Dynamic Geometric Data Structures via Shallow Cuttings

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

Abstract

We present new results on a number of fundamental problems about dynamic geometric data structures:

1. We describe the first fully dynamic data structures with sublinear amortized update time for maintaining (i) the number of vertices or the volume of the convex hull of a 3D point set, (ii) the largest empty circle for a 2D point set, (iii) the Hausdorff distance between two 2D point sets, (iv) the discrete 1-center of a 2D point set, (v) the number of maximal (i.e., skyline) points in a 3D point set. The update times are near $n^{11/12}$ for (i) and (ii), $n^{7/8}$ for (iii) and (iv), and $n^{2/3}$ for (v). Previously, sublinear bounds were known only for restricted "semi-online" settings [Chan, SODA 2002].
2. We slightly improve previous fully dynamic data structures for answering extreme point queries for the convex hull of a 3D point set and nearest neighbor search for a 2D point set. The query time is $O(\log^2 n)$, and the amortized update time is $O(\log^4 n)$ instead of $O(\log^5 n)$ [Chan, SODA 2006; Kaplan et al., SODA 2017].
3. We also improve previous fully dynamic data structures for maintaining the bichromatic closest pair between two 2D point sets and the diameter of a 2D point set. The amortized update time is $O(\log^4 n)$ instead of $O(\log^7 n)$ [Eppstein 1995; Chan, SODA 2006; Kaplan et al., SODA 2017].

2012 ACM Subject Classification Theory of computation → Computational geometry; Theory of computation → Data structures design and analysis

Keywords and phrases dynamic data structures, convex hulls, nearest neighbor search, closest pair, shallow cuttings

Digital Object Identifier 10.4230/LIPIcs.SoCG.2019.24

Related Version A full version of this paper is available at `https://arxiv.org/abs/1903.08387`.

Funding *Timothy M. Chan*: Work supported in part by NSF Grant CCF-1814026.

Acknowledgements I thank Sariel Har-Peled for discussions on other problems that indirectly led to the results of this paper. Thanks also to Mitchell Jones for discussions on range searching for points in convex position.

1 Introduction

Background. *Dynamic* data structures that can support insertions and deletions of data have been a fundamental topic in computational geometry since the beginning of the field. For example, in 1981 an early landmark paper by Overmars and van Leeuwen [25] presented a fully dynamic data structure for *2D convex hulls* with $O(\log n)$ query time and $O(\log^2 n)$ update time; the $\log^2 n$ bound was later improved in a series of work [7, 6, 12] for various basic types of hull queries, e.g., finding extreme points along given directions.

One of the key results in the area is the author's fully dynamic data structure for *3D convex hulls* [10], which was the first to achieve polylogarithmic query and update time for basic types of hull queries. The original solution required $O(\log^2 n)$ query time for extreme point queries, and $O(\log^6 n)$ amortized update time. (A previous solution by Agarwal and

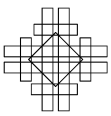

Matoušek [4] had $O(n^\varepsilon)$ query or update time for an arbitrarily small constant $\varepsilon > 0$.) Recently Kaplan et al. [21] noted a small modification of the data structure, improving the update time to $O(\log^5 n)$. The result has numerous applications, including dynamic *2D nearest or farthest neighbor search* (by the standard lifting map). Another application is dynamic *2D bichromatic closest pair* (i.e., computing $\min_{p\in P} \min_{q\in Q} \|p-q\|$ for two planar point sets P and Q) or dynamic *2D diameter* (i.e., computing $\max_{p\in P} \max_{q\in P} \|p-q\|$ for a planar point set P): Eppstein [18] gave a clever, general technique reducing dynamic closest/farthest pair problems to dynamic nearest/farthest neighbor search, which increased the update time by a $\log^2 n$ factor; when combined with the above, this yielded an $O(\log^7 n)$ update time bound.

For many other problems, polylogarithmic update time appears more difficult, and getting sublinear update time is already challenging. For example, in SoCG 2001, the author [8] obtained a dynamic data structure for the *width* of a 2D point set with $O^*(\sqrt{n})$ amortized update time.[1] (Part of the difficulty is that the width problem is neither "decomposable" nor "LP-type".) Sublinear update time is known for a few other assorted geometric problems, such as *dynamic connectivity* for the intersection graph of geometric objects [14].

In SODA 2002, the author [9] explored still more challenging dynamic geometric problems, including maintaining

- **(i)** the number of vertices and facets of a 3D convex hull, or its volume,
- **(ii)** the *largest empty circle* for a 2D point set (with center restricted to be inside a fixed triangle),
- **(iii)** the *Hausdorff distance* for 2D point sets P and Q (i.e., computing $\max_{q\in Q} \min_{p\in P} \|p-q\|$ for two planar point set), and
- **(iv)** the *discrete 1-center* of a 2D point set P (i.e., computing $\min_{q\in P} \max_{p\in P} \|p-q\|$).

The paper [9] obtained sublinear results only for the *insertion-only* case and the *off-line* case (where we are given the entire update sequence in advance), or a generalization of both – the *semi-online* case (as defined by Dobkin and Suri [17], where we are given the deletion time of an element when it is inserted). The update time bounds were $O^*(n^{7/8})$ for (i) and (ii), and $O^*(n^{5/6})$ for (iii) and (iv).

None of these four problems are "decomposable". In particular, problem (i) is nontrivial since known methods such as [10] for 3D convex hull queries do not maintain the global hull explicitly, unlike Overmars and van Leeuwen's original data structure for 2D convex hulls. Problem (ii) also seems to require explicit maintenance of a 3D convex hull (lifted from the 2D farthest-point Voronoi diagram). Problems (iii) and (iv) are max-min or min-max problems, and lack the symmetry of min-min and max-max problems that enable Eppstein's technique. For all these problems, the fully dynamic case has remained open.

New results.

1. We present the first fully dynamic data structures with sublinear update time for Problems (i)–(iv). The amortized update time bounds are $O^*(n^{11/12})$ for (i) and (ii), and $O^*(n^{5/6})$ for (iii) and (iv).

   The approach is general enough to be applicable to many more problems; for example, we can maintain the number of maximal or "skyline" points (points that are not dominated by other points) in a 3D point set in $O^*(n^{2/3})$ amortized time.

[1] Throughout the paper, we use the O^* notation to hide small extra factors that are polylogarithmic, or in some cases, $o(n^\varepsilon)$ for an arbitrarily small constant $\varepsilon > 0$.

2. For basic 3D convex hull queries (e.g., extreme point queries) and 2D nearest neighbor search, as mentioned, Kaplan et al. [21] have lowered the amortized update time of the author's fully dynamic data structure [10], from $O(\log^6 n)$ to $O(\log^5 n)$. We describe a further logarithmic-factor improvement, from $O(\log^5 n)$ to $O(\log^4 n)$.
   Although this improvement is admittedly small, the importance of the result stems from its many applications [10]; for example, we can now compute the *convex (or onion) layers* of a 3D point set in $O(n \log^4 n)$ time, and the *k-level* in an arrangement of planes in 3D in $O(n \log n + f \log^4 n)$ time where f is the output size.

3. For bichromatic closest pair and diameter in 2D, combining Eppstein's technique [18] with the above new result on dynamic nearest neighbor search already gives a slightly improved amortized update time of $O(\log^6 n)$. We describe a further, more substantial improvement that eliminates the two extra logarithmic factors caused by Eppstein's technique [18]. The new update time bound is $O(\log^4 n)$.
   Dynamic bichromatic closest pair has applications to other problems. For example, we can now maintain the Euclidean minimum spanning tree of a 2D point set with $O(\log^6 n)$ amortized update time by using another reduction of Eppstein [18] combined with known results for dynamic minimum spanning trees for graphs [20].

Techniques. The common thread in all of our new methods is the use of *shallow cuttings*: Let H be a set of n hyperplanes in $\mathbb{R}^d$. The *level* of a point q refers to the number of hyperplanes of H strictly below q. A (k, K)*-shallow cutting* is a collection of cells covering all points of level at most k, such that each cell intersects at most K hyperplanes. The *conflict list* H_Δ of a cell Δ refers to the subset of all hyperplanes of H intersecting Δ.

Matoušek [23] proved the existence of shallow cuttings with small number of cells. Specifically, in 3D, the main lemma can be stated as follows:[2]

▶ **Lemma 1.** (Shallow Cutting Lemma) *Given a set H of n planes in $\mathbb{R}^3$ and a parameter $k \in [1, n]$, there exists a $(k, O(k))$-shallow cutting with $O(n/k)$ cells, where each cell is a "downward" tetrahedron containing $(0, 0, -\infty)$. The cutting, together with the conflict lists of all its cells, can be constructed in $O(n \log n)$ time.*

The construction time was first shown by Ramos [26] with a randomized algorithm. Later, Chan and Tsakalidis [15] obtained the first $O(n \log n)$-time deterministic algorithm.

To see how static shallow cuttings may be useful for dynamic geometric data structures, observe that most of the problems considered here are related to the lower envelope of a dynamic set of planes in $\mathbb{R}^3$ (via duality or the standard lifting transformation). Usually, the bottleneck lies in deletions rather than insertions. Basically, a shallow cutting provides a compact implicit representation of the $(\le k)$-level, which is guaranteed to cover the lower envelope even when up to k deletions have occurred.

A further idea behind all our solutions is to classify planes into two types, those that intersect few cells of the shallow cutting, and those that intersect many cells. The latter type of planes may be bad in slowing down updates, but the key observation is that there can't be too many bad elements.

The new sublinear solutions to Problems (i)–(iv), described in Sections 2–3, are obtained by incorporating the shallow cutting idea with the previous techniques from [9], based on periodic rebuilding. The entire solution is conceptually not complicated at all, and the

[2] Matoušek's original formulation in $\mathbb{R}^d$ states the existence of a $(k, n/r)$-shallow cutting with $O(r^{\lfloor d/2 \rfloor}(1 + kr/n)^{\lceil d/2 \rceil})$ cells.

description for Problem (i) fits in under two pages, assuming the availability of known range searching structures. As are typical in other works on data structures with sublinear update time with "funny" exponents, parameters are judiciously chosen to balance several competing costs.

The shallow cutting idea has actually been exploited before in dynamic data structures for basic 3D convex hull queries: Agarwal and Matoušek [4] used shallow cuttings recursively (which caused some loss of efficiency), while the author [10] used a hierarchy of shallow cuttings, for logarithmically many values of k. The above application of shallow cuttings to Problems (i)–(iv) is even more elementary – we only need a single cutting. (This makes it all the more embarassing that the idea was missed till now.)

For basic 3D convex hull queries and 2D nearest neighbor search, our improvement is less innovative. Described in Section 4 (which can be read independently of the previous sections), it is based on the author's original data structure [10], with Kaplan et al.'s logarithmic-factor improvement [21], plus one extra idea to remove a second logarithmic factor: the main observation is that Chan and Tsakalidis's algorithm for shallow cuttings [15] already constructs an entire hierarchy of $O(\log n)$ cuttings in $O(n \log n)$ time, not just a single cutting. However, the hierarchy needed for the data structure in [10] requires some planes be pruned as we go from one cutting to the next, so Chan and Tsakalidis's algorithm cannot be applied immediately. Still, we show that some nontrivial but technical changes (as explained in the appendix) can fix the problem.

For 2D bichromatic closest pair and diameter, our $\log^2 n$-factor improvement, described in Section 5, is a bit more interesting. We still do not know how to improve Eppstein's general reduction [18] from dynamic closest pair to dynamic nearest neighbor search, but intuitively the blind combination of Eppstein's technique with the author's dynamic data structure for 2D nearest neighbor search seems wasteful, since both share some commonalities (both are sophisticated variants of the *logarithmic method* [5], and both handle deletions via re-insertions of elements into smaller subsets). To avoid the redundancy, we show how to directly modify our dynamic data structure for 2D nearest neighbor search to solve the dynamic 2D bichromatic closest pair problem. The resulting modification completely bypasses Eppstein's "conga line" structure [18, 19], and turns out to cause no increase to the $O(\log^4 n)$ bound.

2 Dynamic 3D Convex Hull Size

We begin with our new sublinear-time fully dynamic data structure for maintaining the number of vertices/facets of the convex hull of a dynamic 3D point set. The solution is based on the use of shallow cuttings (Lemma 1) and the author's previous semi-online data structure [9].

▶ **Theorem 2.** *We can maintain the number of vertices, edges, and facets for the convex hull of a dynamic set of n points in $\mathbb{R}^3$, in general position, with $O^*(n)$ preprocessing time and $O^*(n^{11/12})$ amortized insertion and deletion time.*

Proof. It suffices to maintain the number of convex hull facets, which determines the number of vertices and edges (assuming general position). It suffices to compute the number of upper hull facets, since by symmetry we can compute the number of lower hull facets. We describe our solution in dual space, where the problem is to compute the number of vertices in $\mathrm{LE}(H)$ for a dynamic set H of n planes in $\mathbb{R}^3$.

Let k and s be parameters to be set later. We divide the update sequence into phases of k updates each. We maintain a decomposition of the set H into a deletion-only set H_0 and a small set H_{bad} of "bad" planes.

Preprocessing for each phase. At the beginning of each phase, we construct a $(k, O(k))$-shallow cutting Γ of H with $O(n/k)$ cells, together with all their conflict lists, by Lemma 1. We set

$$H_0 = \{h \in H : h \text{ intersects at most } n/s \text{ cells}\} \quad \text{and} \quad H_{\text{bad}} = H - H_0.$$

Since the total conflict list size is $O(n/k \cdot k) = O(n)$, we have $|H_{\text{bad}}| = O(s)$.

Let V_0 and E_0 be the set of vertices and edges of the portion of $\text{LE}(H_0)$ covered by Γ, respectively. There are $O(k)$ such vertices and edges per cell of Γ, and hence, $|V_0|, |E_0| = O(n/k \cdot k) = O(n)$. We preprocess V_0 and E_0 in $O^*(n)$ time by known range searching and intersection searching techniques, so that

- we can count the number of points in V_0 inside a query tetrahedron in $O^*(n^{2/3})$ time (this is 3D simplex range searching) [22, 11, 2];
- we can count the number of line segments in E_0 intersecting a query triangle in $O^*(n^{3/4})$ time (as noted in [9], we can first solve the case of lines and query halfplanes in $\mathbb{R}^3$ using semialgebraic range searching [3] in Plücker space, and then extend the solution for line segments and query triangles by a multi-level data structure [2]).

These data structures can support insertions and deletions of points in V_0 and line segments in E_0 in $O^*(1)$ time each. In addition, we preprocess H_0 in a known dynamic lower envelope data structure in $O^*(n)$ time, to support ray shooting queries in $\text{LE}(H_0)$ in $O^*(1)$ time and deletions in $O^*(1)$ time (e.g., see [4] or Section 4). The total preprocessing time per phase is $O^*(n)$. Amortized over k updates, the cost is $O^*(n/k)$.

Inserting a plane h. We just insert h to the list H_{bad}. Note that $|H_{\text{bad}}| = O(s + k)$ at all times, since there are at most k insertions per phase.

Deleting a plane h from H_{bad}. We just remove h from the list H_{bad}.

Deleting a plane h from H_0. We consider each cell $\Delta \in \Gamma$ intersected by h, and compute $\text{LE}((H_0)_\Delta)$ from scratch in $O(k \log k)$ time (since $|(H_0)_\Delta| = O(k)$). As the number of cells intersected by h is at most n/s, this computation requires $O^*(kn/s)$ total time. The sets V_0 and E_0 undergo at most $O(kn/s)$ changes, and their associated data structures can be updated in $O^*(kn/s)$ time.

Computing the answer. To compute the number of vertices of $\text{LE}(H) = \text{LE}(H_0 \cup H_{\text{bad}})$, we first construct $\text{LE}(H_{\text{bad}})$ in $O((s + k) \log(s + k))$ time, and triangulate all its $O(s + k)$ faces. For each triangle τ in this triangulation:

- we count the number of vertices of V_0 that lie directly below τ, in $O^*(n^{2/3})$ time; and
- we count the number of edges of E_0 that intersect τ, in $O^*(n^{3/4})$ time.

We sum up all these counts. In addition, for each edge of $\text{LE}(H_{\text{bad}})$, we test whether it intersects $\text{LE}(H_0)$ by ray shooting in $O^*(1)$ time, and increment the count if true. For each vertex of $\text{LE}(H_{\text{bad}})$, we test whether it is underneath $\text{LE}(H_0)$ by vertical ray shooting in $O^*(1)$ time, and increment the count if true. Note that $\text{LE}(H)$ is covered by Γ at all times, since there are at most k deletions per phase. The overall count thus gives the answer. The total time to compute the answer is $O^*((s + k)n^{3/4})$.

Analysis. The overall amortized update time is

$$O^*(n/k + kn/s + (s+k)n^{3/4}).$$

The theorem follows by setting $s = k^2$ and $k = n^{1/12}$. ◀

The preprocessing time can be made $O(n \log n)$ and space made $O(n)$ by increasing the update time by an n^ε factor, via known trade-offs for range/intersection searching (with larger-degree partition trees). The method can be deamortized, using existing techniques [24].

The same method can be adapted to maintain the sum or maximum of $f(v)$ over all vertices v of $\mathrm{LE}(H)$, for a general class of functions f. Instead of range counting, we store the set V_0 of points for range sum or range maximum queries (which have similar complexity as range counting). For the set E_0 of line segments, the base level of its multi-level data structure requires data structures $\mathcal{S}_L$ for each canonical subset L of lines in $\mathbb{R}^3$, so that we can return the sum or maximum of $f(\ell \cap h)$ over all $\ell \in L$ for a query plane h in $O^*(|L|^\alpha)$ time, supporting insertions and deletions in L in $O^*(1)$ time. If $\alpha \le 3/4$, the final time bound of our algorithm remains $O^*(n^{11/12})$.

▶ **Theorem 3.** *We can maintain the volume of the convex hull for a dynamic set of n points in $\mathbb{R}^3$, with $O^*(n)$ preprocessing time and $O^*(n^{11/12})$ amortized insertion and deletion time.*

Proof. Let o be a fixed point sufficiently far below all the input points. It suffices to maintain the sum of the volume of the tetrahedra $op_1p_2p_3$ over all upper hull facets $p_1p_2p_3$, since by symmetry we can maintain a similar sum for lower hull facets and subtract. We map each point p to its dual plane h_p. Then the problem fits in the above framework, with $f(v)$ equal to the volume of the tetrahedron $op_1p_2p_3$ for a vertex v defined by the planes $h_{p_1}, h_{p_2}, h_{p_3}$. For a fixed line ℓ defined by the planes h_{p_1} and h_{p_2}, observe that $f(\ell \cap h_p)$ is a linear function over the 3 coordinates of p, since the volume of op_1p_2p can be expressed as a determinant. (This assumes that op_1p_2 is oriented clockwise, which we can ensure at the base level of the multi-level data structure.) Thus, we can implement the base structures $\mathcal{S}_L$ with $\alpha = 0$, by simply summing the 4 coefficients of the associated linear functions over all $\ell \in L$. ◀

▶ **Theorem 4.** *We can maintain the largest empty circle of a dynamic set of n points in $\mathbb{R}^2$, under the restriction that the center lies inside a given triangle Δ_0, with $O^*(n)$ preprocessing time and $O^*(n^{11/12})$ amortized insertion and deletion time.*

Proof. By the standard lifting transformation, map each input point $p = (a, b) \in \mathbb{R}^2$ to the plane h_p with equation $z = -2ax - 2by + a^2 + b^2$ in $\mathbb{R}^3$. Add 3 near-vertical planes along the edges of Δ_0. The largest empty circle problem reduces to finding a vertex $v = (x, y, z)$ of the lower envelope of these planes, maximizing $f(v) = x^2 + y^2 + z$. For a fixed line ℓ, observe that $f(\ell \cap h_{(a,b)})$ is a fixed-degree rational function (ratio of two polynomials) in the 2 variables a and b. We can implement the base structures $\mathcal{S}_L$ with $\alpha = 2/3$, by known techniques for semialgebraic range searching in 3D [4] (applied to the graphs of these bivariate functions). ◀

We can obtain sublinear update time bounds for other similar problems, e.g., maintaining the minimum/maximum-area Delaunay triangle of a dynamic 2D point set. Another application is computing the number of maximal points, also called "skyline points" (which are points not dominated by other points), in a dynamic 3D point set:

▶ **Theorem 5.** *We can maintain the number of maximal points in a dynamic set P of n points in $\mathbb{R}^3$, with $O^*(n)$ preprocessing time and $O^*(n^{2/3})$ amortized insertion and deletion time.*

Proof. The maximal points are vertices of the upper envelope of orthants $(-\infty, a] \times (-\infty, b] \times (-\infty, c]$ over all input points $(a, b, c) \in P$ (the upper envelope is an orthogonal polyhedron). As is well known, an analogue of the shallow cutting lemma holds for such orthants in 3D (in fact, there is a transformation that maps such orthants to halfspaces in 3D); for example, see [13]. The same method can thus be adapted. In fact, it can be simplified. The data structure for V_0 is for orthogonal range searching [16], which has $O^*(1)$ query and update time. The data structure E_0 is not needed. The overall update time becomes

$$O^*(n/k + kn/s + (s + k)).$$

The theorem follows by setting $s = k^2$ and $k = n^{1/3}$. ◀

We can similarly maintain the volume of a union of n boxes in $\mathbb{R}^3$ in the case when all the boxes have a common corner point at the origin (this is called the *hypervolume indicator* problem) with $O^*(n^{2/3})$ update time (previously, an $O^*(\sqrt{n})$ bound was known only in the semi-online setting [9]).

3 Dynamic 2D Hausdorff Distance

The method in Section 2 can also be adapted to solve the dynamic 2D Hausdorff distance problem:

▶ **Theorem 6.** *We can maintain the Hausdorff distance between two dynamic sets P and Q of at most n points in $\mathbb{R}^2$, with $O^*(n)$ preprocessing time and $O^*(n^{8/9})$ amortized insertion and deletion time.*

Proof. By the standard lifting transformation, map each point $p = (a, b) \in P$ to the plane h_p with equation $z = -2ax - 2by + a^2 + b^2$ in $\mathbb{R}^3$. Let H be the resulting set of planes. For each point $q \in Q$, let $\lambda_H(q)$ denote the point on $\mathrm{LE}(H)$ at the vertical line at q. The problem is to find the maximum of $f(\lambda_H(q))$ over all $q \in Q$, where $f(x, y, z) = x^2 + y^2 + z$, for a dynamic set H of at most n planes and a dynamic set Q of at most n points.

Let k and s be parameters to be set later. We divide the update sequence into phases of k updates each. We maintain a decomposition of the set H into a deletion-only set H_0 and a small set H_{bad} of "bad" planes, and a decomposition of the set Q into a deletion-only set Q_0 and a small set Q_{bad} of "bad" points.

Preprocessing for each phase. At the beginning of each phase, we construct a $(k, O(k))$-shallow cutting Γ of H with $O(n/k)$ cells, together with all their conflict lists, by Lemma 1. We further subdivide the cells to ensure that each cell contains at most k points of Q in its xy-projection; this can be done by $O(n/k)$ additional vertical plane cuts, so the number of cells remains $O(n/k)$. We set

$$H_0 = \{h \in H : h \text{ intersects at most } n/s \text{ cells}\} \quad \text{and} \quad H_{\mathrm{bad}} = H - H_0.$$

Since the total conflict list size is $O(n/k \cdot k) = O(n)$, we have $|H_{\mathrm{bad}}| = O(s)$.

We set $Q_0 = Q$. We compute $\lambda_{H_0}(q)$ for all $q \in Q$ in $O(n \log n)$ time. Let Λ_0 be the subset of points in $\{\lambda_{H_0}(q) : q \in Q\}$ covered by Γ. We preprocess the point set Λ_0 in known 3D simplex range searching data structures [22, 11, 2] in $O^*(n)$ time, to support the following queries in $O^*(n^{2/3})$ time:

- compute the maximum of $f(v)$ over all points $v \in \Lambda_0$ inside a query tetrahedron;

- compute the maximum of $f(\lambda_{\{h_p\}}(x,y))$ over all points $v=(x,y,z)\in\Lambda_0$ inside a query tetrahedron for a query plane h_p; note that maximizing $f(\lambda_{\{h_p\}}(x,y))$ is equivalent to maximizing the distance from (x,y) to p (so we can use a 2-level data structure, combining simplex range searching with 2D farthest neighbor searching).

The data structures can support insertions and deletions of points in Λ_0 in $O^*(1)$ time each. In addition, we preprocess H_0 in a known dynamic lower envelope data structure in $O^*(n)$ time, to support ray shooting queries in $\mathrm{LE}(H_0)$ in $O^*(1)$ time and deletions in $O^*(1)$ time (e.g., see [4] or Section 4).

Inserting a plane h to H or a point q to Q. We just insert h to the list H_{bad} or q to the list Q_{bad}. Note that $|H_{\mathrm{bad}}|=O(s+k)$ and $|Q_{\mathrm{bad}}|=O(k)$ at all times.

Deleting a plane h from H_{bad} or a point q from Q_{bad}. We just remove h from the list H_{bad} or q from the list Q_{bad}.

Deleting a point q from Q_0. We just remove $\lambda_{H_0}(q)$ from the set Λ_0 in $O^*(1)$ time.

Deleting a plane h from H_0. We consider each cell $\Delta\in\Gamma$ intersected by h, and compute $\lambda_{(H_0)_\Delta}(q)$ for all $q\in Q$ in the xy-projection of Δ from scratch in $O(k\log k)$ time (since Δ is intersected by $O(k)$ planes in H and contains $O(k)$ points of Q in its xy-projection). As the number of cells intersected by h is at most n/s, this computation takes $O^*(kn/s)$ total time. The set Λ_0 undergoes at most $O(kn/s)$ changes, and its associated data structures can be updated in $O^*(kn/s)$ time.

Computing the answer. To compute the maximum of $f(\lambda_H(q))$ over all $q\in Q$, we first construct $\mathrm{LE}(H_{\mathrm{bad}})$ in $O((s+k)\log(s+k))$ time, and triangulate all its $O(s+k)$ faces. For each triangle τ in this triangulation:

- We compute the maximum of $f(v)$ over all $v=(x,y,z)\in\Lambda_0$ that lie directly below τ, in $O^*(n^{2/3})$ time. Note that for all such v, the $\lambda_H(x,y)=\lambda_{H_0}(x,y)=v$.
- We let h be the plane through τ and compute the maximum of $f(\lambda_{\{h\}}(x,y))$ over all $v=(x,y,z)\in\Lambda_0$ that lie directly above τ, in $O^*(n^{2/3})$ time. Note that for all such v, $\lambda_H(x,y)=\lambda_{\{h\}}(x,y)$.

In addition, for each $q\in Q_{\mathrm{bad}}$, we compute $\lambda_H(q)$ by vertical ray shooting in $\mathrm{LE}(H_0)$ and $\mathrm{LE}(H_{\mathrm{bad}})$ in $O^*(1)$ time; we take the maximum of $f(\lambda_H(q))$ for these points. Note that $\mathrm{LE}(H)$ is covered by Γ at all times, since there are at most k deletions per phase. The overall maximum thus gives the answer. The total time to compute the answer is $O^*((s+k)n^{2/3})$.

Analysis. The overall amortized update time is

$$O^*(n/k+kn/s+(s+k)n^{2/3}).$$

The theorem follows by setting $s=k^2$ and $k=n^{1/9}$. ◀

We can similarly solve the dynamic 2D discrete 1-center problem, by switching lower with upper envelopes and maximum with minimum:

▶ **Theorem 7.** *We can maintain the discrete 1-center of a dynamic set of n points in $\mathbb{R}^2$, with $O^*(n)$ preprocessing time and $O^*(n^{8/9})$ amortized insertion and deletion time.*

It is possible to slightly improve the $O^*(n^{8/9})$ bound to $O^*(n^{5/6})$ in the preceding two theorems: the key observation is that the point set Λ_0 is in convex position, and in the convex-position case, the $O^*(n^{2/3})$ query time for 3D simplex range searching can be improved to $O^*(\sqrt{n})$, as shown by Sharir and Zaban [27]. (The same observation also improves the author's previous $O^*(n^{5/6})$ result to $O^*(n^{3/4})$ in the semi-online setting.)

It remains open whether the dynamic Hausdorff distance and discrete 1-center problem in dimensions $d \geq 3$ can similarly be solved in sublinear time. The author's previous paper [9] gave an $O^*(n^{1-1/(d+1)(\lceil d/2 \rceil+1)})$-time algorithm but only in the semi-online setting. In higher dimensions, the size of shallow cuttings becomes too large for the approach to be effective.

4 Dynamic 3D Convex Hull Queries

In this section, we present a slightly improved data structure for extreme point queries for a dynamic 3D convex hull, by combining the author's previous data structure [10] (as refined by Kaplan et al. [21]) with a modification of Chan and Tsakalidis's algorithm for constructing a hierarchy of shallow cuttings [15].

To describe the latter, we need a definition: Given a set H of n planes in $\mathbb{R}^3$ and a collection Γ_{in} of cells, a Γ_{in}-*restricted* (k, K)-*shallow cutting* is a collection Γ_{out} of cells covering $\{p \in \mathbb{R}^3 : p \text{ is covered by } \Gamma_{\text{in}} \text{ and has level at most } k\}$, such that each cell in Γ_{out} intersects at most K planes. We note that Chan and Tsakalidis's algorithm, with some technical modifications, can prove the following lemma. (The proof requires knowledge of Chan and Tsakalidis's paper, and is deferred to the full paper.)

▶ **Lemma 8.** *There exist constants b, c, and c' such that the following is true: For a set H of at most n planes in $\mathbb{R}^3$ and a parameter $k \in [1, n]$, given a $(-\infty, cbk)$-shallow cutting*[3] *Γ_{in} with at most $c'n/(bk)$ downward cells, together with their conflict lists, we can construct a Γ_{in}-restricted (k, ck)-shallow cutting Γ_{out} with at most $c'n/k$ downward cells, together with their conflict lists, in $O(n + (n/k)\log(n/k))$ deterministic time.*

We now redescribe the author's previous data structure [10] for 3D extreme point queries, with slight changes to incorporate Lemma 8. The redescription uses a recursive form of the logarithmic method [5], which should be a little easier to understand than the original description.

▶ **Theorem 9.** *We can maintain a set of n points in $\mathbb{R}^3$, with $O(n \log n)$ preprocessing time, $O(\log^2 n)$ amortized insertion time, and $O(\log^4 n)$ amortized deletion time, so that we can answer find the extreme point of the convex hull along any query direction in $O(\log^2 n)$ time.*

Proof. We describe our solution in dual space, where we want to answer vertical ray shooting queries for $\text{LE}(H)$, i.e., find the lowest plane of H at a query vertical line, for a dynamic set H of n planes in $\mathbb{R}^3$.

Preprocessing. Our preprocessing algorithm is given by the pseudocode below (ignoring trivial base cases), with the constants b, c, c' from Lemma 8:[4]

preprocess(H):
1. $H_0 = H$, $\Gamma_0 = \{\mathbb{R}^3\}$, $\ell = \log_b n$
2. for $i = 1, \ldots, \ell$ do {
3. $\quad \Gamma_i =$ a Γ_{i-1}-restricted $(n/b^i, cn/b^i)$-shallow cutting of H_{i-1} with at most $c'b^i$ cells
4. $\quad H_i = H_{i-1} - \{h \in H : h \text{ intersects more than } 2cc'\ell \text{ cells of } \Gamma_1 \cup \cdots \cup \Gamma_i\}$
5. $\quad$ for each $\Delta \in \Gamma_i$, compute the conflict list $(H_i)_\Delta$ and initialize $k_\Delta = 0$
}

[3] In a $(-\infty, k)$-shallow cutting, the cells are not required to cover any particular region.
[4] Line 4 is where Kaplan et al.'s improvement lies [21]. The original data structure from [10] basically had $H_i = H_{i-1} - \{h \in H : h \text{ intersects more than } 2cc'\ell \text{ cells of } \Gamma_i\}$.

6. preprocess H_ℓ for static vertical ray shooting
7. $H_{\text{bad}} = H - H_\ell$
8. preprocess(H_{bad})

Note that Γ_{i-1} is a $(-\infty, cn/b^{i-1})$-shallow cutting of H_{i-2}, and consequently a $(-\infty, cn/b^{i-1})$-shallow cutting of H_{i-1}, since $H_{i-1} \subseteq H_{i-2}$. Given Γ_{i-1} and its conflict lists, we can thus apply Lemma 8 to compute Γ_i and its conflict lists, in $O(n + b^i \log b^i)$ time. The total time for lines 1–5 is $O(n \log n + \sum_{i=1}^{\ell} b^i \log b^i) = O(n \log n)$. Line 6 takes $O(n \log n)$ time (by a planar point location method [16]).

We claim that $|H_{\text{bad}}| \le n/2$. To see this, consider each $h \in H_{\text{bad}}$. Let i be the index with $h \in H_{i-1} - H_i$. Then h intersects more than $2cc'\ell$ cells of $\Gamma_1 \cup \cdots \cup \Gamma_i$; send a charge from h to each of these cells. Each cell in Γ_j receives charges only from planes in H_{j-1} that intersect the cell. Thus, the total number of charges is at least $2cc'\ell|H_{\text{bad}}|$ and is at most $\sum_{j=1}^{\ell} cn/b^j \cdot c'b^j = cc'\ell n$. The claim follows. The preprocessing time thus satisfies the recurrence $P(n) \le P(n/2) + O(n \log n)$, which gives $P(n) = O(n \log n)$.

Inserting a plane h. We simply insert h to H_{bad} recursively. When $|H_{\text{bad}}|$ reaches $3n/4$, we rebuild the data structure for H. It takes $\Omega(n)$ updates for a rebuild to occur. The amortized insertion time thus satisfies the recurrence $I(n) \le I(3n/4) + O(P(n)/n) = I(3n/4) + O(\log n)$, which gives $I(n) = O(\log^2 n)$.

Deleting a plane h. The deletion algorithm is as follows:

delete(H, h):
1. for $i = 1, \ldots, \ell$ do
2. for each $\Delta \in \Gamma_i$ with $h \in (H_i)_\Delta$ do {
3. increment k_Δ
4. if $k_\Delta \ge n/b^{i+1}$ then
5. for all $h \in (H_i)_\Delta$ that are still in H but not yet in H_{bad}, insert h to H_{bad}

   }
6. if $h \in H_{\text{bad}}$ then delete(H_{bad}, h)

Let i be the largest index with $h \in H_i$. Then h intersects at most $2cc'\ell = O(\log n)$ cells of $\Gamma_1 \cup \cdots \cup \Gamma_i$. Thus, in each deletion, lines 3–5 are executed $O(\log n)$ times.

In lines 3–5, it takes n/b^{i+1} increments of k_Δ to cause the $|(H_i)_\Delta| \le cn/b^i$ planes to be inserted to H_{bad}. Thus, each increment triggers $O(1)$ amortized number of insertions to H_{bad}, and so a deletion triggers $O(\log n)$ amortized number of insertions to H_{bad}. The amortized deletion time thus satisfies the recurrence $D(n) \le D(3n/4) + O(\log n) I(3n/4) = D(3n/4) + O(\log^3 n)$, which gives $D(n) = O(\log^4 n)$.

Answering the query for a vertical line q. We first answer the query for the static set H_ℓ in $O(\log n)$ time (by planar point location); if the returned plane has already been deleted, ignore the answer. We then recursively answer the query for H_{bad}, and return the lowest of all the planes found. The query time satisfies the recurrence $Q(n) \le Q(3n/4) + O(\log n)$, which gives $Q(n) = O(\log^2 n)$.

Correctness of the query algorithm. To prove correctness, let h^* be the lowest plane at q and $v^* = h^* \cap q$. If $h^* \in H_{\text{bad}}$, correctness follows by induction. So, assume that $h^* \notin H_{\text{bad}}$.

If v^* is covered by Γ_ℓ, say, by the cell $\Delta \in \Gamma_\ell$, then either v^* is on $\text{LE}(H_\ell)$, in which case the algorithm would have correctly found h^*, or some plane in $(H_\ell)_\Delta$ has been deleted from H, in which case all active planes of $(H_\ell)_\Delta$, including h^*, would have been inserted to H_{bad}.

Otherwise, let i be an index such that v^* is not covered by Γ_i but is covered by Γ_{i-1}, say, by the cell $\Delta \in \Gamma_{i-1}$. Since Γ_i is a Γ_{i-1}-restricted $(n/b^i, cn/b^i)$-shallow cutting of H_{i-1}, it follows that v^* must have level more than n/b^i in H_{i-1}. In order for v^* to be the answer, the more than n/b^i planes of H_{i-1} below v^* must have been deleted from H. But then all active planes of $(H_{i-1})_\Delta$, including h^*, would have been inserted to H_{bad}. ◀

By the standard lifting transformation, we obtain:

▶ **Corollary 10.** *We can maintain a set of n points in $\mathbb{R}^2$, with $O(n \log n)$ preprocessing time, $O(\log^2 n)$ amortized insertion time, and $O(\log^4 n)$ amortized deletion time, so that we can answer find the nearest neighbor to any query point in $O(\log^2 n)$ time.*

The space usage in the above data structure is $O(n \log n)$, but can be improved to $O(n)$, by following an idea mentioned in [10] (due to Afshani): instead of storing conflict lists explicitly, generate conflict lists on demand by using a known optimal (static) linear-space data structure for halfspace range reporting [1].

Following [10], we can use the same dynamic data structure to answer other basic types of 3D convex hull queries, e.g., *gift wrapping queries* (finding the two tangents of the hull with a query line outside the hull) in $O(\log^2 n)$ time and *line-intersection queries* (intersecting the hull with a query line) in $O(\log^4 n \log^{O(1)} \log n)$ time. The latter corresponds to *3D linear programming queries* in dual space. The dynamic data structure can be adapted to maintain the *smallest enclosing circle* of a 2D point set. Following [12], the dynamic data structure can also be adapted to answer 3D halfspace range reporting queries.

5 Dynamic 2D Bichromatic Closest Pair

We now adapt the data structure in Section 4 to solve the dynamic 2D bichromatic closest pair problem:

▶ **Theorem 11.** *We can maintain the closest pair between two dynamic sets P and Q of at most n points in $\mathbb{R}^2$, with $O(n \log n)$ preprocessing time, $O(\log^2 n)$ amortized insertion time, and $O(\log^4 n)$ amortized deletion time.*

Proof. By the standard lifting transformation, map each input point $p = (a, b)$ to the plane h_p with equation $z = -2ax - 2by + a^2 + b^2$ in $\mathbb{R}^3$. Let $H = \{h_p : p \in P\}$. For each point $q \in Q$, let $\lambda_H(q)$ denote the point on $\text{LE}(H)$ at the vertical line at q. Let $J = \{h_q : q \in Q\}$. For each point $p \in P$, define $\lambda_J(p)$ similarly. We want to compute the minimum of $f(\lambda_H(q))$ over all $q \in Q$, where $f(x, y, z) = x^2 + y^2 + z$, which is equivalent to the minimum of $f(\lambda_J(p))$ over all $p \in P$.

Preprocessing. We maintain a global heap, whose minimum gives the answer. We modify the preprocess(H) algorithm in Section 4:

preprocess(H, J):

1. run lines 1–7 of the preprocess(H) algorithm on H
2. for each $h_q \in J$, add $f(\lambda_{H_\ell}(q))$ to the heap
3. run lines 1–7 of the preprocess(H) algorithm but with H's replaced by J's
4. for each $h_p \in H$, add $f(\lambda_{J_\ell}(p))$ to the heap
5. preprocess($H_{\text{bad}}, J_{\text{bad}}$)

As in Section 4, the preprocessing time satisfies the recurrence $P(n) \le P(n/2) + O(n \log n)$, which gives $P(n) = O(n \log n)$.

Inserting a plane h_p to H. We recursively insert h_p to H_{bad}. We also compute $\lambda_{J_\ell}(p)$ in $O(\log n)$ time (by planar point location), and add $f(\lambda_{J_\ell}(p))$ to the heap.

When $|H_{\text{bad}}|$ or $|J_{\text{bad}}|$ reaches $3n/4$, we rebuild the data structure for H and J. It takes $\Omega(n)$ updates for a rebuild to occur. The amortized insertion time thus satisfies the recurrence $I(n) \le I(3n/4) + O(\log n) + O(P(n)/n) = I(3n/4) + O(\log n)$, which gives $I(n) = O(\log^2 n)$.

Inserting a plane h_q to J. Symmetric to the above.

Deleting a plane h_p from H. We run lines 1–5 of the delete(H, h) algorithm in Section 4 (with $h = h_p$). In the heap, we remove all entries $f(\lambda_{H_\ell}(q))$ that has $\lambda_{H_\ell}(q) = \lambda_{\{h_p\}}(q)$. If $h_p \in H_{\text{bad}}$, we further recursively delete h_p from H_{bad}. We also remove $f(\lambda_{J_\ell}(p))$ from the heap.

For the analysis, we can charge removals of entries from the heap to their corresponding insertions, by amortization. The amortized deletion time thus satisfies the recurrence $D(n) \le D(3n/4) + O(\log n)I(3n/4) = D(3n/4) + O(\log^3 n)$, which gives $D(n) = O(\log^4 n)$.

Deleting a plane h_q from J. Symmetric to the above.

Correctness. Let p^*q^* be the closest pair with $p^* \in P$ and $q^* \in Q$. If both $h_{p^*} \in H_{\text{bad}}$ and $h_{q^*} \in J_{\text{bad}}$, correctness follows by induction. Otherwise, assume without loss of generality that $h_{p^*} \notin H_{\text{bad}}$. (The case $J_{q^*} \notin J_{\text{bad}}$ is symmetric.) Let $v^* = \lambda_H(q^*)$. The rest of the correctness argument is essentially identical to that in Section 4:

If v^* is covered by Γ_ℓ, say, by the cell $\Delta \in \Gamma_\ell$, then either v^* is on $\text{LE}(H_\ell)$, in which case the algorithm would have included $f(\lambda_H(q^*))$ in the heap, or some plane in $(H_\ell)_\Delta$ has been deleted from H, in which case all active planes of $(H_\ell)_\Delta$, including h_{p^*}, would have been inserted to H_{bad}.

Otherwise, let i be an index such that v^* is not covered by Γ_i but is covered by Γ_{i-1}, say, by the cell $\Delta \in \Gamma_{i-1}$. Since Γ_i is a Γ_{i-1}-restricted $(n/b^i, cn/b^i)$-shallow cutting of H_{i-1}, it follows that v^* must have level more than n/b^i in H_{i-1}. In order for v^* to be the answer, the more than n/b^i planes of H_{i-1} below v^* must have been deleted from H. But then all active planes of $(H_{i-1})_\Delta$, including h_{p^*}, would have been inserted to H_{bad}. ◀

We can similarly solve the diameter problem, by replacing min with max and lower with upper envelopes:

▶ **Theorem 12.** *We can maintain the diameter of a dynamic set of n points in $\mathbb{R}^2$, with $O(n \log n)$ preprocessing time, $O(\log^2 n)$ amortized insertion time, and $O(\log^4 n)$ amortized deletion time.*

References

1. Peyman Afshani and Timothy M. Chan. Optimal halfspace range reporting in three dimensions. In *Proc. 20th ACM-SIAM Sympos. Discrete Algorithms (SODA)*, pages 180–186, 2009. URL: `http://dl.acm.org/citation.cfm?id=1496770.1496791`.
2. Pankaj K. Agarwal. Simplex range searching and its variants: A review. In M. Loebl, J. Nešetril, and R. Thomas, editors, *Journal through Discrete Mathematics*. Springer, to appear.
3. Pankaj K. Agarwal and Jiří Matoušek. On range searching with semialgebraic sets. *Discrete Comput. Geom.*, 11:393–418, 1994. `doi:10.1007/BF02574015`.
4. Pankaj K. Agarwal and Jiří Matoušek. Dynamic half-space range reporting and its applications. *Algorithmica*, 13(4):325–345, 1995. `doi:10.1007/BF01293483`.
5. Jon Louis Bentley and James B. Saxe. Decomposable searching problems I: static-to-dynamic transformation. *J. Algorithms*, 1(4):301–358, 1980. `doi:10.1016/0196-6774(80)90015-2`.
6. Gerth Stølting Brodal and Riko Jacob. Dynamic planar convex hull. In *Proc. 43rd Sympos. Found. Comput. Sci. (FOCS)*, pages 617–626, 2002. `doi:10.1109/SFCS.2002.1181985`.

7 Timothy M. Chan. Dynamic planar convex hull operations in near-logarithmic amortized time. *J. ACM*, 48(1):1–12, 2001. Preliminary version in FOCS 1999. doi:10.1145/363647.363652.

8 Timothy M. Chan. A fully dynamic algorithm for planar width. *Discrete Comput. Geom.*, 30(1):17–24, 2003. Preliminary version in SoCG 2001. doi:10.1007/s00454-003-2923-8.

9 Timothy M. Chan. Semi-online maintenance of geometric optima and measures. *SIAM J. Comput.*, 32(3):700–716, 2003. Preliminary version in SODA 2002. doi:10.1137/S0097539702404389.

10 Timothy M. Chan. A dynamic data structure for 3-d convex hulls and 2-d nearest neighbor queries. *J. ACM*, 57(3):16:1–16:15, 2010. Preliminary version in SODA 2006. doi:10.1145/1706591.1706596.

11 Timothy M. Chan. Optimal partition trees. *Discrete Comput. Geom.*, 47(4):661–690, 2012. Preliminary version in SoCG 2010. doi:10.1007/s00454-012-9410-z.

12 Timothy M. Chan. Three problems about dynamic convex hulls. *Int. J. Comput. Geom. Appl.*, 22(4):341–364, 2012. Preliminary version in SoCG 2011. doi:10.1142/S0218195912600096.

13 Timothy M. Chan, Kasper Green Larsen, and Mihai Pătraşcu. Orthogonal range searching on the RAM, revisited. In *Proc. 27th ACM Sympos. Comput. Geom. (SoCG)*, pages 1–10, 2011. doi:10.1145/1998196.1998198.

14 Timothy M. Chan, Mihai Pătraşcu, and Liam Roditty. Dynamic connectivity: Connecting to networks and geometry. *SIAM J. Comput.*, 40(2):333–349, 2011. Preliminary version in FOCS 2008. doi:10.1137/090751670.

15 Timothy M. Chan and Konstantinos Tsakalidis. Optimal deterministic algorithms for 2-d and 3-d shallow cuttings. *Discrete Comput. Geom.*, 56(4):866–881, 2016. Preliminary version in SoCG 2015. doi:10.1007/s00454-016-9784-4.

16 Mark de Berg, Otfried Cheong, Marc J. van Kreveld, and Mark H. Overmars. *Computational Geometry: Algorithms and Applications*. Springer, 3rd edition, 2008. URL: http://www.worldcat.org/oclc/227584184.

17 David P. Dobkin and Subhash Suri. Maintenance of geometric extrema. *J. ACM*, 38(2):275–298, 1991. doi:10.1145/103516.103518.

18 David Eppstein. Dynamic Euclidean minimum spanning trees and extrema of binary functions. *Discrete Comput. Geom.*, 13:111–122, 1995. doi:10.1007/BF02574030.

19 David Eppstein. Fast hierarchical clustering and other applications of dynamic closest pairs. *ACM Journal of Experimental Algorithmics*, 5:1, 2000. doi:10.1145/351827.351829.

20 Jacob Holm, Eva Rotenberg, and Christian Wulff-Nilsen. Faster fully-dynamic minimum spanning forest. In *Proc. 23rd European Sympos. Algorithms (ESA)*, pages 742–753, 2015. doi:10.1007/978-3-662-48350-3_62.

21 Haim Kaplan, Wolfgang Mulzer, Liam Roditty, Paul Seiferth, and Micha Sharir. Dynamic planar Voronoi diagrams for general distance functions and their algorithmic applications. In *Proc. 28th ACM–SIAM Sympos. Discrete Algorithms (SODA)*, pages 2495–2504, 2017. doi:10.1137/1.9781611974782.165.

22 Jiří Matoušek. Efficient partition trees. *Discrete Comput. Geom.*, 8:315–334, 1992. doi:10.1007/BF02293051.

23 Jiří Matoušek. Reporting points in halfspaces. *Comput. Geom. Theory Appl.*, 2:169–186, 1992. doi:10.1016/0925-7721(92)90006-E.

24 Mark H. Overmars. *The Design of Dynamic Data Structures*, volume 156 of *Lecture Notes in Computer Science*. Springer, 1983. doi:10.1007/BFb0014927.

25 Mark H. Overmars and Jan van Leeuwen. Maintenance of configurations in the plane. *J. Comput. Syst. Sci.*, 23(2):166–204, 1981. doi:10.1016/0022-0000(81)90012-X.

26 Edgar A. Ramos. On range reporting, ray shooting and k-level construction. In *Proc. 15th Sympos. Comput. Geom. (SoCG)*, pages 390–399, 1999. doi:10.1145/304893.304993.

27 Micha Sharir and Shai Zaban. Output-sensitive tools for range searching in higher dimensions. Manuscript, 2013. URL: http://www.cs.tau.ac.il/~michas/shai.pdf.