ParGeo: A Library for Parallel Computational Geometry

Yiqiu Wang ✉ ⓘ
MIT CSAIL

Rahul Yesantharao ✉ ⓘ
MIT CSAIL

Shangdi Yu ✉ ⓘ
MIT CSAIL

Laxman Dhulipala ✉ ⓘ
University of Maryland, College Park

Yan Gu ✉ ⓘ
University of California, Riverside

Julian Shun ✉ ⓘ
MIT CSAIL

──── **Abstract** ────

This paper presents ParGeo, a multicore library for computational geometry. ParGeo contains modules for fundamental tasks including kd-tree based spatial search, spatial graph generation, and algorithms in computational geometry.

We focus on three new algorithmic contributions provided in the library. First, we present a new parallel convex hull algorithm based on a reservation technique to enable parallel modifications to the hull. We also provide the first parallel implementations of the randomized incremental convex hull algorithm as well as a divide-and-conquer convex hull algorithm in $\mathbb{R}^3$. Second, for the smallest enclosing ball problem, we propose a new sampling-based algorithm to quickly reduce the size of the data set. We also provide the first parallel implementation of Welzl's classic algorithm for smallest enclosing ball. Third, we present the BDL-tree, a parallel batch-dynamic kd-tree that allows for efficient parallel updates and k-NN queries over dynamically changing point sets. BDL-trees consist of a log-structured set of kd-trees which can be used to efficiently insert, delete, and query batches of points in parallel.

On 36 cores with two-way hyper-threading, our fastest convex hull algorithm achieves up to 44.7x self-relative parallel speedup and up to 559x speedup against the best existing sequential implementation. Our smallest enclosing ball algorithm using our sampling-based algorithm achieves up to 27.1x self-relative parallel speedup and up to 178x speedup against the best existing sequential implementation. Our implementation of the BDL-tree achieves self-relative parallel speedup of up to 46.1x. Across all of the algorithms in ParGeo, we achieve self-relative parallel speedup of 8.1–46.61x.

1 Introduction

Computational geometry algorithms have important applications in various domains, including computer graphics, robotics, computer vision, and geographic information systems [30, 47]. It is important to provide users with libraries of efficient computational geometry algorithms that they can easily use in their own higher-level applications. Furthermore, due to the growing sizes of data sets that need to be processed today, and the ubiquity of parallel

(multicore) machines, it is beneficial to use parallel algorithms to speed up computations. In this paper, we present the ParGeo library for parallel computational geometry, which includes a rich set of parallel algorithms for geometric problems and data structures, including kd-trees, k-nearest neighbor search, range search, well-separated pair decomposition, Euclidean minimum spanning tree, spatial sorting, and geometric clustering. ParGeo also contains a collection of geometric graph generators, including k-nearest neighbor graphs and various spatial networks. Algorithms from ParGeo can either run sequentially, or run using parallel schedulers such as OpenMP, Cilk, or ParlayLib.

While there exist numerous libraries for computational geometry, most of them are not designed for parallel processing. For example, Libigl [38] is a library that specializes in the construction of discrete differential geometry operators and finite-element matrices. However, only some aspects of Libigl take advantage of parallelism. In comparison, the algorithms and implementations of ParGeo are designed for parallelism, and target a different set of problems. CGAL (Computational Geometry Algorithms Library) [2] is a well-known library of computational geometry algorithms that includes a wide range of algorithms, but most implementations are not parallel. Batista et al. [15] targeted a few important algorithms, including spatial sorting, box intersection, and Delaunay triangulation for shared-memory parallel processing, with code in CGAL. In comparison, ParGeo targets similar classes of problems as CGAL, but *all* of our implementations are highly parallel. PMP [51], Cinolib [43], and Tetwild [37] are libraries for polygonal and polyhedron meshes, tackling different problems from ParGeo. MatGeom [5] is a library for sequential geometric computing with MATLAB. The Problem Based Benchmark Suite [50, 12] is a multicore benchmark suite that has some overlap in algorithms with ParGeo. LEDA [44] is a library of data structures and algorithms for sequential combinatorial and geometric processing. ArborX [42] is a parallel library for spatial search.

In this paper, in addition to providing an overview of work on ParGeo, we describe new parallel algorithms implemented in ParGeo for convex hull, smallest enclosing ball, and batch-dynamic kd-tree that we developed. For convex hull, we develop new parallel algorithms for both $\mathbb{R}^2$ and $\mathbb{R}^3$, where our key algorithmic novelty is a reservation technique to enable parallel modifications to the hull. For smallest enclosing ball, we propose a new sampling-based algorithm based on Larsson et al.'s [41] approach to quickly reduce the size of the data set. We also provide the first parallel implementation of the classic randomized incremental algorithm [27]. For kd-trees, we develop the BDL-tree, a new parallel data structure that supports batch-dynamic operations (construction, insertions, and deletions) as well as exact k-NN queries. BDL-trees consist of a set of exponentially growing kd-trees and perform batched updates in parallel.

To demonstrate the efficiency of our proposed algorithms and library, we perform a comprehensive set of experiments on synthetic and real-world geometric data sets, and compare the performance across our parallel implementations as well as optimized sequential baselines. On 36 cores with two-way hyper-threading, our best convex hull implementations achieve up to 44.7x (42.8x on average) self-relative speedup and up to 559x (325x on average) speedup against the best existing sequential implementation for $\mathbb{R}^2$, and up to 24.9x (11.81x on average) self-relative speedup and up to 124x (61.4x on average) speedup against the best existing sequential implementation for $\mathbb{R}^3$. Our sampling-based smallest enclosing ball algorithm achieves up to 27.1x (20.08x on average) self-relative speedup and up to 178x (109x on average) speedup against the best existing sequential implementation for $\mathbb{R}^2$ and $\mathbb{R}^3$. Our BDL-tree achieves self-relative speedup of up to $35.4\times$ ($30.0\times$ on average) for construction, up to $35.0\times$ ($28.3\times$ on average) for batch insertion, up to $33.1\times$ ($28.5\times$ on
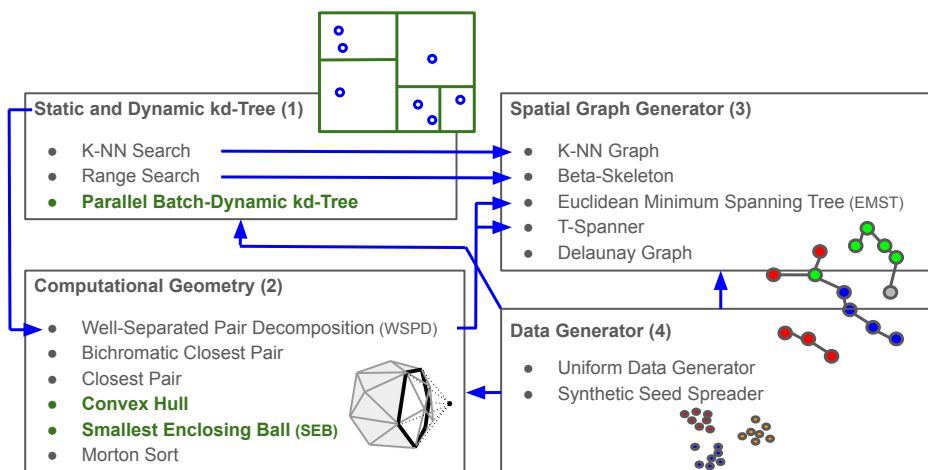
Figure 1 The figure shows an overview of modules in ParGeo. An arrow indicates that a component is used inside another component. In this paper, we present new algorithms and techniques for the modules highlighted in green.

average) for batch deletion, and up to $46.1\times$ ($40.0\times$ on average) for full k-NN. Finally, across all implementations in ParGeo, we achieve self-relative parallel speedup of 8.1–46.6x (on average 23.2x).

2   The ParGeo Library

Our main goal in designing ParGeo was to enable reusable and efficient parallel implementations of geometric algorithms and data structures. We present an overview of the modules of ParGeo in Figure 1, highlighting how the modules interact with each other. ParGeo contains efficient multicore implementations of static and batch-dynamic kd-trees (Module (1)). The code supports kd-tree based spatial search, including k-nearest neighbor and range search. Our code is optimized for fast kd-tree construction by performing the split in parallel (either by spatial median or by object median), and performing the queries in a data-parallel fashion. which we will introduce in Section 5.

ParGeo contains a module for parallel computational geometry algorithms (Module (2)). Our kd-tree can be used to generate a well-separated pair decomposition [26] (WSPD), which can in turn be used to compute the hierarchical DBSCAN [56], ParGeo contains parallel implementations for the bichromatic closest pair, closest pair, convex hull, smallest enclosing ball, and Morton sorting.

In addition, ParGeo contains a collection of geometric graph generators (Module (3)) for point data sets. Our kd-tree's k-NN search is used to generate the k-NN graph, and the range search is used to generate the β-skeleton graph [40]. Our WSPD generated from the kd-tree can also be used to compute the Euclidean minimum spanning tree [25, 56], and spanners [26]. ParGeo also generates the Delaunay graph.

ParGeo contains a point data generator module (Module (4)) for which can generate uniformly distributed data sets, and clustered data sets of varying densities [33]. These data sets are used for benchmarking the other modules.

As shown in Table 1, on a machine with 36 cores with two-way hyper-threading, ParGeo achieves self-relative parallel speedups of 8.1–46.61x (23.15x on average) on a uniformly distributed data set, across all of the benchmarks. In the subsequent sections, we present three new algorithmic contributions provided in the library.

Implementation	T_1	T_{36h}	Speedup
kd-tree Build (2d)	5.51	0.43	12.70x
kd-tree Build (5d)	8.39	0.89	9.40x
kd-tree k-NN (2d)	31.45	0.68	46.34x
kd-tree Range Search (2d)	17.14	0.37	46.61x
Batch-dynamic kd-tree Construction (5d)	6.70	0.60	10.70x
Batch-dynamic kd-tree Insert (5d)	8.80	1.10	8.10x
Batch-dynamic kd-tree Delete (5d)	29.20	1.20	23.90x
WSPD (2d)	6.72	0.24	27.63x
EMST (2d)	33.02	1.58	20.86x
Convex Hull (2d)	0.38	0.0088	43.13x
Convex Hull (3d)	2.36	0.097	24.36x
Smallest Enclosing Ball (2d)	0.053	0.0033	16.30x
Smallest Enclosing Ball (5d)	0.13	0.014	9.54x
Closest Pair (2d)	10.35	0.52	19.90x
Closest Pair (3d)	28.00	2.32	12.07x
k-NN Graph (2d)	37.89	1.46	25.99x
Delaunay Graph (2d)	55.91	2.03	27.53x
Gabriel Graph (2d)	59.61	1.99	29.99x
β-skeleton Graph (2d)	113.27	3.20	35.37x
Spanner (2d)	27.19	2.15	12.67x

■ **Table 1** Runtimes (seconds) and parallel speedups (T_1/T_{36h}) for PARGEO implementations on uniform hypercube data sets of varying dimensions and 10 million points. T_1 and T_{36h} denote the single-threaded and the 36-core hyper-threaded times, respectively. For batch-dynamic kd-tree updates, each batch contains 10% of the data set.
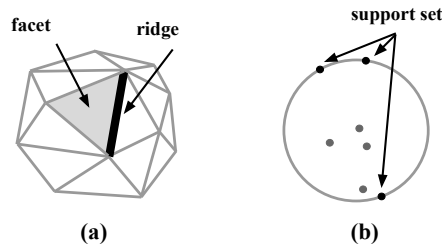

■ **Figure 2** (a) A facet and a ridge of a convex hull in $\mathbb{R}^3$. (b) The support of the smallest enclosing ball in $\mathbb{R}^2$.

3   Convex Hull

The convex hull of a set of points P in $\mathbb{R}^d$ is the smallest convex polyhedron containing P. It is common to represent the convex hull using a set of ***facets***. The boundary of two facets is a ***ridge***. For example, in $\mathbb{R}^3$, assuming the points are in general position (no four points are on the same plane), each facet is a triangle, and each ridge is a line that borders two facets (see Figure 2(a)).

     The randomized incremental algorithm and the quickhull algorithm are the most widely used algorithms for solving convex hull in practice. The randomized incremental algorithm for $\mathbb{R}^d$ was proposed by Clarkson and Shor [27]. Given a point data set P in $\mathbb{R}^d$, the randomized incremental algorithm first constructs a d-simplex, a generalization of a tetrahedron in d-dimensions as the initial hull. Then, the algorithm adds the points to the polyhedron in a random order, updating the hull if necessary. In practice, the quickhull algorithm [36, 14], another incremental algorithm, is often used. Unlike the randomized algorithm, the quickhull algorithm processes a point that is furthest from a facet, which enables the hull to be expanded more quickly. The quickhull algorithm is by far one of the most common implementations for convex hull due to its simplicity and efficiency [4, 6, 7, 1, 3, 2]. There have also been works that study parallel implementations of quickhull, but they are either limited to $\mathbb{R}^2$ [45, 52], or do not return the exact convex hull for $\mathbb{R}^3$ [53, 55]. Recently, Blelloch et al. [24] proposed a new
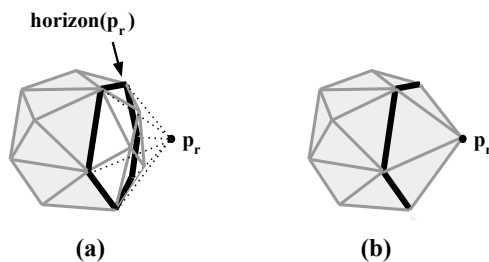
■ **Figure 3** Illustration of adding a visible point p_r to the convex hull. (a) shows the convex hull prior to the addition of p_r. The visible facets are in white, while the non-visible facets are in gray. The thicker line segments correspond to the horizon. (b) shows the convex hull after adding p_r with newly created facets.

randomized incremental algorithm that is highly parallel in theory. However, the algorithm does not seem to be practical due to numerous data structures required for bookkeeping.

In this section, we describe our new parallel reservation-based algorithm. Our algorithm is able to express both the randomized incremental convex hull algorithm and the quickhull algorithm. Specifically, unlike a sequential incremental algorithm that adds one point per round, we add multiple points in parallel per round. We resolve conflicts caused by the parallel insertion using a reservation technique. We also apply a general parallelization technique based on divide-and-conquer, which in combination with our parallel incremental algorithm, leads to faster implementations in practice.

Parallel Reservation-Based Algorithm. Our parallel reservation-based algorithm can be implemented as either a randomized incremental algorithm or a quickhull algorithm. We will first introduce the overall structure of the algorithm. Then, we will describe the details about the implementations, and compare with existing approaches. We will base our description in the context of $\mathbb{R}^3$ for the sake of clarity, but the algorithm can be extended to $\mathbb{R}^d$ for any constant integer $d \geq 2$.

We first give a high-level overview of the algorithm. Given an ordered set of points $P = \{p_1, p_2, \ldots, p_n\}$, we let $P_r = \{p_1, p_2, \ldots, p_r\}$ be the prefix of P of size r, and $CH(P_r)$ be the convex hull on P_r. We start the construction by first arbitrarily selecting four points from P that do not lie on the same plane and putting them at the beginning of P, forming a tetrahedron $CH(P_4)$. Then, the algorithm proceeds iteratively, but on each round, rather than inserting just p_r to form $CH(P_r)$, we process a batch of points in parallel. On each round, let each point outside of $CH(P_{r-1})$ be called a ***visible point***. We first select a batch of visible points, and try to add them to $CH(P_{r-1})$ in parallel in the same round.

The key challenge of this approach is that some of these points cannot be processed in parallel due to concurrent modifications on the shared structures of the convex polyhedron. We use a reservation algorithm to resolve these conflicts, such that we only process the points that modify disjoint facets of the polyhedron. Specifically, each point will perform a priority write [49] with its ID to reserve all of its visible facets. Points that have their ID written to all of its visible facets are *successful*. We then process the successful points in parallel by enabling them to make modifications to $CH(P_{r-1})$. At the end of the round, in parallel, we filter out points that are no longer visible. The algorithm will terminate when there are no more visible points.

We now describe the algorithm in more detail. Figure 3 illustrates the processing of a visible point p_r. We denote a facet as a ***visible facet*** of p_r if point p_r is in the half space away from the center of the convex hull. We first retrieve the set of visible facets of p_r via facets stored in it. The visible facets of p_r form a closed region, whose boundary is a set of ridges, known as the ***horizon***. We delete the visible facets from $CH(P_{r-1})$, and replace
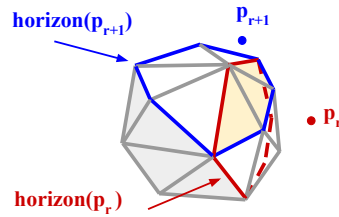
■ **Figure 4** This figure illustrates the attempt to add p_r and p_{r+1} in parallel. The visible points and horizons of p_r and p_{r+1} are in red and blue, respectively. The visible facets to either visible points are in white/yellow, while the other facets are in gray. The overlap of the three visible facets between the p_r and p_{r+1} is in yellow.

```
1   Input: 3-dimensional points P, batch size r
2   Output: 3-dimensional convex hull
3   CH := initialize with 4 points
4   while (P is not empty):
5     Q := a batch of size r of visible points in P
6     par_for (q in Q): /* reservation */
7       for (f in q.visibleFacets):
8         WriteMin(&f.reservation, q.id)
9     par_for (q in Q): /* check reservation */
10      for (f in q.visibleFacets):
11        q.success &&= (f.reservation == q.id)
12    par_for (q in Q): /* process successful points */
13      if (q.success):
14        delete q's visible facets
15        create new facets of q
16        update CH
17    P := ParallelPack(P, visible)
```

■ **Figure 5** Pseudocode for the parallel reservation-based convex hull algorithm (which includes the randomized incremental algorithm and the quickhull algorithm).

them with new facets, where each new facet is formed by adding two ridges from a horizon ridge to p_r.

Because of the structural changes to the convex hull that occur when adding a visible point, concurrent structural changes can cause data races, which need to be avoided. We show an example of the conflict in Figure 4, where we are attempting to add two visible points p_r and p_{r+1} in parallel. As shown in the figure, the closed regions formed by the visible facets of each visible point overlap with each other in three facets, which are highlighted in yellow. Should the two visible points be processed in parallel, the resulting polyhedron may not be well-defined due to data races. When processed sequentially, p_{r+1}'s visible facets would have been different, involving newly created facets by p_r.

Our reservation algorithm allows only a subset of the visible points that update disjoint facets of the convex hull to be processed in parallel on each round. At a high level, we use the lexicographical order of the visible points to determine the priority in processing a facet (a smaller ID has higher priority). In the example shown in Figure 4, since p_r has a smaller ID than p_{r+1}, the three conflicting facets can only be processed by p_r in that round. The pseudocode for the algorithm is shown in Figure 5. P is processed iteratively until it is empty (Line 4). We allocate an extra data field in each facet for performing reservations (Lines 6–8). For each visible point in parallel, we iterate through its visible facets and use a parallel priority write (**WriteMin**) to write its ID to the facets' "reservation" fields. Then on Lines 9–11, we determine which visible points successfully reserved all of its facets. Again, in

parallel for each visible point, we check each of its visible facets for a successful reservation by comparing the value of the reservation field with its token. The reservation of a visible point is only successful if its ID is stored in all of its visible facets. Then, on Lines 12–16, we process the visible points whose reservations are successful, adding them to the hull and updating the appropriate data structures. Finally, on Line 17, we process the points in P such that those remaining as visible points are packed to replace the original P, and the non-visible points are discarded. Note that the visible points that succeeded in the reservation are no longer visible points because they are now part of the convex hull. Some of the remaining points will also no longer be visible points due to the growth of the convex hull.

We use a simple and fast data structure to keep track of the visibility relationship between the visible points and the facets. At each step of the algorithm, when a visible point is processed, it needs to identify the set of visible facets. On the other hand, for the facets undergoing structural changes, they need to identify and redistribute their visible points to new facets. To find the set of visible facets of p_r, it is inefficient to iterate through all of the facets of $CH(P_{r-1})$. While existing approaches [30] keep track of the visibility between visible points and *all* of their visible facets, we found such an approach to be slow becaseu each vertex is associated with multiple facets, making the cost of storing and updating the data structure high. We only store the reference of an arbitrary visible facet to each visible point, from which we use a local breadth-first search to retrieve all of the visible facets only when needed. For storing the visible points in the facets, we assign each point to one of its visible facets. During point redistribution, we gather the points stored in each visible facet into an array, and in parallel distribute each point to a new visible facet replacing the original visible facet. Each such point also stores a reference to this visible facet.

Our reservation-based algorithm can be used to implement the parallel randomized incremental algorithm or the quickhull algorithm for convex hull. For the randomized incremental algorithm, we randomly permute the input points at the beginning, and on each round attempt to add a prefix of the permuted points to the convex hull. For the quickhull algorithm, on each round, we instead select a set of points furthest from a subset of facets. We describe the implementation of the two algorithms in greater detail in Appendix A. Our reservation-based algorithm is inspired by the idea of "deterministic reservations" from Blelloch et al. [22], who introduce this approach to implement parallel algorithms for other problems. We also show the work overhead of doing reservations compared to the sequential algorithm is small in Appendix B.

Parallel Divide-and-Conquer. We adopt a common parallelization strategy using divide-and-conquer, which calls our reservation-based algorithm as a subroutine. Some early convex hull algorithms are based on divide-and-conquer, notably, the algorithm by Preparata and Hong [46]. The algorithm splits the input into two spatially disjoint subsets by a mid-point along one of the axis, recursively computes the convex hull on each subset, and then merges the results together. Later work [10, 29, 11] extended this approach to the parallel setting. However, most of these approaches rely on complicated subroutines to merge convex hulls, which are not practical and have not been implemented, to the best of our knowledge.

We implement a practical divide-and-conquer algorithm by partitioning the input into $c \cdot numProc$ equal subsets, where c is a small constant and $numProc$ is the number of processors. For each subset, the convex hull of the subset is computed by a single processor using the sequential quickhull algorithm, but run in parallel across the different subsets. Then, the vertices of the outputs of the subproblems are collected to form a new input, from which the final convex hull is computed using our reservation-based parallel algorithm described earlier.

Point Culling via Pseudohull Computation. We also implement a multicore variant of Tang et al.'s pseudohull heuristic [54], originally proposed for the GPU. Starting from an initial tetrahedra, we recursively grow each facet into three new facets, using the furthest point from the facet, similar to the quickhull algorithm. The visible points associated with the facet are redistributed to the new facets. This results in a polyhedron, and the points in the interior of the polyhedron will not be part of the convex hull. Therefore, we can prune away the points inside the polyhedron and compute the convex hull on the rest of the points.

There are several differences in our implementation from Tang et al.'s algorithm. Our implementation executes the recursive calls on different facets asynchronously in parallel, whereas Tang et al.'s implementation maps the algorithm to the GPU architecture by pre-allocating space for the facets and visible points, and runs the algorithm in an iterative manner in lock-step. Specifically, successively generated facets and points points associated them are updated by multiple threads in parallel in each iteration. We use a parallel maximum-finding routine to find the furthest point of each facet in each call. Rather than growing the pseudohull until there are no more visible points as done by Tang et al., we set a threshold on the number of points associated with a facet, below which we stop growing the pseudohull. This prevents stack overflow on large and skewed data sets due to too many recursive calls, and the extra unpruned points do not contribute significantly to the work of the final computation of the convex hull. At the end of pruning, we use our parallel reservation-based quickhull algorithm to compute the final hull on the remaining points, whereas Tang et al. uses a sequential implementation.

4   Smallest Enclosing Ball

The smallest enclosing ball of P in $\mathbb{R}^d$ is the smallest d-sphere containing P. It is well known that the smallest enclosing ball is unique and defined by a ***support set*** of $d + 1$ points on the surface of the ball (see Figure 2(b)).

Welzl [57] showed that by using a randomized incremental algorithm, the smallest enclosing ball can be computed in $O(n)$ time in expectation for constant d. The algorithm iteratively expands the support set of the ball by adding points in a random order until the ball contains all of the points. The algorithm was later improved by Gartner [35] with practical optimizations for speed and robustness. Larsson et al. [41] proposed practical parallel algorithms that use a new method for expanding the support set, and their implementations work on both CPUs and GPUs. Later, Blelloch et al. [23] proposed a parallel algorithm based on Welzl's algorithm, but without any implementations.

In this section, we describe our new algorithms for the smallest enclosing ball problem based on Larsson et al.'s approach [41]. We propose a sampling-based algorithm to quickly reduce the size of the data set. We also provide the first parallel implementation of Welzl's classic algorithm.

Given a ball B, we define ***visible points*** to be points that lie outside of B. Existing approaches for computing the smallest enclosing ball focus on expanding the support set in an iterative manner, and output the enclosing ball when there are no more visible points. Welzl's algorithm expands the support set by adding points in a random order [57]. In comparison, Larsson et al.'s approach scans the input to search for good support sets in a round-based manner. In $\mathbb{R}^3$, Larsson's algorithm divides the space into eight orthants centered at the center of B. On each round, the input is scanned to find the furthest visible points in each orthant. B is then updated to the next intermediate solution using the existing support set of B and the new visible points found during the scan. The algorithm iterates

```
1   Input: d-dimensional points P, batch size c
2   Output: d-dimensional smallest enclosing ball
3   B = ball()
4   /* Sampling phase */
5   scanned = 0
6   while (scanned < n):
7     hasOutlier, support =
8       orthantScan(P[scanned:min(scanned+c,n)-1],B)
9     scanned += c
10    if (!hasOutlier):
11      break /* current sample does not violate B */
12    else
13      B = constructBall(support)
14  /* Final computation phase */
15  while (hasOutlier):
16    hasOutlier, support = orthantScan(P, B)
17    if (!hasOutlier):
18      return B
19    else
20      B = constructBall(support)
```

Figure 6 Pseudocode for the parallel sampling-based algorithm for smallest enclosing ball.

until there are no more visible points. It is parallelized within each round by performing the scan on the input in parallel.

Sampling-Based Algorithm. We find each iteration in Larsson et al.'s algorithm to be unnecessarily expensive due to having to scan the entire data set on every round. Our approach is to use a sampling heuristic to first obtain a good initial ball, inspired by Welzl's randomized algorithm. Specifically, we use small random samples to obtain good estimates of the support set at a negligible cost.

We show the pseudocode of our algorithm in Figure 6. Our sampling-based algorithm consists of two phases: the sampling phase (Line 5–13) and the final compute phase (Line 15–20). First, we initialize the ball using a few arbitrary points (Line 3). Then, we iterate through a random permutation of the input to take multiple samples (Line 5–13). On each iteration, we scan through a constant-sized segment of the unseen part of the input, which is equivalent to a random sample. We perform an orthant scan similar to Larsson's approach. Our implementation of orthant scan will return a new estimate of the support set based on the sample, and a boolean *hasOutlier* indicating whether the sample contains visible points with respect to the current smallest enclosing ball B (Line 7). We recompute B using the new support set. If there are visible points in the current sample, we will continue the sampling process with our new B. If there are no visible points in the sample, the support set likely contains most of the points, and so we terminate sampling and move on to the next phase. Now, with a good estimate of the optimal smallest enclosing ball, we run Larsson's orthant scan to compute the final smallest enclosing ball (Line 15–20). The sampling phase allows us to generate good support sets without having to scan the entire input.

We parallelize the orthant scan, which is the most expensive operation of the algorithm. Specifically, we divide the input array to orthant scan into blocks, and process each block sequentially, but in parallel across different blocks. Afterward, the extrema for the orthants obtained from the blocks are merged, and a new support set is computed on these points and the existing support set of B.

We parallelize the orthant scan, which is the most expensive operation of the algorithm. Specifically, we divide the input array to orthant scan into blocks, and process each block

sequentially, but in parallel across different blocks. Afterward, the extrema for the orthants obtained from the blocks are merged, and a new support set is computed on these points and the existing support set of B.

Parallel Welzl's Algorithm and Optimizations. We also implemented and optimized the parallel version of Welzl's algorithm described by Blelloch et al. [23]. Welzl's sequential algorithm uses a random permutation of the input P and processes the points one by one. If the algorithm encounters a visible point p_i with respect to the current bounding ball B, B is recomputed on P_i, the prefix of points up until p_i, using recursive calls to the algorithm. Blelloch et al.'s parallel algorithm also uses a random permutation of P. Across iterations, the algorithm processes prefixes of P of exponentially increasing size. If the prefix contains at least one visible point, the earliest visible point p_i is identified, and B is recomputed on prefix P_i by recursively calling the parallel algorithm. Each prefix is processed in parallel.

We implement the algorithm with some practical optimizations. When there are numerous visible points in the prefix, the work of the parallel algorithm will increase significantly, because each time a visible point is discovered, the points after the visible point in the same prefix will have to be reprocessed in the next round. Therefore, given that there will be more visible points in the initial rounds when the prefix size is small (< 500000), we process these prefixes sequentially by calling Welzl's sequential algorithm. This also reduces the amount of overhead from parallel primitives, since there is limited parallelism for small prefixes.

In addition, we extend existing optimizations of Welzl's sequential algorithm to the parallel setting. We implement the move-to-front heuristic [57], which upon encountering a visible point, moves the visible point to the front of P, so that it will be processed earlier in recursive calls, reducing the number of subsequent visible points. We also parallelize the pivoting heuristic proposed by Gartner [35]. In this heuristic, upon encountering a visible point, rather than processing the visible point directly, we search P for a *pivot point* furthest away from the center of the current B, and use the pivot point to compute the new B instead of the visible point. We use a parallel maximum-finding algorithm to identify the pivot point.

5   Parallel Batch-Dynamic kd-tree

The kd-tree, first proposed by Bentley [16], is a binary tree data structure that arranges and holds spatial data to speed up spatial queries. At each node, the data set is split into two using an axis-aligned hyperplane along a dimension, until the node holds a small constant number of points. kd-trees are used in a wide range of applications, such as in databases, machine learning, data compression, and cluster analysis.

In this section, we introduce the BDL-tree, a parallel batch-dynamic kd-tree implemented using the logarithmic method [17, 18]. Our BDL-trees build on ideas from the Bkd-Tree by Procopiuc et al. [48] and the cache-oblivious kd-tree by Agarwal et al. [9]. The logarithmic method [17, 18] for converting static data structures into dynamic ones is a very general idea. At a high level, the idea is to partition the static data structure into multiple structures with exponentially growing sizes (powers of 2). Then, inserts are performed by only rebuilding the smallest structure necessary to account for the new points. In the specific case of the kd-tree, a set of N_s static kd-trees is allocated, with capacities $[2^0, 2^1, \ldots, 2^{N_s-1}]$, as well as an extra buffer tree with size 2^0. Then, when an insert is performed, the insert cascades up from the buffer tree, rebuilding into the first empty tree with all the points from the lower trees. If desired, the sizes of all of the trees can be multiplied by a buffer size X, which is a constant that is tuned for performance.

We implement the underlying static kd-trees in an BDL-tree using the van Emde Boas

(a) Static tree 0 is full.

(b) Static tree 1 is full and buffer tree has 1 point.

(c) Static trees 0 and 1 are full and buffer tree has 2 points.

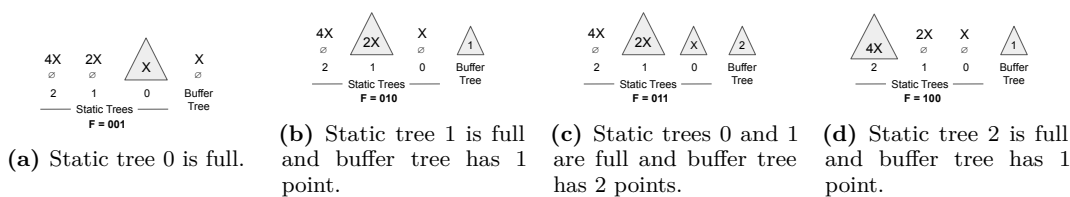**(d)** Static tree 2 is full and buffer tree has 1 point.

Figure 7 A BDL-tree in various configurations with $X > 2$; starting from (a), inserting $X + 1$ points gives (b), then inserting $X + 1$ points gives (c), and then inserting $X - 1$ points gives (d).

(vEB) [13, 31, 9] recursive layout. Agarwal et al. [9] show that this memory layout can be used with kd-trees to make traversal cache-oblivious. We provide more details of the static tree structure, and parallel algorithms for the construction, deletion, and k-NN search in Appendix C.1.

Parallel Batch Insertion. Batch insertions are performed in the style of the logarithmic method [17, 18], with the goal of maintaining the minimum number of full trees within BDL-tree. Thus, upon inserting a batch P of points, we rebuild larger trees if it is possible using the existing points and the newly inserted batch. We use a bitmask to determine which static trees in the structure to destroy and reconstruct after each insertion. Specifically, we build a bitmask F of the current set of full static trees. Given the buffer kd-tree size X, we add the value $\lfloor |P|/X \rfloor$ to F when a point set P is inserted, after which the bitwise difference with the previous F indicates which trees need to be changed. We gather the points in the trees to be destroyed, and with P, we construct a subset of new trees in parallel. As an implementation detail, note that we first add $|P| \mod X$ points to the buffer kd-tree—if we fill up the buffer kd-tree, then we gather the X points from it and treat them as part of P, effectively increasing the size of P by X. Refer to Figure 7 for an example of this insertion method ($X > 2$ in this example). In Figure 7(a), the BDL-tree contains X points, giving a bitmask of $F = 1$ (because only the smallest tree is in use). If we insert $X + 1$ points, then we put one point in the buffer tree and compute $F_{new} = 1 + \lfloor \frac{X}{X} \rfloor = 2$, and so we have to deconstruct static tree 0 and build static tree 1, as shown in Figure 7(b). Then, if we insert $X + 1$ points again, then we again put one point in the buffer tree and compute $F_{new} = 2 + \lfloor \frac{X}{X} \rfloor = 3$, and so we simply construct tree 0 on the X new points (leaving tree 1 intact), as seen in Figure 7(c). Finally, if we then insert $X - 1$ points, this would fill the buffer up, and so we take 1 point from the buffer and insert X points; then, $F_{new} = 3 + \lfloor \frac{X}{X} \rfloor = 4$, and so we deconstruct trees 0 and 1, and construct tree 2, as seen in Figure 7(d). We include a more detailed explanation of the algorithm in Appendix C.2, and explain the batch deletion algorithm in Appendix C.3.

Data-Parallel k-NN. In the data-parallel k-NN implementation, we parallelize over S, the set of points to search for nearest neighbors for. First, we allocate a k-NN buffer for each of the points in S. Then, iterating over each of the non-empty trees in the BDL-tree sequentially, we call the data-parallel k-NN subroutine on the tree, passing in the set S of points and the k-NN buffers. Because we reuse the same set of k-NN buffers for each k-NN call (note that each k-NN call is internally parallel), we end up with the k-nearest neighbors of the entire pointset for each point in S. We include a more detailed explanation in Appendix C.4.

6 Experimental Evaluation

Data Sets. We use several types of synthetic data sets. The first is **Uniform** (**U**), consisting of points distributed uniformly at random inside a hypercube with side length $\sqrt{n}$, where n is the number of points. The second type **InSphere** (**IS**) is similar to the first, but the points are distributed in a hypersphere instead. We also use **_OnSphere_** (**OS**) and **_OnCube_** (**OC**) data
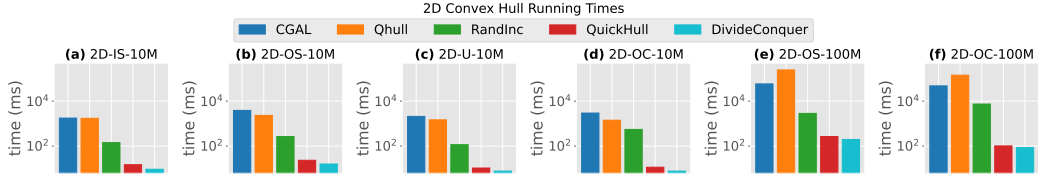
Figure 8 Running times of implementations across different data sets for 2-dimensional convex hull on 36 cores with 2-way hyper-threading.
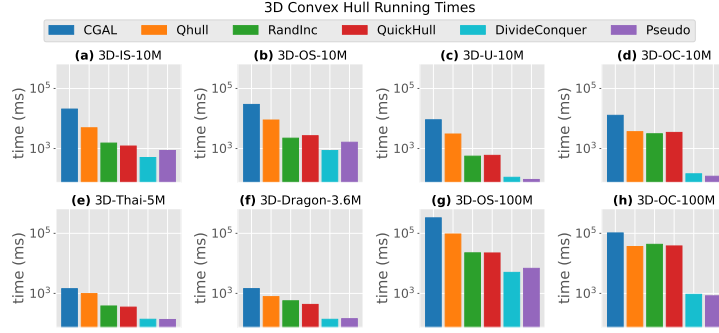
Figure 9 Running times of implementations across different data sets for 3-dimensional convex hull on 36 cores with 2-way hyper-threading.

sets, where points are uniformly distributed on the surface of a hypersphere and a hypercube, respectively. The surfaces have a thickness equal to 0.1 times the diameter or side length of the sphere or cube. We name the data sets in the format of **Dimension-Name-Size**.

We also use the following real-world data sets from the Stanford 3D Scanning Repository [8]: **3D-Thai-5M** is a 3-dimensional point data set of size 4999996 from a scanned thai-statue; and **3D-Dragon-3.6M** is a 3-dimensional point data set of size 3609600 from a scanned statue of a dragon.

Testing Environment. The experiments are run on an AWS c5.18xlarge instance with 2 Intel Xeon Platinum 8124M CPUs (3.00 GHz), for a total of 36 two-way hyper-threaded cores and 144 GB RAM. We compile our benchmarks with the `g++` compiler (version 9.3.0) with the `-O3` flag, and use ParlayLib [20] for parallelism.

6.1 Convex Hull

We test the following implementations for convex hull (our new implementations are underlined). All implementations are for both $\mathbb{R}^2$ and $\mathbb{R}^3$.

- **CGAL**: sequential C++ implementation of quickhull in CGAL [2].
- **Qhull**: sequential C++ implementation of quickhull [6] by Barber et al. [14].
- **RandInc**: our implementation of the parallel randomized incremental algorithm described in Section 3.
- **QuickHull**: for $\mathbb{R}^2$, it is a simple recursive parallel algorithm [19], and we use the implementation in PBBS [50]; for $\mathbb{R}^3$, we use our parallel quickhull algorithm described in Section 3.
- **Pseudo**: our implementation of the pseudoHull heuristic proposed by Tang et al. [54] for 3-dimensional convex hull described in Section 3. The final stage of the computation uses our *quickHull* algorithm for $\mathbb{R}^3$.
- **DivideConquer**: our divide-and-conquer algorithm described in Section 3.

In Figures 8 and 9, We show a comparison of running times across different methods using 36 cores with two-way hyper-threading. Our implementations achieve significant speedup

compared to existing sequential implementations. Our fastest parallel implementations achieve speedups of 190–559x (325x on average) over *CGAL* for 2-dimensional convex hull, and speedups of 10.5–124x (61.4x on average) over *CGAL* for 3-dimensional convex hull. Our fastest parallel implementations have speedups of 147–1673x (605x on average) over 2-dimensional *Qhull*, and speedups of 5.68–43.8x (19.9x on average) over 3-dimensional *Qhull*. When run using a single thread, our parallel implementations achieve speedups of 3.26–12.4x and 1.31–5.05x over *CGAL* for 2 and 3 dimensions, respectively; and 3.39–47.6x and 0.99–2.06x speedups over *Qhull* for 2 and 3 dimensions, respectively.

For $\mathbb{R}^2$, *DivideConquer* is always the fastest method due to having high scalability from processing many independent subproblems in parallel. For $\mathbb{R}^3$, the fastest two methods are *DivideConquer* and *Pseudo*. We observe that on data sets with a larger output size, *Pseudo* is slower than *DivideConquer* (Figures 9(a), (b), and (g)). This is because the final computation after pruning takes longer given that there are a higher number of remaining points after pruning. For instance, the number of remaining points for *3D-IS-10M* and *3D-U-10M* after pruning are 83669 and 2316, respectively, and *Pseudo* is relatively slower on the former. We observe that *RandInc* and *QuickHull* take relative longer compared with the fastest methods for data sets with a smaller output size (Figures 9(c)–(e) and (h)). This is caused by higher contention during the reservation of facets, since there are fewer facets on the intermediate hull. For example, for *3D-IS-10M* and *3D-U-10M*, the output sizes are 14163 and 423, respectively. During the computation, *3D-U-10M* exposes fewer facets for reservation, leading to a lower success rate during the reservations.

DivideConquer achieves the best parallel speedup (42.78x and 16.55x on average for $\mathbb{R}^2$ and $\mathbb{R}^3$, respectively). This is because the bulk of the time is spent in computing independent convex hulls across different threads. On the other hand, the incremental algorithms, *RandInc* and *QuickHull*, demonstrate lower scalability because of load imbalance caused by the different amounts of work for each conflict point being processed in parallel.
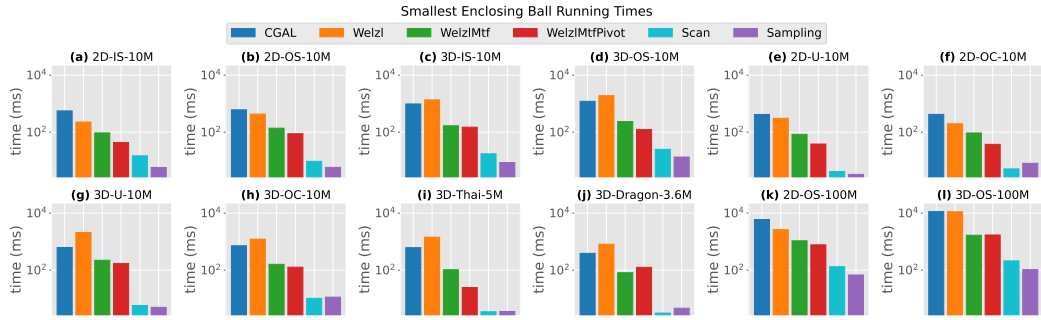
6.2 Smallest Enclosing Ball

We test the following implementations for smallest enclosing ball (our new implementations are underlined). All implementations work for both $\mathbb{R}^2$ and $\mathbb{R}^3$.

- **CGAL**: sequential C++ implementation of Welzl's algorithm in CGAL [2].
- **Orthant-scan**: our implementation of Larsson et al.'s orthant-scan algorithm [41].
- **Sampling**: our parallel sampling algorithm described in Section 4.
- **Welzl**: our parallel implementation of Welzl's algorithm described in Section 4.
- **WelzlMtf**: the same as *Welzl*, but with the move-to-front heuristic [10].
- **WelzlMtfPivot**: the same as *Welzl*, but with both the move-to-front and the pivoting heuristic [35].

For smallest enclosing ball, we show the comparison across implementations using 36 cores with two-way hyper-threading in Figure 10. Our fastest parallel implementations have speedups of 70–178x (109x on average) over *CGAL*. On one thread, our fastest implementations achieve speedup of 2.81–7.05x (4.96x on average) over *CGAL*.

Our sampling-based method is the fastest for eight out of the twelve data sets, whereas *Orthant-scan* without sampling is the fastest for the other four. We observe that the sampling phase on average scans only about 5% of the data set, and results in up to 2.55x (1.47x on average) speedup compared to just running *Orthant-scan*. Comparing across different implementations of Welzl's algorithms, we see that the move-to-front, and the pivoting heuristic implemented in parallel consistently improve the running times. Specifically, *WelzlMtf* is 2.09–13.9x faster than *Welzl*, and *WelzlMtfPivot* is 3.4–58.6x faster than *Welzl*.

Figure 10 Running times of implementations across different data sets for smallest enclosing ball on 36 cores with 2-way hyper-threading.

We also see that *Sampling* and *Orthant-scan* are 4.63–34.8x and 2.96–40.3x faster than *WelzlMtfPivot*, respectively.

6.3   BDL-tree

We designed a set of experiments to investigate the performance and scalability of BDL-tree and compare it to two baselines that we also implemented. **B1** is a baseline where the kd-tree is rebuilt on each batch insertion and deletion in order to maintain balance. This allows for improved query performance (as the tree is always perfectly balanced) at the cost of slowing down updates. **B2** is another baseline that inserts points directly into the existing tree structure without recalculating the splits. This results in very fast updates at the cost of potentially skewed trees (which slows down query performance). **BDL** is our BDL-tree described in Section 5. We consider splitting the points based on either using the object median (median among the points along a dimension) or the spatial median (splitting the space along a dimension in half).

Construction. Figure 11(a) shows the scalability of the throughput on the 7D-U-10M data set. As we can see from the results, **BDL** achieves similar or better performance both serially and in parallel than both **B1** and **B2**, and has similar or better scalability than both. With the object median splitting, it achieves up to $34.8\times$ self-relative speedup, with an average self-relative speedup of $28.4\times$. We also note that the single-threaded runtimes are faster with spatial median splitting than with object median splitting. This is because spatial median only involves splitting points at each level compared with finding the median for object median, hence it is less expensive to compute; however, we also note that the scalability for spatial median is lower because there is less work to distribute among parallel threads. The construction of B2 is significantly slower than that of B1, because a separate memory buffer is allocated at each leaf node in B2 to allow for future insertions. The construction of the BDL-tree is faster than both B1 and B2 because splitting the construction across multiple trees while keeping the number of elements the same reduces the total work, and provides ample parallelism when running on multiple threads.

Batch Insertion. In this benchmark, we measure the performance of our batch insertion implementation as compared to the baselines. We measure the time required to insert 10 batches each containing 10% of the points in the data set into an initially empty tree for each of our two baselines as well as our BDL-tree.

From Figure 11(b), we see that **B2** achieves the best performance on batched insertions—this is due to the fact that it does not perform any extra work to maintain balance and simply directly inserts points into the existing spatial structure. **BDL** achieves the second-best performance—this is due to the fact that it does not have to rebuild the entire tree on every
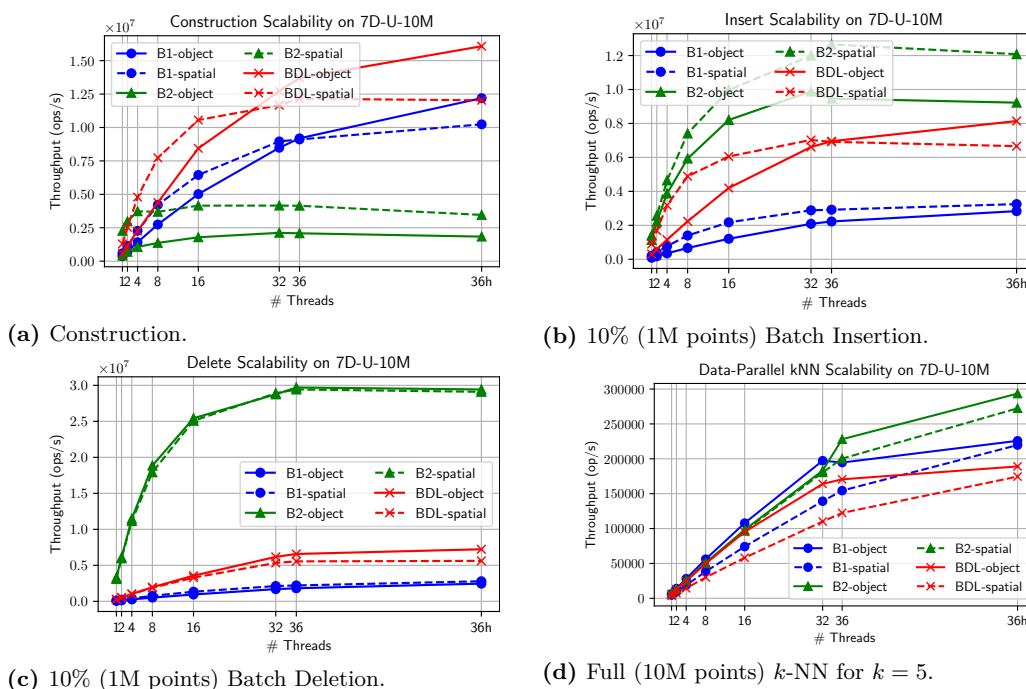
(a) Construction.

(b) 10% (1M points) Batch Insertion.

(c) 10% (1M points) Batch Deletion.

(d) Full (10M points) k-NN for $k = 5$.

Figure 11 Plot of throughput (operations per second) of batch operations over thread count for both object and spatial median implementations for the 7D-U-10M data set. The prefix of the implementation name refers to the median splitting heuristic. "36h" corresponds to 36 cores with two-way hyper-threading.

insert, but amortizes the rebuilding work across the batches. Finally, **B1** has the worst performance, as it must fully rebuild on every insertion. Similar to construction, we note that spatial median splitting performs better in the serial case but has lower scalability. With object median splitting, **BDL** achieves parallel self-relative speedup of up to $35.5\times$, with an average self-relative speedup of $27.2\times$.

Batch Deletion. We measure the time required to delete 10 batches each containing 10% of the points in the data set from an initially full tree for each of our two baselines as well as the BDL-tree. From Figure 11(c), we observe that **B2** has vastly superior performance—it does almost no work other than tombstoning the deleted points so it is extremely efficient. Next, we see that **BDL** has the second-best performance, as it amortizes the rebuilding across the batches, rather than having to rebuild across the entire point set for every delete. Finally, **B1** has the worst performance as it rebuilds on every delete. With object median splitting, **BDL** achieves parallel speedup of up to $33.1\times$, with an average speedup of $28.5\times$.

Data-Parallel k-NN. We measure the performance and scalability of our k-NN implementation as compared to the baselines. As shown in in Figure 11(d), the results show that **B1** and **B2** have similar performance. Furthermore, they are both faster than BDL-tree. This is to be expected, because the k-NN operation is performed directly over the tree after it is constructed over the entire data set in a single batch. Thus, both baselines will consist of fully balanced trees and will be able to perform very efficient k-NN queries. On the other hand, **BDL** consists of a set of multiple trees, which adds overhead to the k-NN operation, as it must be performed separately on each of these individual trees. In Appendix D, we show that when the trees are constructed via a set of batch insertions rather than all at once, the performance of **B2** suffers significantly due to the tree being unbalanced.

Comparison with Zd-tree. We compared with the Zd-tree recently proposed by Blelloch and Dobson [21]. The Zd-tree data structure combines the approach of a kd-tree and Morton

ordering of the data set, and supports parallel batch-dynamic insertions and deletions, and k-NN. The implementation currently only supports 2 and 3 dimensional data sets, whereas our implementation is not restricted to 2 and 3 dimensions. We tested their implementation on 3D-U-10M. Using all threads, their implementation takes 0.12 seconds to construct, and an average of 0.026 and 0.024 seconds for insertion and deletion of 10% of the data points, and takes 1.65 seconds for k-NN. Our BDL-tree implementation is $3.3\times$, $23.1\times$, and $45.83\times$ slower, for construction, insertion, and deletion, respectively, but achieves roughly the same speed for k-NN search. The reason is that the Morton sort used in their implementation is fast and highly optimized for 2 and 3 dimensions; however, extending this technique to higher dimensions would result in overheads due to more bits needed for the Morton ordering.

7   Conclusion

In this paper, we presented ParGeo, a multicore library for computational geometry containing modules for fundamental tasks including kd-tree based spatial search, spatial graph generation, and algorithms in computational geometry. We also presented new parallel algorithms, implementations, and optimizations for convex hull, smallest enclosing ball, and parallel batch-dynamic kd-tree. We performed a comprehensive experimental study showing that our new implementations achieve significant speedups over prior work and obtain high parallel scalability.

References

1   C++ implementation of the 3d quickhull algorithm. `https://github.com/akuukka/quickhull`.

2   The computational geometry algorithms library. `https://www.cgal.org/`.

3   Header only 3d quickhull in c99. `https://github.com/karimnaaji/3d-quickhull`.

4   A header-only C implementation of the quickhull algorithm for building n-dimensional convex hulls and Delaunay meshes. `https://github.com/leomccormack/convhull_3d`.

5   Matlab geometry toolbox for 2d/3d geometric computing. `https://github.com/mattools/matGeom`.

6   Qhull. `http://www.qhull.org/`.

7   Quickhull3d: A robust 3d convex hull algorithm in Java. `https://www.cs.ubc.ca/~lloyd/java/quickhull3d.html`.

8   The Stanford 3d scanning repository. `http://graphics.stanford.edu/data/3Dscanrep/`.

9   Pankaj K. Agarwal, Lars Arge, Andrew Danner, and Bryan Holland-Minkley. Cache-oblivious data structures for orthogonal range searching. In *Proceedings of the Symposium on Computational Geometry*, page 237–245, 2003.

10 A. Aggarwal, B. Chazelle, L. Guibas, C. Ó'Dúnlaing, and C. Yap. Parallel computational geometry. *Algorithmica*, 3(1):293–327, March 1988.

11 Nancy M. Amato and Franco P. Preparata. The parallel 3d convex hull problem revisited. *International Journal of Computational Geometry & Applications*, 2(02):163–173, 1992.

12 Daniel Anderson, Guy E. Blelloch, Laxman Dhulipala, Magdalen Dobson, and Yihan Sun. The problem-based benchmark suite (PBBS), v2. In *Proceedings of the ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP)*, page 445–447, 2022.

13 Lars Arge, Gerth Stølting Brodal, and Rolf Fagerberg. Cache-oblivious data structures. In *Handbook of Data Structures and Applications*, pages 545–565. Chapman and Hall/CRC, 2018.

14 C. Bradford Barber, David P. Dobkin, and Hannu Huhdanpaa. The quickhull algorithm for convex hulls. *ACM Trans. Math. Softw.*, 22(4):469–483, December 1996.

15 Vicente H.F. Batista, David L. Millman, Sylvain Pion, and Johannes Singler. Parallel geometric algorithms for multi-core computers. *Computational Geometry*, 43(8):663–677, 2010.

16 Jon Louis Bentley. Multidimensional binary search trees used for associative searching. *Commun. ACM*, 18(9):509–517, September 1975.

17 Jon Louis Bentley. Decomposable searching problems. *Information Processing Letters*, 8(5):244–251, 1979.

18 Jon Louis Bentley and James B Saxe. Decomposable searching problems I. static-to-dynamic transformation. *Journal of Algorithms*, 1(4):301–358, 1980.

19 Guy E Blelloch. *Vector models for data-parallel computing*, volume 2. MIT Press Cambridge, 1990.

20 Guy E. Blelloch, Daniel Anderson, and Laxman Dhulipala. ParlayLib - a toolkit for parallel algorithms on shared-memory multicore machines. In *Proceedings of the ACM Symposium on Parallelism in Algorithms and Architectures*, page 507–509, 2020.

21 Guy E. Blelloch and Magdalen Dobson. Parallel nearest neighbors in low dimensions with batch updates. In *Proceedings of the Symposium on Algorithm Engineering and Experiments*, pages 195–208, 2022.

22 Guy E. Blelloch, Jeremy T. Fineman, Phillip B. Gibbons, and Julian Shun. Internally deterministic parallel algorithms can be fast. In *Proceedings of the ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 181–192, 2012.

23 Guy E. Blelloch, Yan Gu, Julian Shun, and Yihan Sun. Parallelism in randomized incremental algorithms. 2020.

24 Guy E. Blelloch, Yan Gu, Julian Shun, and Yihan Sun. Randomized incremental convex hull is highly parallel. In *Proceedings of the ACM Symposium on Parallelism in Algorithms and Architectures*, page 103–115, 2020.

25 Paul B. Callahan and S. Rao Kosaraju. Faster algorithms for some geometric graph problems in higher dimensions. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 291–300, 1993.

26 Paul B. Callahan and S. Rao Kosaraju. A decomposition of multidimensional point sets with applications to k-nearest-neighbors and n-body potential fields. *J. ACM*, 42(1):67–90, 1995.

27 Kenneth L. Clarkson and Peter W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom*, 4:387–421, 1989.

28 Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. *Introduction to algorithms*. MIT press, 2009.

29 N. Dadoun and D.G. Kirkpatrick. Parallel construction of subdivision hierarchies. *Journal of Computer and System Sciences*, 39(2):153–165, 1989.

30 Mark de Berg, Marc van Kreveld, Mark Overmars, and Otfried Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, second edition, 2000.

31 Erik D Demaine. Cache-oblivious algorithms and data structures. *Lecture Notes from the EEF Summer School on Massive Data Sets*, 8(4):1–249, 2002.

32 Jerome H Friedman, Jon Louis Bentley, and Raphael Ari Finkel. An algorithm for finding best matches in logarithmic expected time. *ACM Transactions on Mathematical Software (TOMS)*, 3(3):209–226, 1977.

33 Junhao Gan and Yufei Tao. DBSCAN revisited: Mis-claim, un-fixability, and approximation. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, page 519–530, 2015.

34 Mingcen Gao, Thanh-Tung Cao, Ashwin Nanjappa, Tiow-Seng Tan, and Zhiyong Huang. GHull: A GPU algorithm for 3d convex hull. *ACM Trans. Math. Softw.*, 40(1), October 2013.

35 B. Gärtner. Fast and robust smallest enclosing balls. In *European Symposium on Algorithms*, 1999.

36 Jonathan S. Greenfield. A proof for a quickhull algorithm. 1990.

37 Yixin Hu, Qingnan Zhou, Xifeng Gao, Alec Jacobson, Denis Zorin, and Daniele Panozzo. Tetrahedral meshing in the wild. *ACM Trans. Graph.*, 37(4):60:1–60:14, July 2018.

38 Alec Jacobson, Daniele Panozzo, et al. libigl: A simple C++ geometry processing library, 2018. https://libigl.github.io/.

39 Joseph JaJa. *Introduction to Parallel Algorithms*. Addison-Wesley Professional, 1992.

40 David G. Kirkpatrick and John D. Radke. A framework for computational morphology. In *Computational Geometry*, volume 2 of *Machine Intelligence and Pattern Recognition*, pages 217–248. 1985.

41 Thomas Larsson, Gabriele Capannini, and Linus Källberg. Parallel computation of optimal enclosing balls by iterative orthant scan. *Computers & Graphics*, 56:1–10, 2016.

42 D. Lebrun-Grandié, A. Prokopenko, B. Turcksin, and S. R. Slattery. ArborX: A performance portable geometric search library. *ACM Trans. Math. Softw.*, 47(1), December 2020.

43 Marco Livesu. cinolib: a generic programming header only C++ library for processing polygonal and polyhedral meshes. *Transactions on Computational Science XXXIV*, 2019. https://github.com/mlivesu/cinolib/.

44 Kurt Mehlhorn and Stefan Näher. Leda: A platform for combinatorial and geometric computing. *Commun. ACM*, 38(1):96–102, January 1995.

45 S. Näher and Daniel Schmitt. A framework for multi-core implementations of divide and conquer algorithms and its application to the convex hull problem. In *Canadian Conference on Computational Geometry (CCCG)*, 2008.

46 F. P. Preparata and S. J. Hong. Convex hulls of finite sets of points in two and three dimensions. *Commun. ACM*, 20(2):87–93, February 1977.

47 Franco P. Preparata and Michael I. Shamos. *Computational Geometry: An Introduction*. Springer-Verlag, 1985.

48 Octavian Procopiuc, Pankaj K Agarwal, Lars Arge, and Jeffrey Scott Vitter. Bkd-tree: A dynamic scalable kd-tree. In *International Symposium on Spatial and Temporal Databases*, pages 46–65, 2003.

49 Julian Shun, Guy E. Blelloch, Jeremy T. Fineman, and Phillip B. Gibbons. Reducing contention through priority updates. In *Proceedings of the ACM Symposium on Parallelism in Algorithms and Architectures*, page 152–163, 2013.

50 Julian Shun, Guy E. Blelloch, Jeremy T. Fineman, Phillip B. Gibbons, Aapo Kyrola, Harsha Vardhan Simhadri, and Kanat Tangwongsan. Brief announcement: The problem based benchmark suite. In *Proceedings of the ACM Symposium on Parallelism in Algorithms and Architectures*, page 68–70, 2012.

51 Daniel Sieger and Mario Botsch. The polygon mesh processing library, 2020. http://www.pmp-library.org.

52 D Srikanth, Kishore Kothapalli, R Govindarajulu, and P Narayanan. Parallelizing two dimensional convex hull on NVIDIA GPU and Cell BE. In *International Conference on High Performance Computing (HiPC)*, pages 1–5, 2009.

53 Ayal Stein, Eran Geva, and Jihad El-Sana. Cudahull: Fast parallel 3d convex hull on the gpu. *Computers & Graphics*, 36(4):265–271, 2012. Applications of Geometry Processing.

54 Min Tang, Jie yi Zhao, Ruo feng Tong, and Dinesh Manocha. GPU accelerated convex hull computation. *Computers & Graphics*, 36(5):498–506, 2012.

55 Stanley Tzeng and John D Owens. Finding convex hulls using quickhull on the GPU. *arXiv preprint arXiv:1201.2936*, 2012.

56 Yiqiu Wang, Shangdi Yu, Yan Gu, and Julian Shun. Fast parallel algorithms for Euclidean minimum spanning tree and hierarchical spatial clustering. In *Proceedings of the International Conference on Management of Data*, page 1982–1995, 2021.

57 Emo Welzl. Smallest enclosing disks (balls and ellipsoids). In *New Results and New Trends in Computer Science*, pages 359–370, 1991.

A   Convex Hull

Randomized Incremental Algorithm. Our reservation-based algorithm can be used to implement the parallel randomized incremental algorithm. At the start of the algorithm, we randomly permute P. Line 5 will then take a prefix of the remaining points in P to process. On each round, we choose a prefix of $c \cdot numProc$ visible points to perform the reservation, where c is a small constant and $numProc$ is the number of processors. Compared with the existing parallelization approach by Blelloch et al. [24], our approach is much simpler because we avoid the use of complicated data structures.

Quickhull Algorithm. Our reservation-based algorithm also applies to the quickhull algorithm. Specifically, on Line 5 of the algorithm, we select a set of visible points that are furthest from their respective visible facets. The number of visible points that we choose to process on each round is again $c \cdot numProc$.

The 3-dimensional quickhull algorithm is one of the most widely used convex hull algorithms in practice, and so we compare with some of the existing approaches based on quickhull. Since concurrent insertions of visible points creates data races, existing 3 dimensional implementations compromise either correctness or scalability. For Stein et al.'s CudaHull algorithm [53], on each round, the algorithm chooses the furthest point for each facet, and replaces each facet with three new facets, which is done in parallel on the GPU. However, such an approach does not produce a convex polyhedron. Therefore, the implementation uses the CPU to fix the concave artifacts produced via ridge rotation at the end of each round. However, Gao et al. [34] pointed out that certain artifacts cannot be fixed by Stein et al.'s algorithm, leaving concavities in the final polyhedron. Tang et al.'s GPU-based heuristic for convex hull [54]. It first generates a "pseudohull" polyhedron using a quickhull-like algorithm, in which the points are removed. The convex hull is then computed on the remaining points sequentially on the CPU. A clear drawback is that the last phase of the algorithm is not parallel, causing a scalability bottleneck for certain data sets. In comparison, our approach computes a correct convex hull while also achieving high parallel scalability. The correctness is because for each visible point, our algorithm considers all the visible facets and replaces them with new facets. Meanwhile, Stein et al.'s algorithm only considers one visible facet for each visible point, and replaces it directly with three facets, giving rise to concavities during the process.

B   Overhead of Reservation

Our parallel algorithms are work-efficient since it does the same amount of asymptotic work as the sequential counterparts for both the randomized incremental and quickhull algorithms. In addition, given that the expected number of visible facets associated with each visible point is constant [30], the amount of extra work done to perform the reservations on Lines 5–11 is a constant multiplicative factor in each round. When the number of facets is low, instead of using reservations, we process only a single point per round, and we choose the point from the facet that is visible to the most visible points, which maximizes the volume increase of the convex hull.

Figure 12 shows an empirical comparison on the amount of overhead incurred by the reservation algorithm. Specifically, the comparison is between the reservation-based quickhull algorithm and our optimized sequential quickhull algorithm, both running on one thread, for data sets containing 10 million points in 3 dimensions (described in Section 6). The purpose of running on one thread is to measure the amount of work without parallelism. The
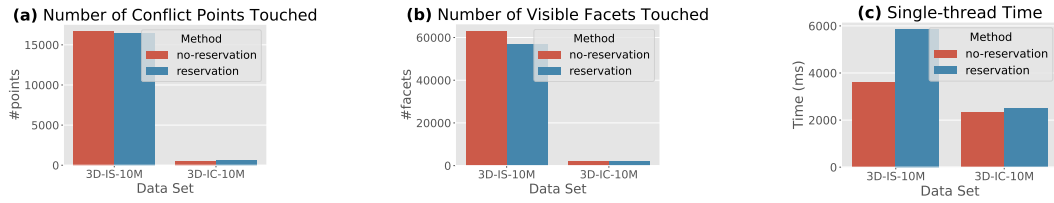
Figure 12 The plots show the overhead of reservation compared with without reservation. (a) and (b) show the number of visible points and facets touched by the algorithms, respectively. (c) shows the single-thread running time of the algorithms.

comparison is based on the number of visible points and facets touched during the algorithm as well as the running time. As we can see from Figures 12(a) and (b), the reservation-based algorithm does not necessarily cause more points or facets to be touched during the algorithm as a majority of the reservations succeed. For example, for the *3D-IS-10M* data set, the number of visible points and facets touched is similar to that of the non-reservation algorithm. For the *3D-IC-10M* data set, the reservation-based approach actually touches fewer visible points and facets, due to the different order in which the visible points are selected between the two algorithms. For both data sets, the reservation algorithm incurs some overhead in doing the work of reservations, as shown by the single-threaded running times in Figure 12(c); however the increase in running time is modest. This overhead is reasonable since it enables parallelism, as we will show in our experiments.

C   Detailed BDL-tree

This section presents more details on operations supported by BDL-tree. We analyze our algorithms using the work-depth model [28, 39]. The **work** of an algorithm is the total number of operations used and the **depth** is the length of the longest sequential dependence (i.e., the parallel running time).

C.1   Static Tree Parallel Algorithms

C.1.1   Parallel vEB Construction

The algorithm for parallel construction of the cache-oblivious kd-tree is shown in Algorithm 1. The function itself is recursive, and so the top level $\textsc{BuildvEB}_\textsc{S}$ function allocates space on line 2 and calls the recursive function $\textsc{BuildvEBRecursive}_\textsc{S}$. Refer to Figure 13 for a graphical representation of this construction.
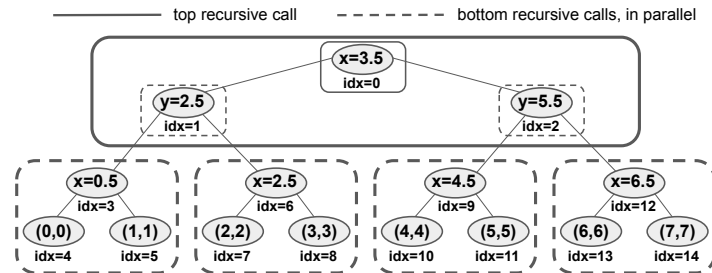
Figure 13 Constructing a vEB kd-tree in parallel over 8 2-dimensional points. Note that the top 3 nodes are placed before the remaining 4 bottom subtrees are built in parallel.

The recursive function $\textsc{BuildvEBRecursive}_\textsc{S}$ maintains state with 5 parameters: a point set Q, a node index idx, a splitting dimension c, the number of levels to build l, and whether it is building the top or bottom of the tree (indicated by t). On line 5, we check for

▢ **Algorithm 1** Parallel vEB-layout kd-tree Construction

Input: Point Set P

Output: kd-tree over P, laid out with the vEB layout on a contiguous memory array of size $2|P| - 1$.

1: **procedure** BUILDVEB$_S$(P)
2:      Allocate $2|P| - 1$ nodes in contiguous memory. The tree nodes will be laid out in this space.
3:      BUILDVEBRECURSIVE$_S$(P, 0, 0, $\lfloor \log(|P|) \rfloor$+1, BOTTOM)
4: **procedure** BUILDVEBRECURSIVE$_S$(Q, idx, c, l, t)
   idx: current node index in the memory array
   c: current dimension to split on
   l: number of levels to build
   t: whether we are building the top or bottom of a tree
5:      If we hit the base case $n = 1$, then we construct a node at idx. If t is TOP, then we perform a parallel median partition on Q in dimension c and record this split as an internal node. Otherwise, we create a leaf node that represents the points in Q.
6:      Compute $l_b = \lceil \lceil \frac{l+1}{2} \rceil \rceil$ and $l_t = l - l_b$ (vEB layout).
7:      Recursively build the top half of the tree with BUILDVEBRECURSIVE$_S$(Q, idx, c, l_t, TOP).
8:      Compute $idx_b = idx + 2^{l_t} - 1$ as the offset where the top half of the tree was just laid out.
9:      Construct the 2^{l_t} lower subtrees in parallel with BUILDVEBRECURSIVE$_S$(Q_i, idx_i, $(c + n_t) \mod d$, l_b, t) where Q_i is the subarray of points that are held by the parent of this subtree and idx_i is the index at which this subtree is to be placed (precomputed with a parallel prefix sum).

---

the base case—if the number of levels to build is 1, then we have to construct a node. If this is the top of a tree, then this node will be an internal node, so we perform a parallel median partition in dimension c and save it as an internal node. On the other hand, if this is the bottom of the tree, we construct a leaf node that holds all the points in Q. Lines 6–9 form the recursive step. In accordance with the exponential layout [9], we have to first construct the top "half" of the tree and then the bottom "half". Therefore, on line 6, we compute the number of levels l_b in the bottom portion as the hyperceiling[1] of $\frac{l+1}{2}$ and the remaining number of levels l_t in the top portion of the tree as $l - l_b$. On line 7, we recursively build the top half of the tree. Then, on line 8, we note that because the top half of the tree is a complete binary tree with l_t levels, it will use $2^{l_t} - 1$ nodes. Therefore, we compute $idx_b = idx + 2^{l_t} - 1$, the node index where the bottom half of the tree should start because the trees are laid out consecutively in memory. Finally, on line 9, we construct each of the 2^{l_t} subtrees that fall under the top half of the tree, each with l_b levels. Each of these trees falls into a distinct segment of memory in the array, and so we can perform this construction in parallel across all of the subtrees by precomputing the starting index idx_i for each of the 2^{l_t} subtrees.

We trace this process on an example in Figure 13, in which BUILDVEB$_S$ is called on a set P of 8 points. This spawns a call to BUILDVEBRECURSIVE$_S$($P[0 : 8]$, 0, 0, 4, BOTTOM). On line 6, we will compute $l_b = 2$ and $l_t = 2$, and on line 7, we spawn a recursive call to

---

[1] The hyperceiling of n, denoted as $\lceil \lceil n \rceil \rceil$ is the smallest power of 2 that is greater than or equal to n, i.e., $2^{\lceil \log n \rceil}$.

BuildvEBRecursive$_S$($P[0:8]$, 0, 0, 2, top). This call is shown as the solid box around the top 3 nodes in Figure 13. In this call, we will hit one further level of recursion before laying out the 3 nodes in indices 0, 1, 2. Then, the original recursive call will proceed to line 8, where it will compute $idx_b = 3$ as the index to begin laying out the $2^{l_t} = 4$ bottom subtrees. Finally, on line 9, we will precompute that the starting indices for the 4 bottom subtrees are $(idx_0, idx_1, idx_2, idx_3) = (3, 6, 9, 12)$. This results in 4 parallel recursive calls, shown in the 4 lower dashed boxes in Figure 13. Each of these recursive calls internally has one more level of recursion to lay out their 3 nodes.

▶ **Theorem 1.** *The cache-oblivious kd-tree with a vEB layout can be constructed over n points in $O(n \log n)$ work and $O(\log n \log \log n)$ depth.*

Proof. The work bound is obtained by observing that there are $O(\log n)$ levels in the fully-constructed tree, and the median partition at each level takes $O(n)$ work, giving a total of $O(n \log n)$ work. For the depth bound, at each recursive step we first build an upper tree with size $O(\sqrt{n})$, and then construct the lower trees in parallel, each with size $O(\sqrt{n})$. Further, we use an $O(\log n)$-depth prefix sum to compute idx_i at every level except the base case and an $O(\log n \log \log n)$-depth median partition in the base case. Overall, this results in $O(\log n \log \log n)$ depth.   ◀

After the vEB-layout kd-tree is constructed, it can be queried as a regular kd-tree—the only difference is the physical layout of the nodes in memory. The correctness of this recursive algorithm can be seen through induction on the number of levels. In particular, we form two inductive hypotheses:

- BuildvEBRecursive$_S$(Q, idx, c, l, top) creates a contiguous, fully-balanced binary tree with l levels rooted at memory location idx. Furthermore, this binary tree consists of internal kd-tree nodes that equally split the point set Q in half at each level.
- BuildvEBRecursive$_S$(Q, idx, c, $\lceil \log |Q| \rceil + 1$, bottom) creates a contiguous kd-tree with l levels rooted at memory location idx.

The base cases, with $l = 1$, for these inductive hypotheses are explicitly given on line 5. Then, the inductive step follows easily by noting that the definition of hyperceiling implies that the recursive calls on line 9 are all sized such that $l_b = \lceil \log |Q|_i \rceil + 1$.

C.1.2 Static Tree Parallel Deletion

The algorithm for parallel deletion from a single kd-tree is shown in Algorithm 2. The function itself is recursive, so the top level Erase$_S$ calls the subroutine EraseRecursive$_S$ on the root node on line 2.

The recursive function EraseRecursive$_S$ acts on one node at a time, represented by the index idx. On line 4, it checks for the base case—if the current node is a leaf node, it simply performs a linear scan to mark any points in the leaf node that are also in Q as deleted. Then, it returns NULL if the entire leaf was emptied; otherwise, it returns the current node idx. Lines 5–7 represent the recursive case. First, on line 5, we perform a parallel partition of Q around the current node's splitting hyperplane. We refer to the lower partition as Q_l and the upper partition as Q_r. On line 6, we recurse on the left and right subtrees in parallel, passing Q_l to the left subtree and Q_r to the right. Finally, line 7 updates the tree structure. We always ensure that every node has 2 children in order to flatten any unnecessary tree traversal. The return value of EraseRecursive$_S$ indicates the node that should take the place of idx in the tree (potentially the same node)—a return value of NULL indicates that the entire subtree rooted at idx was removed. So, if both the left and right child are

---

Algorithm 2 Parallel kd-Tree Deletion

Input: Point Set P

1: **procedure** ERASE$_S$(P)
2:     ERASERECURSIVE$_S$(P, 0)

3: **procedure** ERASERECURSIVE$_S$(Q, idx)
   idx: current node index
4:     If the current node is a leaf node, mark any points in the leaf node that are also in Q as deleted. If all of the points in the current leaf are deleted, return NULL. Otherwise, return the current idx.
5:     Otherwise, perform a parallel partition on Q around the split represented by the current node. Let Q_l, Q_r be the resulting left and right arrays, respectively, after the partition.
6:     Then, recurse on the children in parallel with ERASERECURSIVE$_S$(Q_l, idx_l), ERASERECURSIVE$_S$(Q_r, idx_r), where idx_l and idx_r are the IDs of the left and right children, respectively.
7:     If neither of the recursive calls return NULL, reset the left and right children to be the results of these calls and return the current node. If both of the recursive calls return NULL, return NULL. If one of the recursive calls returns NULL and the other does not, return the non-NULL node.

---

removed, then we can remove the current node as well by returning NULL. On the other hand, if neither the left or right child are removed, then the subtree is still intact, and we simply reset the left and right child pointers of the current node and return the current node idx, indicating that it was not removed. Finally, if exactly one of the children was removed, then we remove the current node as well and let the remaining child connect directly to its grandparent—in this way, we remove an unnecessary internal splitting node. We do this by simply returning the non-NULL child, signaling that it will take the place of the current node in the kd-tree.

▶ **Theorem 2.** *Deleting a batch of B points from a single kd-tree constructed over n points can be done in $O(B \log n)$ work and $O(\log B \log n)$ depth in the worst case.*

Proof. We can see the work bound by noting that each of the B points traverse down $O(\log n)$ levels as part of the algorithm. For the depth, note that in the worst-case the parallel partition at each level operates over $O(B)$ points at each level. Because parallel partition has logarithmic depth, this would result in a worst-case $O(\log B)$ depth at each of the $O(\log n)$ levels, giving the overall depth of $O(\log B \log n)$.   ◀

C.1.3   Data-Parallel k-NN

We execute our k-NN searches in a data-parallel fashion by parallelizing across all of the query points in a batch. The k-NN search for each point is executed serially. We implement a "k-NN buffer", a data structure that maintains a list of the current k-nearest neighbors and provide quick insert functionality to test and insert new points if they are closer than the existing set. The data structure maintains an internal buffer of size $2k$. To insert a point, it simply adds that point to the end of the buffer. If the buffer is filled up, then it uses a serial selection algorithm to partition the buffer around the k-th nearest element and clears out the remaining k elements. This achieves a serial amortized $O(1)$ runtime (because the selection partition step is $O(k)$ and is only performed for every k insertions).

To implement batched k-NN on the kd-tree, we perform a k-NN search for each individual

---

■ **Algorithm 3** Parallel BDL-tree Batch Insertion

1: **procedure** INSERT$_L$(P)
2:      Build an integer bitmask F that represents the static trees within the logarithmic tree structure that are currently filled using 1's, and the trees that are empty using 0's.
3:      Compute $F_{new} = F + \frac{|P|}{X}$, where X is the buffer tree size. This is the new bitmask of trees that should be filled.
4:      Based on the difference between F and F_{new}, determine which trees should be combined into larger trees.
5:      Gather the relevant points and construct all the new trees in parallel using BUILDVEB$_S$ (or BUILD$_S$ for the buffer tree).

---

point in parallel across all the points. We now describe the k-NN method (KNN$_S$) for a single point p. We first allocate a k-NN buffer for the point. Then, we recursively descend through the kd-tree searching for the leaf that p falls into. When we find this leaf, we add all of the points in the leaf to the k-NN buffer. Then, as the recursion unfolds, we check whether the k-NN buffer has k points. If it does not, we add all the points in the sibling of the current node to the k-NN buffer to try to fill up the buffer with nearby points as quickly as possible to improve our estimate of the k-th nearest neighbor. Otherwise, we use the current distance of the k-th nearest neighbor to prune subtrees in the tree. In particular, if the bounding box of the current subtree is entirely contained within the distance of the k-th nearest neighbor, we add all points in the subtree to the k-NN buffer. If the bounding box is entirely disjoint, then we prune the subtree. Finally, if they intersect, we recurse on the subtree.

▶ **Theorem 3.** *For a constant k, k-NN queries over a batch of B points can be performed over a single kd-tree containing n points in worst-case $O(Bn)$ work and worst-case $O(n)$ depth.*

Proof. In the worst-case, we have to search the entire tree, of size $O(n)$, resulting in total work of $O(Bn)$ (due to the amortized $O(1)$ insert cost for k-NN buffers) and depth of $O(n)$, as the queries are done in parallel over the batch, but each search is serial. ◀

As noted by Bentley [16] and Friedman et al. [32], the work for a single nearest-neighbor query on a kd-tree is empirically found to be logarithmic in n, so the experimental runtime and scalability are much better than suggested by the worst-case bounds.

C.2   Parallel Insertion

Insertions are performed in the style of the logarithmic method [17, 18], with the goal of maintaining the minimum number of full trees within BDL-tree. Thus, upon inserting a batch B of points, we rebuild larger trees if it is possible using the existing points and the newly inserted batch. This is implemented as shown in Algorithm 3, and depicted in Figure 7.

First, on line 2, we build a bitmask F of the current set of full static trees in the logarithmic structure. Then, on line 3, because the buffer kd-tree has size X, we can add $|P|/X$ to F to compute a new bitmask F_{new} of full trees that would result if we added $|P|$ points to the tree structure. As an implementation detail, note that we first add $|P|$ mod X points to the buffer kd-tree—if we fill up the buffer kd-tree, then we gather the X points from it and treat them as part of P, effectively increasing the size of P by X. Then, on line 4, taking the bitwise difference between these two bitmasks gives the set of trees that should be consolidated into new larger trees—specifically, any tree that is set in F_{new} but not in F must be constructed from trees that are set in F but not in F_{new}. After determining which trees should be combined into new trees, on line 5 we construct all the new trees in

---

Algorithm 4 Parallel BDL-tree Batch Deletion

Input: Point Set P

1: **procedure** $\text{ERASE}_\text{L}(P)$
2:    In parallel, delete P from each of the underlying trees which is nonempty by calling $\text{ERASE}_\text{S}(P)$ on each of these trees.
3:    In parallel, gather the points from any trees that drop to below half of their original capacity into a set R.
4:    Call $\text{INSERT}_\text{L}(R)$ to reinsert these points into the log-tree structure.

---

parallel—in parallel for each new tree to be constructed, we deconstruct and gather all the points from trees that are being combined into it and then we construct the new tree over these points and any additional required points from P using Algorithm 1.

Refer to Figure 7 for an example of this insertion method (suppose for this example that $X > 2$). In Figure 7(a), the BDL-tree contains X points, giving a bitmask of $F = 1$ (because only the smallest tree is in use). If we insert $X + 1$ points, then we put one node in the buffer tree and compute $F_{new} = 1 + \frac{X}{X} = 2$, and so we have to deconstruct static tree 0 and build static tree 1, as shown in Figure 7(b). Then, if we insert $X + 1$ points again, then we again put one point in the buffer tree and compute $F_{new} = 2 + \frac{X}{X} = 3$, and so we simply construct tree 0 on the X new points (leaving tree 1 intact), as seen in Figure 7(c). Finally, if we then insert $X - 1$ points, we note that this would fill the buffer up, so we take 1 point from the buffer and insert X points; then, $F_{new} = 3 + \frac{X}{X} = 4$, and so we deconstruct trees 0, 1 and construct tree 2, as seen in Figure 7(d).

We refer the readers to Appendix C.3 for the parallel algorithm for batch deletion.

C.3 BDL-tree Parallel Deletion

When deleting a batch of points, the goal is to maintain balance within the subtrees. Thus, if any subtree decreases to less than half of its full capacity, we move all the points down to a smaller subtree in order to maintain balance. As seen in Algorithm 4, this is implemented as a three-step process.

On line 2, we call a parallel bulk erase subroutine on each of the individual trees in parallel in order to actually erase the points from the trees. On line 3, we scan the trees in parallel and collect the points from all trees which have been depleted to less than half of their original capacity. Finally, on line 4, we use the INSERT_L routine to reinsert these points into the structure.

We provide the bounds for parallel insertion and deletion below.

▶ **Theorem 4.** *Given an BDL-tree that was created using only batch insertions and deletions, each batch of B updates takes $O(B \log^2(n+B))$ amortized work and $O(\log(n + B) \log \log(n + B))$ depth, where n is the number of points in the tree before applying the updates.*

Proof. We first argue the work for only performing insertions starting from an empty BDL-tree. In the worst case, points are added to the structure one by one. Then, as similar to the analysis by Bentley [17], the total work incurred is given by noting that the number of times the i'th tree is rebuilt when inserting m points one by one is $O(2^{\log m - i})$. Then, summing the total work gives $O(\sum_{i=0}^{\log m} 2^i i 2^{\log m - i}) = O(m \log^2 m)$, where we use the work bound from Theorem 1. After inserting a batch of size B, we have $n + B$ points in the BDL-tree, and so the amortized work for the batch is $O(B \log^2(n + B))$. Now, if deletions occurred prior to a batch insertion, and the current BDL-tree has n points, there still must have been n previous insertions (since we started with an empty data structure), and so the work of this batch
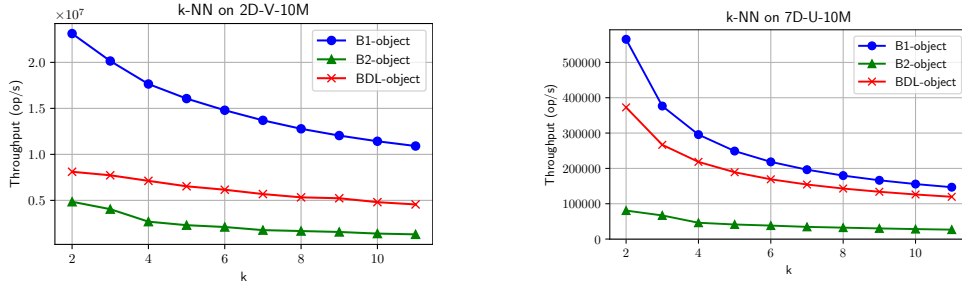
(a) k-NN on 2D-V-10M.

(b) k-NN on 7D-U-10M.

Figure 14 Plots of k-NN throughput (operations per second) vs. k using all 36 cores with hyper-threading, for the 2D-V-10M and 7D-U-10M data sets.

can still be amortized against those n insertions. We now argue the depth bound. When a batch inserted into the tree, the points from smaller trees can be gathered in worst-case $O(\log(n+B))$ depth (if all the points must be rebuilt) and the rebuilding process takes worst case $O(\log(n+B)\log\log(n+B))$ depth, using the result from Theorem 1.

The initial step of deleting the batch of points from each of the underlying kd-trees incurs $O(B\log^2 n)$ work (there are $O(\log n)$ kd-trees, each taking work $O(B\log n)$) and depth $O(\log B \log n)$. Then, collecting the points that need to be reinserted can be done in worst-case depth $O(\log(n+B))$ and the reinsertion takes $O(\log(n+B)\log\log(n+B))$ depth, from before. Overall, the depth is $O(\log(n+B)\log\log(n+B))$. The amortized work for reinserting points in trees that are less than half full is $O(B\log^2 n)$, as every point we reinsert can be charged to a deletion of another point, either from this batch or from a previous batch. This is because for a tree that is half full, there must be at least as many deletions from the tree as the number of points remaining in the tree. ◀

C.4 Data-Parallel k-NN

In the data-parallel k-NN implementation, we parallelize over the set of points given to search for nearest neighbors. First, we allocate a k-NN buffer for each of the points in S. Then, for each of the non-empty trees in BDL-tree, we call the data-parallel k-NN subroutine on the individual tree, passing in the set S of points and the set of k-NN buffers. Because we reuse the same set of k-NN buffers for each underlying k-NN call, we eventually end up with the k-nearest neighbors across all of the individual trees for each point in S.

D k-NN Performance on Incrementally Constructed Trees

In this experiment, we benchmark the throughput of the k-NN operation on 36 cores with hyper-threading as k varies from 2 to 11. For all three trees, we perform the k-NN operation after building the tree from a set of batch insertions, with a batch size of 5% of the data set, until the entire data set is inserted. The results are shown for the 2D VisualVar data set in Figure 14(a) and for the 7D Uniform data set in Figure 14(b). In both data sets, we see that **B1** has the best k-NN performance, followed closely by **BDL**. **B2** has significantly worse performance—this is because the construction of the tree was performed with a set of batch insertions, rather than a single construction over the entire data set, and so the tree ends up unbalanced and the k-NN query performance suffers.