threads. Furthermore, since we observed wildly varying times with higher k, we only included times for $k < 10$ in our experiments.

4 Experiments

In this section, we provide experimental results which show that 1) our algorithms perform well under many types of scaling and across different architectures, and 2) our algorithms outperform every implementation we test against.

4.1 Experimental Setup.
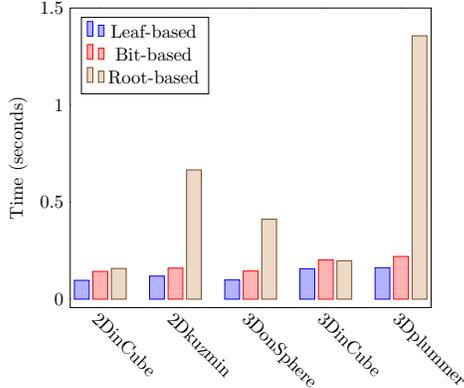
Figure 3: A bar chart showing the performance of our three k-nearest neighbor algorithms on different datasets. All datasets are 10 million points, and times reported are for $k = 1$.

Machines. We ran most all of our experiments on a 72-core Dell R930 with 4x Intel(Ⓡ) Xeon(Ⓡ) E7-8867 v4 (18 cores, 2.4GHz and 45MB L3 cache), and 1Tbyte memory. With hyperthreading the total number of threads is 144. To check whether the results were robust across machines, we also ran one set of experiments on a 4-socket AMD machine with 32 physical cores in total, each running at 2.4 GHz, with 2-way hyperthreading, a 6MB L3 cache per socket, and 200 GB of main memory.

Test data. After testing our algorithms on some real-world image data, we discovered, similarly to Connor and Kumar [22], that uniformly random points perform very similarly to real-world datasets. To facilitate testing at various sizes we therefore use a few distributions of random points in 2 and 3 dimensions, on sizes up to 100 million points. The point sets we used are listed in Figure 3. The 2D and 3DinCube datasets are points picked uniformly at random in a square and cube, respectively. The points in the 3DonSphere dataset are selected on the 2D surface of a sphere in 3 space. This is meant to represent various graphics applications where the point sets are on a 2D surface embedded in 3D. The 3Dplummer distribution uses the Plummer model [1], which is based on the study of galaxies, and highly dense in at the center, becoming very sparse on the outside. The 2Dkuzmin distribution is a similarly skewed distribution in two dimensions.

As can be seen from the various statistics, the Plummer and Kuzmin distributions are significantly more skewed than the others. Indeed. these distributions do not have bounded expansion constant. The performance is slower due to the fact that a few small points are extremely far away from most of the points, which are in a dense cluster at the center. This causes the tree to be unbalanced and the searches for the nearest neighbors of these far points to be expensive. However the overall time is hardly affected by the skewed distribution

for our leaf and bit-based implementations (Figure 3) showing the algorithms are robust under quite skewed distributions. As expected, the depth of the trees for the uniform distributions in a cube are just the logarithm of the number of points.

Algorithms tested. We ran three classes of experiments: (1) generating a k-nearest neighbor graph on a set of n points, (2) building a k-nearest neighbor query structure on n points followed by dynamic queries on a different set of n points, and (3) batch insertions for a total of n points (after the insertion). The results from (2) can be found in the full version of our paper; they are not significantly different from building the k-nearest neighbor graph.

Altogether we tested 9 variants of the algorithms: our parallelized version of Chan's algorithm, the CGAL algorithm, four variants of STANN (KNN, KNNG, parlayKNN and parlayKNNG), and three variants of our algorithm (leaf-based, root-based, and bit-based). The parlayKNN and parlayKNNG are our modified versions of Connor and Kumar's algorithms. Since our modified versions are always significantly faster, we only report numbers for our versions.

For all implementations, we sort by Morton order before querying. This is to ensure that all algorithms are getting the same benefit of locality in the tree when querying. The experiments on batch insertion only use our algorithm since the others do not support dynamic updates.

4.2 Leaf vs. Root Based. In Figure 3, we show the performance of our three search algorithms for finding the k-nearest neighbor graph on varying datasets and $k = 1$. Figure 3 shows the same result using a different measurement: the average and maximum number of nodes visited during a query. One takeaway from the figures is that even though the leaf-based method takes $O(n)$ work, and the bit-based method takes $O(n \log n)$ work, the prior is only slightly faster. This is because the constant in the $O(k \log k)$ term for a search from the leaf is much larger than the constant in the $O(\log n)$ search from the root. Another takeaway is that for the Kuzmin and Plummer distributions, when starting from the leaf using the root-based algorithm (Algorithm 2) makes an enormous difference.

4.3 k-Nearest Neighbor Graphs. The results of our experiments for generating the k-nearest neighbor graph can be found in Figure 4.

Varying dataset size. (Figure 4(a) and (b).) We measured the total time per point (that is, to build the tree and perform the query) by dividing the total time to build and search by the number of points. As discussed
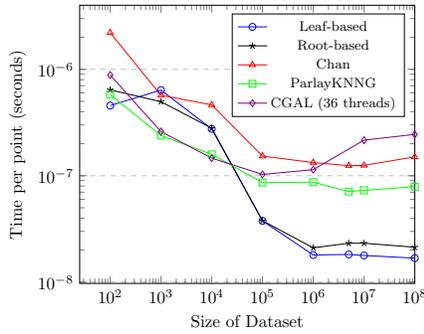
(a) Time required to calculate nearest neighbors as the size of the dataset increases. Calculated by dividing the total time by the number of points queried.
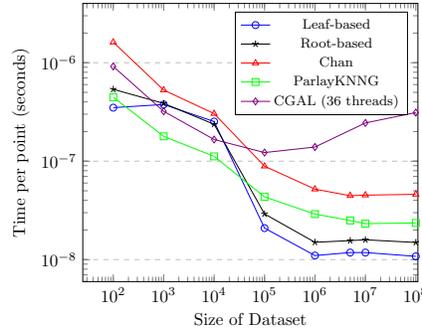
**(b)** The same as (a) but with a 2D dataset drawn randomly from a square instead of a 3D dataset.
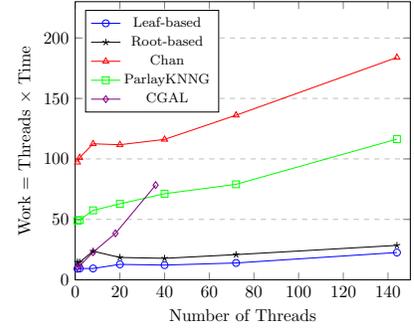
(c) Total work (threads × time) required to build a tree of 10 million points, then build the nearest neighbor graph of the point set. Shown as the number of threads vary.
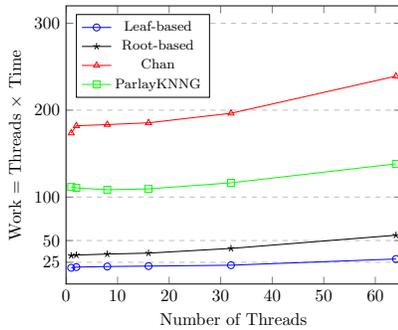
(d) The same measurements as (c), but on the 32-core AMD machine.
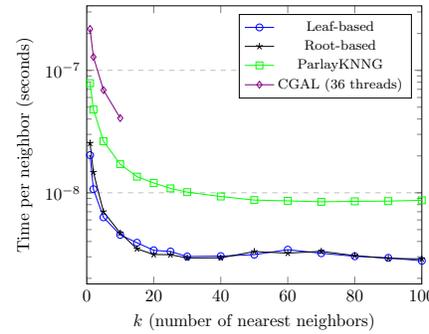
(e) Time required to calculate a neighbor as the number of neighbors k increases. Calculated by dividing the total time by k times the number of queries.
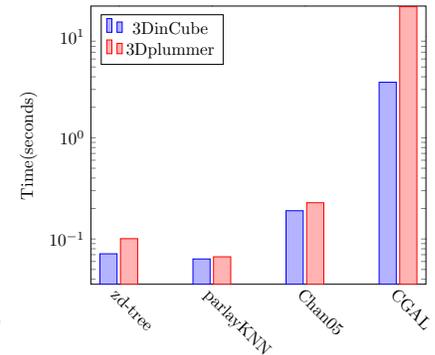
(f) The bars represent the total time each algorithm takes to build the data structure, for points drawn randomly from a 3D cube, and points drawn from a Plummer distribution.

Figure 4: Statistics related to non-dynamic queries. Unless otherwise stated, the size of the dataset is 10 million, the number of nearest neighbors $k = 1$, experiments were performed on 144 threads on a 72-core Dell R930, and data points are drawn randomly from a 3D cube.

in Section 3, we took measures to limit the number of cache misses where possible.

Work efficiency. (Figure 4(c) and (d).) Experiments showed that our algorithms performed significantly less work than our competitors as the number of threads increased to 144. To show that we maintain work efficiency on different architectures, we also ran the same experiments on a 32-core AMD machine.

Varying k. (Figure 4(e).) Our results on varying numbers of neighbors show that our algorithms remain fast and scalable.

Tree building. (Figure 4(f).) To illustrate that the tree-building step itself is efficient (except in the case of CGAL as explained in Section 3.2), we show time required to build the data structure for the 3DinCube and 3Dplummer distributions.

4.4 Dynamic Updates. We test the efficiency of our batch-dynamic updates by measuring the time required per update as the number of updates in the batch increases. Figure 5 shows the time taken per point as the size of the batch increases, with both insertions and deletions shown. The figure shows a drastic change in time, spanning almost four orders of magnitude from 10^{-4} seconds for a single update to 10^{-8} seconds per update for a batch of 5 million. The first period of decrease as the batch size increases to 10^4 or 10^5 can be explained by parallelism—this is the point at which the parallel sort and the parallel recursion down the tree begin to save significant time. The fact that time continues to decrease even after the size grows large enough to see the full effects of parallelism can be attributed to the work efficiency of the batch-dynamic updates, as shown in Theorem 2.3.
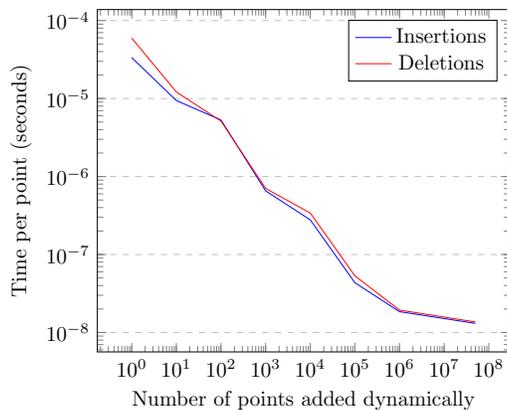
Figure 5: Time required per point as the number of points in the batch increases. Updates performed on a dataset of 10 million random points inside a 2D cube.

4.5 Code availability. Our implementation is part of the publicly available Problem-Based Benchmark Suite [38].

5 Conclusion

In this work, we presented the zd-tree, a data structure for k-nearest neighbors that combines the ideas of kd-trees and Morton ordering and supports batch-dynamic updates. We showed that the zd-tree is both theoretically efficient and fast in practice, performing well even on data sets which do not have bounded ratio or bounded expansion. One future experimental direction is to experiment with datasets with higher dimensions, or with using our algorithms as a sub-step in calculating nearest neighbors in high dimensions. Another direction worth exploring is to use the zd-tree or a similar data structure for other problems in low-dimensional geometry, such as closest pair or n-body interactions.

Acknowledgments

References

[1] Aarseth, S.J., Henon, M., & Wielen, R. 1974. Numerical methods for the study of star cluster dynamics. *Astronomy and Astrophysics*, **37**(2), 183–187.

[2] Agarwal, Pankaj K., Fox, Kyle, Munagala, Kamesh, & Nath, Abhinandan. 2016. Parallel Algorithms for Constructing Range and Nearest-Neighbor Searching Data Structures. *Pages 429–440 of: Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS 2016, San Francisco, CA, USA, June 26 - July 01, 2016.* ACM.

[3] Alexa, Marc, Behr, Johannes, Cohen-Or, Daniel, Fleishman, Shachar, Levin, David, & Silva, Cláudio T. 2001. Point Set Surfaces. *Pages 21–28 of: 12th IEEE Visualization Conference, IEEE Vis 2001, San Diego, CA, USA, October 24-26, 2001, Proceedings.* IEEE Computer Society.

[4] Alliez, Pierre, & Fabri, Andreas. 2016. CGAL: the computational geometry algorithms library. *Pages 8:1– 8:8 of: Special Interest Group on Computer Graphics and Interactive Techniques Conference, SIGGRAPH '16, Anaheim, CA, USA, July 24-28, 2016, Courses.* ACM.

[5] Anagnostopoulos, Evangelos, Emiris, Ioannis Z., & Psarros, Ioannis. 2015. Low-Quality Dimension Reduction and High-Dimensional Approximate Nearest Neighbor. *Pages 436–450 of: 31st International Symposium on Computational Geometry, SoCG 2015, June 22-25, 2015, Eindhoven, The Netherlands.* LIPIcs, vol. 34. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

[6] Anagnostopoulos, Evangelos, Emiris, Ioannis Z., & Psarros, Ioannis. 2018. Randomized Embeddings with Slack and High-Dimensional Approximate Nearest Neighbor. *ACM Trans. Algorithms*, **14**(2), 18:1–18:21.

[7] Arora, N. S., Blumofe, R. D., & Plaxton, C. G. 2001. Thread Scheduling for Multiprogrammed Multiprocessors. *Theory of Computing Systems (TOCS)*, **34**(2).

[8] Arya, Sunil, Mount, David M., Netanyahu, Nathan S., Silverman, Ruth, & Wu, Angela Y. 1994. An Optimal Algorithm for Approximate Nearest Neighbor Searching. *Pages 573–582 of: Proceedings of the Fifth Annual ACM-SIAM Symposium on Discrete Algorithms. 23-25 January 1994, Arlington, Virginia, USA.* ACM/SIAM.

[9] Bentley, Jon Louis. 1975. Multidimensional binary search trees used for associative searching. *Commun. ACM*, **18**(9).

[10] Beygelzimer, Alina, Kakade, Sham M., & Langford, John. 2006. Cover trees for nearest neighbor. *Pages 97–104 of:* Cohen, William W., & Moore, Andrew W. (eds), *Machine Learning, Proceedings of the Twenty-Third International Conference (ICML 2006), Pittsburgh, Pennsylvania, USA, June 25-29, 2006.* ACM International Conference Proceeding Series, vol. 148. ACM.

[11] Birn, Marcel, Holtgrewe, Manuel, Sanders, Peter, & Singler, Johannes. 2010. Simple and Fast Nearest Neighbor Search. *Pages 43–54 of: Proceedings of the Twelfth Workshop on Algorithm Engineering and Experiments, ALENEX 2010, Austin, Texas, USA, January 16, 2010.* SIAM.

[12] Blelloch, Guy E., Fineman, Jeremy T., Gu, Yan, & Sun, Yihan. 2020a. Optimal Parallel Algorithms in the Binary-Forking Model. *In: ACM Symposium on Parallelism in Algorithms and Architectures (SPAA).*

[13] Blelloch, Guy E., Anderson, Daniel, & Dhulipala, Laxman. 2020b. ParlayLib - A Toolkit for Parallel Algorithms on Shared-Memory Multicore Machines. *In: ACM Symposium on Parallelism in Algorithms and Architectures (SPAA).*

[14] Blumofe, Robert D., & Leiserson, Charles E. 1999. Scheduling multithreaded computations by work stealing. *J. ACM*, **46**(5), 720–748.

[15] Callahan, Paul B, & Kosaraju, S Rao. 1995. A decomposition of multidimensional point sets with applications to k-nearest-neighbors and n-body potential fields. *J. ACM*, **42**(1).

[16] Chan, Timothy. 2006. *A minimalist's implementation of an approximate nearest neighbor algorithm in fixed dimensions.*

[17] Chan, Timothy M. 2008. Well-separated pair decomposition in linear time? *Inf. Process. Lett.*, **107**(5), 138–141.

[18] Chen, Yewang, Zhou, Lida, Tang, Yi, Singh, Jai Puneet, Bouguila, Nizar, Wang, Cheng, Wang, Hua-zhen, & Du, Ji-Xiang. 2019. Fast neighbor search by using revised k-d tree. *Inf. Sci.*, **472**, 145–162.

[19] Clarenz, Ulrich, Rumpf, Martin, & Telea, Alexandru. 2004. Finite Elements on Point Based Surfaces. *Pages 201–211 of: 1st Symposium on Point Based Graphics, PBG 2004, Zurich, Switzerland, June 2-4, 2004.* Eurographics Association.

[20] Clarkson, Kenneth L. 1994. An Algorithm for Approximate Closest-Point Queries. *Pages 160–164 of: Proceedings of the Tenth Annual Symposium on Computational Geometry, Stony Brook, New York, USA, June 6-8, 1994.* ACM.

[21] Clarkson, Kenneth L. 1999. Nearest Neighbor Queries in Metric Spaces. *Discret. Comput. Geom.*, **22**(1), 63–93.

[22] Connor, Michael, & Kumar, Piyush. 2010. Fast Construction of k-Nearest Neighbor Graphs for Point Clouds. *IEEE Trans. Vis. Comput. Graph.*, **16**(4), 599–608.

[23] Erickson, Jeff. 2003. Nice Point Sets Can Have Nasty Delaunay Triangulations. *Discret. Comput. Geom.*, **30**(1), 109–132.

[24] Fleishman, Shachar, Cohen-Or, Daniel, & Silva, Cláudio T. 2005. Robust moving least-squares fitting with sharp features. *ACM Trans. Graph.*, **24**(3), 544–552.

[25] Gago, Silvia, Schlatter, & Dirk. 2009. Bounded expansion in web graphs. *Commentationes Mathematicae Universitatis Carolinae*, **50**(2), 181–190.

[26] Hu, Linjia, Nooshabadi, Saeid, & Ahmadi, Majid. 2015. Massively parallel KD-tree construction and nearest neighbor search algorithms. *Pages 2752–2755 of: 2015 IEEE International Symposium on Circuits and Systems, ISCAS 2015, Lisbon, Portugal, May 24-27, 2015.* IEEE.

[27] Intel Threading Building Blocks. https://www.threadingbuildingblocks.org.

[28] JáJá, Joseph. 1992. *An Introduction to Parallel Algorithms.* USA: Addison Wesley Longman Publishing Co., Inc.

[29] Karger, David R., & Ruhl, Matthias. 2002. Finding nearest neighbors in growth-restricted metrics. *Pages 741–750 of: Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada.* ACM.

[30] Kazana, Wojciech, & Segoufin, Luc. 2013. Enumeration of first-order queries on classes of structures with bounded expansion. *Pages 297–308 of: Proceedings of the 32nd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS 2013, New York, NY, USA - June 22 - 27, 2013.* ACM.

[31] López-Sastre, Roberto Javier, Oñoro-Rubio, Daniel, Gil-Jiménez, Pedro, & Maldonado-Bascón, Saturnino. 2012. Fast reciprocal nearest neighbors clustering. *Signal Process.*, **92**(1), 270–275.

[32] Mitra, Niloy J., Nguyen, An Thanh, & Guibas, Leonidas J. 2004. Estimating surface normals in noisy point cloud data. *Int. J. Comput. Geom. Appl.*, **14**(4-5), 261–276.

[33] Pajarola, Renato. 2005. Stream-Processing Points. *Pages 239–246 of: 16th IEEE Visualization Conference, IEEE Vis 2005, Minneapolis, MN, USA, October 23-28, 2005, Proceedings.* IEEE Computer Society.

[34] Pauly, Mark, Keiser, Richard, Kobbelt, Leif, & Gross, Markus H. 2003. Shape modeling with point-sampled geometry. *ACM Trans. Graph.*, **22**(3), 641–650.

[35] Ram, Parikshit, & Sinha, Kaushik. 2019. Revisiting kd-tree for Nearest Neighbor Search. *Pages 1378–1388 of: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2019, Anchorage, AK, USA, August 4-8, 2019.* ACM.

[36] Salam, Gavin, & Cacciari, Matteo. 2006. Jet clustering in particle physics, via a dynamic nearest neighbour graph implemented with CGAL. 04.

[37] Segoufin, Luc, & Vigny, Alexandre. 2017. Constant Delay Enumeration for FO Queries over Databases with Local Bounded Expansion. *Pages 20:1–20:16 of: 20th International Conference on Database Theory, ICDT 2017, March 21-24, 2017, Venice, Italy.* LIPIcs, vol. 68. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

[38] Shun, Julian, Blelloch, Guy E., Fineman, Jeremy T., Gibbons, Phillip B., Kyrola, Aapo, Simhadri, Harsha Vardhan, & Tangwongsan, Kanat. 2012. Brief announcement: the problem based benchmark suite. *Pages 68–70 of: 24th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA '12, Pittsburgh, PA, USA, June 25-27, 2012.* ACM.

[39] Singh, Aditi, Subramanya, Suhas Jayaram, Krishnaswamy, Ravishankar, & Simhadri, Harsha Vardhan. 2021. FreshDiskANN: A Fast and Accurate Graph-Based ANN Index for Streaming Similarity Search. *CoRR*, **abs/2105.09613**.

[40] Tangelder, Hans, & Fabri, Andreas. 2020. dD Spatial Searching. *In: CGAL User and Reference Manual*, 5.2 edn. CGAL Editorial Board.

[41] Vaidya, Pravin M. 1986. An optimal algorithm for the All-Nearest-Neighbors Problem. *Pages 117–122 of: 27th Annual Symposium on Foundations of Computer Science, Toronto, Canada, 27-29 October 1986.* IEEE Computer Society.

A Other Nearest Neighbor Implementations

For the purpose of comparison we use three existing implementations of nearest neighbor search: CGAL [4], STANN [22], and Chan [16]. Here we describe some performance issues with their code, and some modifications we made to improve the performance of their code.

Furthermore, we give a brief explanation for why we do not benchmark Delaunay triangulation based methods, since those are sometimes used for computing nearest neighbors in low dimensions. Firstly, the existing implementations for finding Delaunay triangulations were very slow: on 10 million points, the ParlayLib built-in function takes a few seconds, compared to less than .05 seconds to build the kd-tree. Secondly, the literature we found on Delaunay triangulations focused on the 2D case [11], while our experiments focused on the 3D case.

Chan. Chan's code was fully sequential so we needed to parallelize it. Conceptually this is relatively straightforward since the algorithm just requires using a parallel sort instead of a sequential one, and then running the queries in parallel. Indeed the first step of using a parallel sort was easy and we just replaced the C++ STL sort with the ParlayLib sort. The second step was required some work since the code was not thread safe. However, once modified and using the parallel loop from ParlayLib the code achieves very good speedup—about 75-fold speedup on 72 cores with 144 threads (see Section 4 for the full details).

Chan's algorithm only describes how to search from the root, and correspondingly his code only searches from the root. There seems to be no inherent reason that it would not be possible to start at the leaves when generating a nearest neighbor graph, but we did not implement such a variant. We note that even the root-based implementation of our code is significantly faster than Chan's. Chan's code uses arrays to store the points and therefore does not support dynamic updates. In his paper he mentions that his algorithm could support dynamic updates by storing the points in a balanced binary search tree in Morton order. This would require completely rewriting their code. Experiments with Chan's code only deal with the $k = 1$ case since his algorithms did not provide support for higher k.

STANN. STANN includes both a k-nearest neighbor graph (KNNG) function and a k-nearest neighbor (KNN) function. The first finds the k nearest neighbors among a set of points, and the second supports a function to build a tree and a separate function to query a point for its k-nearest neighbors. They supply a parallel version of KNNG, that was parallelized with OpenMP, and only a sequential version of KNN. On our initial tests we were able to get performance on the parallel KNNG that closely match what they report. However, their algorithm did not scale well beyond 16 threads (their numbers agree with this). The main issues is that the algorithm left some components sequential, including the Morton order sort, and initializing various standard template library (STL) vectors. For a larger number of cores this became the bottleneck. We therefore updated their code to use ParlayLib using a parallel sort and replacing uses of STL vectors with ParlayLib sequences, which are initialized

in parallel. We also made a couple other optimizations, including changing the size of the base case of the recursive query from 4 to 10, and using vectors instead of a priority queues to store the nearest neighbors for a point when k is small. These changes made a significant improvement in performance, especially at a larger number of cores, as indicated in Figure **??**.

The STANN KNN code was fully sequential. We therefore parallelized this as well, which required much the same changes as we made to the parallel code (using a parallel sort, making all loops parallel, and using parlay sequences instead of STL vectors). We also run the queries in parallel, which required some minor changes to make their code thread safe. As with Chan's algorithm, STANN stores the points in an array (STL vector) and therefore does not support dynamic updates.

CGAL. CGAL implements a parallel version of their k-nearest neighbor code using the threading building blocks (TBB) [27]. We use their code directly with no modifications. We note that as with the original version of STANN KNNG, their code does not scale well past 16 or so threads. We looked into this and there are several reasons. Perhaps most fundamentally, although in their recursive routine for building the kd-tree they invoke the two recursive calls in parallel, they do the splitting within each node completely sequentially. At the root of the tree this means they do linear work completely sequentially. Indeed from a theoretical point of view their algorithm does a total of $O(n \log n)$ work (assuming uniform input) and has $O(n)$ span, meaning it only has $O(\log n)$ parallelism. Fixing this problem would require a major rewrite of their code.

A second issue is that they allocate their tree nodes by pushing onto the back of a TBB concurrent vector. Although this is thread safe, it requires a lock and becomes a bottleneck on a large number of threads. A similar issue appears to be true in the query. In particular although the code appears to be thread safe giving correct answers when run in parallel, there seems to be contention when there are many threads slowing them all down. This is often caused by some form of memory allocation, as with the build tree, but in this case we were not able to track down the source of the problem. Due to the particularly bad performance beyond 36 threads (which all are on one chip), we only report numbers up to 36 threads. Furthermore, since we observed wildly varying times with higher k, we only included times for $k < 10$ in our experiments.

B Dynamic Queries

In this appendix, we provide data and commentary on our experimental results for non-dynamic queries. The results are presented graphically in Figure 4.

Varying k, work efficiency. (Figure 5(a) and (b).) All the algorithms we tested had similar performance for work efficiency and scalability of the number of nearest neighbors as they did in the k-nearest neighbor graph building case.

Scaling size of dataset. (Figure 5(c).) Our algorithms are particularly robust compared to others when

scaling the size of the dataset. One thing to notice in the relevant panel (c) of Figure 6 is that with the exception of CGAL, all algorithms experience a local minima when the size of the dataset reaches 10^5. While Chan, ParlayKNN, and CGAL's performances begin to slowly increase after this point, our algorithms' do not. The dataset of size 10^5 is approximately where we expect cache misses to start affecting the performance; as explained in Section 3, pre-sorting the data for dynamic queries helps alleviate this problem.

Other datasets. Since the bit-based and root-based algorithms performed very similarly on the random distribution for dynamic queries, we refer to Figure 4(f) for evidence that the bit-based algorithm outperforms the root-based for some distributions.

C Full Proof of Theorem 2.3

Here we extend the results in the main body where X is a general convex space rather than a bounding cube. This is only a concern for Theorem 2.3; Theorems 2.1 and 2.2 did not actually use the assumption on X.

Theorem 2.3 works by giving an upper bound on the number of points in an update that can be inserted into an unbalanced path of length ℓ. To generalize it to arbitrary convex spaces, we need to show that the upper bound a) still holds when the bounding box B of X is not a hypercube and b) when X is not a hypercube or hyperrectangle. The following lemma concerns the former property.

LEMMA C.1. *The number of points in X which can be part of an unbalanced path of length ℓ in T is maximized when the bounding box B of X is a hypercube.*

Proof. This follows from noting that the proof of Theorem 2.1 treated a split of X as a split of an oct- or quad-tree—that is, each split is a d-dimensional split of a bounding cube into 2^d boxes. Thus, if the longest length in the bounding hyperrectangle B of X is normalized to n and another side of B has length $r << n$, there will only be $\log r$ nonempty cuts of X perpendicular to that side. Thus the length of an unbalanced path in T will always be dominated by the length of the longest edge, and B being a bounding rectangle can only result in fewer points in an update traversing an unbalanced chain than if B were a hypercube. □

The second piece needed to apply Theorem 2.3 to general convex spaces X is to verify that the number of points near the boundary of X is still upper bounded when X is not a hyperrectangle.

LEMMA C.2. *Let X be convex with bounded expansion. Consider a subset of X with volume V such that volume V is a fraction f of X's volume. Then V contains at most $\left(\frac{\gamma^2}{\gamma^2+1}\right)^{\log 1/f} n$ points.*

Proof. $O(\log 1/f)$ divisions of X in half are needed to produce a shape of volume V. To maximize the number of points in V, we need to minimize the number of points that can be in the "sparse" half of each cut. The number of

points in the sparse half is dictated by how many individual boxes can be packed into the sparse area and then expanded twice to reach the boundary of the dense area. Thus, the number of points in the dense area is maximized when there is only one such box, i.e. when $d = 1$. This implies that we can upper bound the number of points in V by its upper bound when $d = 1$, i.e. a $\left(\frac{\gamma^2}{\gamma^2+1}\right)^{\log 1/f}$ fraction of the total points. This bound will be sufficient to finish the proof of Theorem 2.3. □

Now we are ready to prove the extended version of Theorem 2.3.

Proof. [Proof of Theorem 2.3] Fix a bounding box B. Now, for all convex shapes with bounding box B and more sides than a hyperrectangle, the surface-area-to-volume ratio decreases from that of a hyperrectangle. Thus, for a given distance to the boundary of X, the volume in X that can be at that distance or closer is only smaller when X has more sides than a hyperrectangle. Since Lemma C.2 tells us that the upper bound given in Lemma 2.5 applies no matter how the volume is arranged, the number of points in an update that can travel down an unbalanced path is no more than it would be if X were a hyperrectangle. Thus, we can take the number of points (and their corresponding cost) in an update to a hyperrectangle as an upper bound for this case.

However, triangles, tetrahedra, and their higher-dimensional counterparts have a higher surface-area-to-volume ratio than the hyperrectangle. Since the dimension is constant and the bounding box B is assumed to be the smallest possible, the surface-area-to-volume ratio of the tetrahedron is only a constant factor larger than that of the hypercube, and thus the result still holds. □
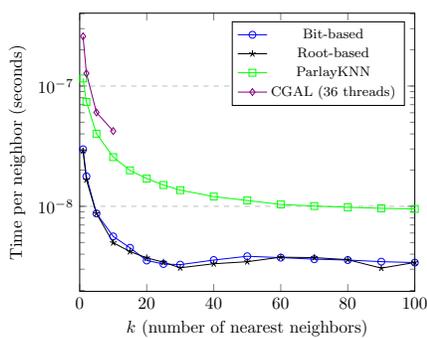
(a) Time required to calculate a neighbor as k increases. Calculated by dividing the total time by k times the number of queries.
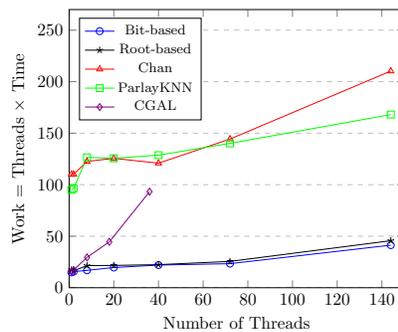
(b) Total work (threads × time) required to build a tree of 10 million points, then dynamically query the same number of points, on varying numbers of threads.
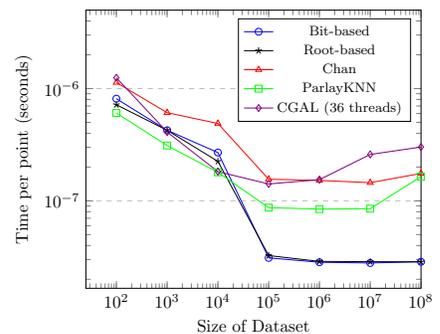
(c) Time required to calculate nearest neighbors as the size of the dataset increases. Calculated by dividing the total time by the number of points queried.

Figure 6: Statistics related to dynamic queries. Unless otherwise stated, the size of the dataset is 10 million, 10 million dynamic queries were performed, the number of nearest neighbors $k = 1$, experiments were performed on 144 threads on a 72-core Dell R930, and data points are drawn randomly from a 3D cube.