Convex hulls of spheres and convex hulls of disjoint convex polytopes[☆]

Menelaos I. Karavelas[a,b,*], Raimund Seidel[c], Eleni Tzanaki[a,b]

[a]*University of Crete, Department of Applied Mathematics, GR-714 09 Heraklion, Greece*
[b]*Foundation for Research and Technology - Hellas, Institute of Applied and Computational Mathematics,*
P.O. Box 1385, GR-711 10 Heraklion, Greece
[c]*Fachrichtung Informatik, Universität des Saarlandes, Postfach 151150, D-66041 Saarbrücken, Germany*

Abstract

Given a set Σ of spheres in $\mathbb{E}^d$, with $d \geq 3$ and d odd, having a constant number of m distinct radii $\rho_1, \rho_2, \ldots, \rho_m$, we show that the worst-case combinatorial complexity of the convex hull of Σ is $\Theta(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor \frac{d}{2} \rfloor})$, where n_i is the number of spheres in Σ with radius ρ_i.

To prove the lower bound, we construct a set of $\Theta(n_1 + n_2)$ spheres in $\mathbb{E}^d$, with $d \geq 3$ odd, where n_i spheres have radius ρ_i, $i = 1, 2$, and $\rho_2 \neq \rho_1$, such that their convex hull has combinatorial complexity $\Omega(n_1 n_2^{\lfloor \frac{d}{2} \rfloor} + n_2 n_1^{\lfloor \frac{d}{2} \rfloor})$. Our construction is then generalized to the case where the spheres have $m \geq 3$ distinct radii.

For the upper bound, we reduce the sphere convex hull problem to the problem of computing the worst-case combinatorial complexity of the convex hull of a set of m disjoint d-dimensional convex polytopes in $\mathbb{E}^{d+1}$, where $d \geq 3$ odd, a problem which is of independent interest. More precisely, we show that the worst-case combinatorial complexity of the convex hull of a set of m disjoint d-dimensional convex polytopes in $\mathbb{E}^{d+1}$ is $O(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor \frac{d}{2} \rfloor})$, where n_i is the number of vertices of the i-th polytope. Using the lower bound construction for the sphere convex hull problem, it is also shown to be tight for all odd $d \geq 3$.

Finally, we discuss how to compute convex hulls of spheres with a constant number of distinct radii, or convex hulls of a constant number of disjoint convex polytopes.

Keywords: high-dimensional geometry, discrete geometry, combinatorial geometry, combinatorial complexity, convex hull, spheres, convex polytopes, disjoint polytopes

2010 MSC: 68U05, 52B05, 52B11, 52C45

1. Introduction and results

Let Σ be a set of n spheres in $\mathbb{E}^d$, $d \geq 2$, where the dimension d is considered constant. We call Π a *supporting* hyperplane of Σ if it has non-empty intersection with Σ and Σ is contained in one of the two closed halfspaces bounded by Π. We call H a *supporting halfspace* of the set Σ if it contains all spheres in Σ and is bounded by a supporting hyperplane Π of Σ. The intersection of all supporting halfspaces of Σ is called the convex hull $CH_d(\Sigma)$ of Σ. The definition of convex hulls detailed above is applicable not only to spheres, but also to any finite set of compact geometric objects in $\mathbb{E}^d$. In the case of points, i.e., if we have a set P of n points in $\mathbb{E}^d$, the worst-case combinatorial complexity[1] of $CH_d(P)$ is known to be $\Theta(n^{\lfloor \frac{d}{2} \rfloor})$. Moreover, there exist worst-case

---

[☆]A short version of this paper has appeared in [1].
[*]Corresponding author
 Email addresses: mkaravel@tem.uoc.gr (Menelaos I. Karavelas), rseidel@cs.uni-sb.de (Raimund Seidel), etzanaki@tem.uoc.gr (Eleni Tzanaki)
 [1]In the rest of the paper, and unless otherwise stated, we use the term "complexity" to refer to "combinatorial complexity".

optimal algorithms for constructing $CH_d(P)$ that run in $O(n^{\lfloor \frac{d}{2} \rfloor} + n \log n)$ time, e.g., see [2–6]. Since the complexity of $CH_d(P)$ may vary from $O(1)$ to $\Theta(n^{\lfloor \frac{d}{2} \rfloor})$, a lot of work has been devoted to the design of output-sensitive algorithms for constructing $CH_d(P)$, i.e., algorithms whose running times depend on the size of the output convex hull $CH_d(P)$, e.g., see [7–16]. For a nice overview of the various algorithms for computing the convex hull of points sets, the interested reader may refer to the paper by Erickson [17], while Avis, Bremner and Seidel [18] have a very nice discussion about the effectiveness of output-sensitive convex hull algorithms for point sets.

Results about the convex hull of non-linear objects are very limited. Aurenhammer [19] showed that the worst-case complexity of the power diagram of a set of n spheres in $\mathbb{E}^d$, $d \geq 2$, is $O(n^{\lceil \frac{d}{2} \rceil})$. A direct consequence of this result is that the worst-case complexity of a single additively weighted Voronoi cell or the convex hull of a set of n spheres is $O(n^{\lceil \frac{d}{2} \rceil})$. Rappaport [20] devised an $O(n \log n)$ algorithm for computing the convex hull of a set of discs on the plane, which is worst-case optimal. Boissonnat $et\ al.$ [21] gave an $O(n^{\lceil \frac{d}{2} \rceil} + n \log n)$ algorithm for computing the convex hull of a set of n spheres in $\mathbb{E}^d$, $d \geq 2$, which is worst-case optimal in three and also in even dimensions, since they also showed that the worst-case complexity of the convex hull of n spheres in $\mathbb{E}^3$ is $\Theta(n^2)$. Finally, their results hold true for the case of homothetic convex objects. Boissonnat and Karavelas [22] settled a conjecture in [21]: they proved that the worst-case complexity of the convex hull of a set of n spheres in $\mathbb{E}^d$, $d \geq 2$, is $\Theta(n^{\lceil \frac{d}{2} \rceil})$, which also implied that the algorithm presented in [21] is optimal for all d. As far as output-sensitive algorithms are concerned, Boissonnat, Cérézo and Duquesne [23] showed how to construct the convex hull of a set of n three-dimensional spheres in $O(nf)$ time, where f is the size of the output convex hull, while Nielsen and Yvinec [24] discussed optimal or almost optimal output-sensitive convex hull algorithms for planar convex objects.

In this paper we consider the problem of determining the complexity of the convex hull of a set of spheres, when the spheres have a constant number of distinct radii. This problem has been posed by Boissonnat and Karavelas [22], and it is meaningful for odd dimensions only: in even dimensions the complexity of both the convex hull of n points and the convex hull of n spheres is $\Theta(n^{\lfloor \frac{d}{2} \rfloor}) = \Theta(n^{\lceil \frac{d}{2} \rceil})$, i.e., the two bounds match.

Consider a set of n spheres Σ in $\mathbb{E}^d$, where $d \geq 3$ and d odd, such that the spheres in Σ have a constant number m of distinct radii $\rho_1, \rho_2, \ldots, \rho_m$. Let n_i be the number of spheres in Σ with radius ρ_i. In this paper we prove that the worst-case complexity of $CH_d(\Sigma)$ is $\Theta(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor \frac{d}{2} \rfloor})$. Our result refines the result in [22] for any odd dimension $d \geq 3$. To better explain our bounds, both qualitatively and quantitatively, we first introduce some terminology. We say that ρ_λ $dominates$ Σ if $n_\lambda = \Theta(n)$. We further say that Σ is $uniquely\ (resp.,\ strongly)\ dominated$, if, for some λ, ρ_λ dominates Σ, and $n_i = o(n)$ (resp., $n_i = O(1)$), for all $i \neq \lambda$. Using this terminology, we can express our results as follows. Firstly, if Σ is strongly dominated, then, from the combinatorial complexity point of view, $CH_d(\Sigma)$ behaves as if we had a set of points, or equivalently a set of spheres with the same radius, since in this case the complexity of $CH_d(\Sigma)$ is $\Theta(n^{\lfloor \frac{d}{2} \rfloor})$. If, however, Σ is dominated by at least two radii, the complexity of $CH_d(\Sigma)$ is $\Theta(n^{\lceil \frac{d}{2} \rceil})$, that is $CH_d(\Sigma)$ behaves as in the generic case, where we impose no restriction on the number of distinct radii in Σ. Finally, if Σ is uniquely dominated (but not strongly dominated), the complexity of $CH_d(\Sigma)$ is $o(n^{\lceil \frac{d}{2} \rceil})$ and $\omega(n^{\lfloor \frac{d}{2} \rfloor})$, i.e., it stands in-between the two extremes above: the complexity of $CH_d(\Sigma)$ is asymptotically larger than the case of points (or when we have spheres with the same radius), and asymptotically smaller than the generic case, where we impose no restriction on the number of distinct radii in Σ.

To establish the lower bound for the complexity of $CH_d(\Sigma)$, we construct a set Σ of $\Theta(n_1 + n_2)$ spheres in $\mathbb{E}^d$, for any odd $d \geq 3$, where n_1 spheres have radius ρ_1 and n_2 spheres have radius $\rho_2 \neq \rho_1$, such that worst-case complexity of $CH_d(\Sigma)$ is $\Omega(n_1 n_2^{\lfloor \frac{d}{2} \rfloor} + n_2 n_1^{\lfloor \frac{d}{2} \rfloor})$. This construction is then generalized to sets of spheres having a constant number of $m \geq 3$ distinct radii. More precisely, we construct a set Σ of $n = \sum_{i=1}^{m} n_i$ spheres, where n_i spheres have radius ρ_i, with the ρ_i's being pairwise distinct, such that the worst-case complexity of $CH_d(\Sigma)$ is $\Omega(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor \frac{d}{2} \rfloor})$.

To prove our upper bound we use a lifting map, introduced in [21], that lifts spheres $\sigma_i = (c_i, r_i)$ in $\mathbb{E}^d$ to points $p_i = (c_i, r_i)$ in $\mathbb{E}^{d+1}$. The convex hull $CH_d(\Sigma)$ is then the intersection of the

hyperplane $\{x_{d+1} = 0\}$ with the Minkowski sum of the convex hull $CH_{d+1}(P)$ and the hypercone λ_0, where P is the point set $\{p_1, p_2, \ldots, p_n\}$ in $\mathbb{E}^{d+1}$, and λ_0 is the lower half hypercone with arbitrary apex, vertical axis and angle at the apex equal to $\frac{\pi}{4}$. When the spheres in Σ have a constant number m of distinct radii, the points of P lie on m hyperplanes parallel to the hyperplane $\{x_{d+1} = 0\}$. In this setting, computing the complexity of $CH_d(\Sigma)$ amounts to computing the complexity of the convex hull of m convex disjoint d-polytopes in $\mathbb{E}^{d+1}$. This observation gives rise to the second major result in this paper, which is of independent interest: given a set $\mathscr{P} = \{\mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_m\}$ of m disjoint convex d-polytopes in $\mathbb{E}^{d+1}$, with $d \geq 3$ and d odd, we show that the worst-case complexity of the convex hull $CH_{d+1}(\mathscr{P})$ is $\Theta(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor \frac{d}{2} \rfloor})$, where n_i is the number of vertices of $\mathcal{P}_i$. For our upper bound we make the boundary of $CH_{d+1}(\mathscr{P})$ simplicial by considering its bottom-vertex triangulation. The resulting complex, denoted by $\partial\widehat{\mathcal{P}}$, is a simplicial combinatorial d-sphere, for which we show that the number of its $(k-1)$-faces, $f_{k-1}(\partial\widehat{\mathcal{P}})$, is $O(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\min\{k-1, \lfloor \frac{d}{2} \rfloor\}} + \sum_{i=1}^{m} n_i^{\min\{k, \lfloor \frac{d}{2} \rfloor\}})$, for all $0 \leq k \leq \lfloor \frac{d+1}{2} \rfloor$; for $\lfloor \frac{d+1}{2} \rfloor < k \leq d+1$ the upper bound for $f_{k-1}(\partial\widehat{\mathcal{P}})$ follows directly from the Dehn-Sommerville equations for $\partial\widehat{\mathcal{P}}$. On the other hand, the lower bound for the complexity of $\mathcal{P}$ follows from the lower bound on the complexity of the convex hull of spheres having m distinct radii. For $d \geq 3$ and d odd, our bound constitutes an improvement over the worst-case complexity of convex hulls of point sets, if a single polytope of $\mathscr{P}$ has $\Theta(n)$ vertices, whereas all other polytopes have $o(n)$ vertices (n is the total number of vertices of all m polytopes), while it matches the worst-case complexity of convex hulls of point sets if at least two polytopes have $\Theta(n)$ vertices.

The rest of our paper is structured as follows: In Section 2 we detail our proof of the upper bound on the worst-case complexity of the convex hull of m disjoint d-polytopes in $\mathbb{E}^{d+1}$, while in Section 3 we discuss how to compute this convex hull. In Section 4 we prove our upper bound on the worst-case complexity of the convex hull of a set of spheres. Next we present our lower bound construction for any odd $d \geq 3$ in two steps: first for sphere sets with two distinct radii and then for sphere sets with $m \geq 3$ distinct radii. We end the section by discussing how this lower bound yields a tight lower bound for the problem of the Section 2. In Section 5 we explain how to modify the algorithm by Boissonnat *et al.* [21] so as to almost optimally compute the convex hull of a set of spheres with a constant number of distinct radii. Finally, in Section 6 we summarize our results and state some open problems.

2. Convex hulls of disjoint convex polytopes

A *convex polytope*, or simply *polytope*, $\mathcal{P}$ in $\mathbb{E}^d$ is the convex hull of a finite set of points P in $\mathbb{E}^d$. A polytope $\mathcal{P}$ can equivalently be described as the intersection of all the closed halfspaces containing P. A *face* of $\mathcal{P}$ is an intersection of $\mathcal{P}$ with hyperplanes for which the polytope is contained in one of the two closed halfspaces determined by the hyperplane. The dimension of a face of $\mathcal{P}$ is the dimension of its affine hull. A *k-face* of $\mathcal{P}$ is a k-dimensional face of $\mathcal{P}$. We consider the polytope a *trivial* face of itself; all the other faces are called *proper* faces. We will use the term *d-polytope* to refer to a polytope the trivial face of which is d-dimensional. For a d-polytope $\mathcal{P}$, the 0-faces of $\mathcal{P}$ are its *vertices*, the $(d-2)$-faces of $\mathcal{P}$ are called *ridges*, while the $(d-1)$-faces are called *facets*. For $0 \leq k \leq d$, we denote by $f_k(\mathcal{P})$ the number of k-faces of $\mathcal{P}$. Note that every k-face F of $\mathcal{P}$ is also a k-polytope whose faces are all the faces of $\mathcal{P}$ contained in F.

A polytope is called *simplicial* if all its proper faces are simplices, where a simplex in $\mathbb{E}^d$ is the convex hull of any $0 \leq k \leq d+1$ affinely independent points in $\mathbb{E}^d$.

A *polytopal complex* $\mathcal{C}$ is a finite collection of polytopes in $\mathbb{E}^d$ such that (i) $\emptyset \in \mathcal{C}$, (ii) if $\mathcal{P} \in \mathcal{C}$ then all the faces of $\mathcal{P}$ are also in $\mathcal{C}$ and (iii) the intersection $\mathcal{P} \cap \mathcal{Q}$ for two polytopes in $\mathcal{C}$ is a face of both $\mathcal{P}$ and $\mathcal{Q}$. The dimension $\dim(\mathcal{C})$ of $\mathcal{C}$ is the largest dimension of a polytope in $\mathcal{C}$. A polytopal complex is called *pure* if all its maximal (with respect to inclusion) faces have the same dimension. We will use the term *d-complex* to refer to a pure polytopal complex whose maximal faces are d-dimensional. A polytopal complex is simplicial if all its faces are simplices. One important class of polytopal complexes arise from polytopes. More precisely, a d-polytope $\mathcal{P}$,

together with all its faces and the empty set, form a polytopal d-complex, denoted by $\mathcal{C}(\mathcal{P})$. The only maximal face of $\mathcal{C}(\mathcal{P})$ is the polytope $\mathcal{P}$ itself. Moreover, all proper faces of $\mathcal{P}$ form a pure polytopal complex, called the *boundary complex* $\mathcal{C}(\partial\mathcal{P})$. The maximal faces of $\mathcal{C}(\partial\mathcal{P})$ are just the facets of $\mathcal{P}$, and its dimension is $\dim(\mathcal{P}) - 1 = d - 1$.

The f-vector $(f_{-1}(\mathcal{P}), f_0(\mathcal{P}), \ldots, f_{d-1}(\mathcal{P}))$ of a d-polytope $\mathcal{P}$ is defined as the $(d+1)$-dimensional vector consisting of the number $f_k(\mathcal{P})$ of k-faces of $\mathcal{P}$, $-1 \leq k \leq d$, where $f_{-1}(\mathcal{P}) = 1$ refers to the empty set. The h-vector $(h_0(\mathcal{P}), h_1(\mathcal{P}), \ldots, h_d(\mathcal{P}))$ of a d-polytope $\mathcal{P}$ is defined as the $(d+1)$-dimensional vector, where $h_k(\mathcal{P}) := \sum_{i=0}^{k} (-1)^{k-i} \binom{d-i}{d-k} f_{i-1}(\mathcal{P})$ for $0 \leq k \leq d$. Thus the h-vector is a linear transform of the f-vector. It turns out that this transform is invertible and the f vector can be expressed as $f_{i-1}(\mathcal{P}) = \sum_{k=0}^{k=d} \binom{d-k}{i-k} h_k(\mathcal{P})$ for $0 \leq i \leq d$. For simplicial polytopes the elements of the f-vector are not independent; they satisfy the so-called *Dehn-Sommerville equations*, which can be written in a very concise form in terms of the h-vector of $\mathcal{P}$: $h_k(\mathcal{P}) = h_{d-k}(\mathcal{P})$, $0 \leq k \leq d$. An important implication of the existence of the Dehn-Sommerville equations is that if we know the face numbers $f_k(\mathcal{P})$ for all $0 \leq k \leq \lfloor \frac{d}{2} \rfloor - 1$, we can determine the remaining face numbers $f_k(\mathcal{P})$ for all $\lfloor \frac{d}{2} \rfloor \leq k \leq d - 1$.

A *simplicial combinatorial d-sphere (resp., d-ball)*, or, simply, *simplicial d-sphere (resp., d-ball)*, is a simplicial complex that is homeomorphic to the d-dimensional sphere (resp., ball). The boundary complex of a simplicial d-polytope is a simplicial $(d-1)$-sphere; the converse is not true in general: there are simplicial 4-spheres that are not *polytopal* (i.e., not realizable as boundary complexes of polytopes). What is of interest for this paper, however, are two facts about simplicial spheres (cf. [25, 26]):

(1) They satisfy the Dehn-Sommerville equations, i.e., for any simplicial $(d-1)$-sphere S we have $h_k(S) = h_{d-k}(S)$, $0 \leq k \leq d$.
(2) The Upper Bound Conjecture holds for simplicial spheres (thus becoming the Upper Bound Theorem for simplicial spheres). More precisely, given an n-vertex simplicial $(d-1)$-sphere S, then for all $-1 \leq k \leq d-1$, we have $f_k(S) \leq f_k(C_d(n)) = O(n^{\min\{k+1, \lfloor \frac{d}{2} \rfloor\}})$, where $C_d(n)$ stands for the cyclic d-polytope with n vertices.

For a d-polytope Q its *bottom-vertex triangulation* $\widehat{Q}$ is a simplicial complex defined on the vertex set of Q as follows (see [27]): If $d \leq 1$ then $\widehat{Q} = Q$. If $d > 1$ let v be the "lowest" vertex of Q (assume that Q is oriented such that all vertices are at distinct "heights"); for each facet F of Q that does not contain v consider each $(d-1)$-simplex Δ in its bottom-vertex triangulation $\widehat{F}$ and include the d-simplex spanned by Δ and v (along with its faces) in $\widehat{Q}$. It is well known that $\widehat{Q}$ forms a simplicial d-ball and its boundary complex $\partial\widehat{Q}$ constitutes a simplicial $(d-1)$-sphere.

Let $\mathscr{P} = \{\mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_m\}$ be a set of m disjoint d-polytopes in $\mathbb{E}^{d+1}$, where $m \geq 2$ and m is constant. We denote by P_i the set of vertices of $\mathcal{P}_i$, by n_i the cardinality of P_i, and by P the union $P = P_1 \cup P_2 \cup \ldots \cup P_m$. We are interested in the number of faces of the bottom-vertex triangulation $\widehat{\mathcal{P}}_i$.

Lemma 1. *For all k, with $0 \leq k \leq d+1$, we have $f_{k-1}(\widehat{\mathcal{P}}_i) = O(n_i^{\min\{k, \lfloor \frac{d}{2} \rfloor\}})$.*

Proof. First note that since $\partial\widehat{\mathcal{P}}_i$ is a simplicial $(d-1)$-sphere with n_i vertices the Upper Bound Theorem implies that $f_{k-1}(\partial\widehat{\mathcal{P}}_i) = O(n_i^{\min\{k, \lfloor \frac{d}{2} \rfloor\}})$ for $0 \leq k \leq d$.

For $k \leq 1$ the claim of the lemma is trivial. For $k > 1$ each $(k-1)$-face of $\widehat{\mathcal{P}}_i$ is either in the boundary complex $\partial\widehat{\mathcal{P}}_i$ or not. The number of such boundary faces is

$$f_{k-1}(\partial\widehat{\mathcal{P}}_i) = O(n_i^{\min\{k, \lfloor \frac{d}{2} \rfloor\}}).$$

Each non-boundary $(k-1)$-face is spanned by the bottom-vertex v of $\mathcal{P}_i$ and a unique $(k-2)$-face in $\partial\widehat{\mathcal{P}}_i$. Thus the number of such non-boundary $(k-1)$-faces is bounded by

$$f_{k-2}(\partial\widehat{\mathcal{P}}_i) = O(n_i^{\min\{k-1, \lfloor \frac{d}{2} \rfloor\}}) = O(n_i^{\min\{k, \lfloor \frac{d}{2} \rfloor\}}),$$

which completes the proof. $\qquad\square$

Let $\mathcal{P} = CH_{d+1}(P)$, and let $\partial\widehat{\mathcal{P}}$ be the simplicial d-complex formed by constructing the bottom-vertex triangulation of $\partial\mathcal{P}$. Clearly, for all $0 \le k \le d+1$, $f_{k-1}(\mathcal{P}) = f_{k-1}(\partial\mathcal{P}) \le f_{k-1}(\partial\widehat{\mathcal{P}})$, so in order to bound the number of faces of $\mathcal{P}$, it suffices to bound the number of faces of $\partial\widehat{\mathcal{P}}$.

Lemma 2. *For all $0 \le k \le d+1$, we have*

$$f_{k-1}(\partial\widehat{\mathcal{P}}) = O\left(\sum_{1 \le i \neq j \le m} n_i n_j^{\min\{k-1, \lfloor \frac{d}{2} \rfloor\}} + \sum_{i=1}^{m} n_i^{\min\{k, \lfloor \frac{d}{2} \rfloor\}}\right).$$

Proof. The bound trivially holds for $k = 0$. Below, we will only consider positive values for k. Furthermore, we will assume, without loss of generality, that $n_1 \ge n_2 \ge \cdots \ge n_m$.

Since $\partial\widehat{\mathcal{P}}$ is a simplicial d-sphere, it suffices to bound the number of $(k-1)$-faces of $\partial\widehat{\mathcal{P}}$ for all $0 \le k \le \lfloor \frac{d+1}{2} \rfloor$; for k with $\lfloor \frac{d+1}{2} \rfloor < k \le d+1$, the bounds follow from the Dehn-Sommerville equations for $\partial\widehat{\mathcal{P}}$.

Let F be a $(k-1)$-face of $\partial\widehat{\mathcal{P}}$. Since $\partial\widehat{\mathcal{P}}$ is simplicial, F is a $(k-1)$-simplex, i.e., it is defined by k vertices in P. Moreover, F intersects each $\widehat{\mathcal{P}}_i$ in a (k_i-1)-face with k_i vertices. This immediately gives the following trivial combinatorial upper bound:

$$f_{k-1}(\partial\widehat{\mathcal{P}}) \le \sum_{k_1+k_2+\cdots+k_m=k} \prod_{i=1}^{m} f_{k_i-1}(\widehat{\mathcal{P}}_i). \tag{1}$$

Let $K = (k_1, k_2, \ldots, k_m)$, $|K| = \sum_{i=1}^{m} k_i$, and denote by $\dim(K)$ the number of non-zero elements of K. Using this notation, relation (1) can be rewritten as:

$$f_{k-1}(\partial\widehat{\mathcal{P}}) \le \sum_{\substack{(0,\ldots,0) \preccurlyeq K \preccurlyeq (\lfloor \frac{d+1}{2} \rfloor, \ldots, \lfloor \frac{d+1}{2} \rfloor) \\ |K|=k}} \prod_{i=1}^{m} f_{k_i-1}(\widehat{\mathcal{P}}_i), \tag{2}$$

where the notation $A \preccurlyeq B$ means that each coordinate of A is smaller or equal than the corresponding coordinate of B. We consider each term in the right-hand side sum of (2) individually, and, in particular, we distinguish between the case where K consists of a single positive element (i.e., $\dim(K) = 1$), and the case where K consists of at least two positive elements (i.e., $\dim(K) \ge 2$).

$\dim(K) = 1$. Let $k_j > 0$, whereas $k_i = 0$, for all $i \neq j$. Clearly, in this case, $k_j = k$. Since $f_{-1}(\widehat{\mathcal{P}}_i) = 1$, $1 \le i \le m$, and recalling that d is odd, we have:

$$\prod_{i=1}^{m} f_{k_i-1}(\widehat{\mathcal{P}}_i) = f_{k_j-1}(\widehat{\mathcal{P}}_j) = f_{k-1}(\widehat{\mathcal{P}}_j) = O(n_j^{\min\{k, \lfloor \frac{d}{2} \rfloor\}}) = O(n_1^{\min\{k, \lfloor \frac{d}{2} \rfloor\}}),$$

where the last two equalities above come from Lemma 1 and the fact that $n_j \le n_1$, for all $j \ge 1$.

$\dim(K) \ge 2$. Let $k_{j_1}, k_{j_2} > 0$, with $j_1 < j_2$. Clearly, $j_2 \ge 2$, which means that $n_{j_2} \le n_2 \le n_1$. In this case we have $f_{k_i-1}(\widehat{\mathcal{P}}_i) = O(n_i^{k_i}) = O(n_1^{k_i})$, for all $i \neq j_2$. Moreover, $f_{k_{j_2}-1}(\widehat{\mathcal{P}}_{j_2}) = O(n_{j_2}^{k_{j_2}}) = O(n_{j_2}^{k_{j_2}-1} n_{j_2}) = O(n_1^{k_{j_2}-1} n_2)$. Hence:

$$\prod_{i=1}^{m} f_{k_i-1}(\widehat{\mathcal{P}}_i) = f_{k_{j_2}-1}(\widehat{\mathcal{P}}_{j_2}) \cdot \prod_{i \neq j_2} f_{k_i-1}(\widehat{\mathcal{P}}_i) = O(n_1^{k_{j_2}-1} n_2) \cdot \prod_{i \neq j_2} O(n_1^{k_i})$$

$$= O(n_1^{|K|-1} n_2) = O(n_1^{k-1} n_2) = O(n_1^{\min\{k-1, \lfloor \frac{d}{2} \rfloor\}} n_2),$$

where we used the fact that $|K| = k$, and that $k - 1 \le \lfloor \frac{d+1}{2} \rfloor - 1 = \lfloor \frac{d}{2} \rfloor$ (since d is odd).

5

We can now split the right-hand side sum in (2) in two parts: the terms for which $\dim(K) = 1$ and the terms for which $\dim(K) \geq 2$. Using the bounds derived above for each term in the sum, and noting that since m is constant the number of terms in the right-hand side sum in (2) is also constant, we deduce that

$$f_{k-1}(\partial \widehat{\mathcal{P}}) = O(n_1^{\min\{k-1, \lfloor \frac{d}{2} \rfloor\}} n_2) + O(n_1^{\min\{k, \lfloor \frac{d}{2} \rfloor\}})$$

$$= O \left(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\min\{k-1, \lfloor \frac{d}{2} \rfloor\}} + \sum_{i=1}^{m} n_i^{\min\{k, \lfloor \frac{d}{2} \rfloor\}} \right). \qquad \square$$

By Lemma 2, and the fact that the number of faces of $\partial \mathcal{P}$ is bounded from above by the number of faces of $\partial \widehat{\mathcal{P}}$, we deduce that the worst-case complexity of the convex hull $CH_{d+1}(\mathscr{P})$ is $O(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor \frac{d}{2} \rfloor})$. As we will see in Subsection 4.4 (see Corollary 12), this bound is asymptotically tight for any odd $d \geq 3$. Hence:

Theorem 3. *Let $\mathscr{P} = \{\mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_m\}$ be a set of a constant number of $m \geq 2$ disjoint d-polytopes in $\mathbb{E}^{d+1}$, where $d \geq 3$ and d is odd. The worst-case complexity of the convex hull $CH_{d+1}(\mathscr{P})$ is $\Theta(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor \frac{d}{2} \rfloor})$, where $n_i = f_0(\mathcal{P}_i)$, $1 \leq i \leq m$.*

Remark 4. *The proof of Lemma 2, and thus the upper bound in Theorem 3, still holds under much weaker assumptions on the polytopes $\mathcal{P}_i$. Their dimension can be at most d, instead of exactly d, and they can even intersect arbitrarily, as long as the intersection of a face of $\mathcal{P}$ with a face of some $\mathcal{P}_i$ is a face of both $\mathcal{P}$ and $\mathcal{P}_i$ (this is, for example, the case if the m polytopes in $\mathscr{P}$ form a polytopal complex in $\mathbb{E}^{d+1}$ of dimension at most d).*

3. Computing the convex hull of disjoint convex polytopes

In view of Theorem 3, when we have two d-polytopes $\mathcal{P}_i$ and $\mathcal{P}_j$, where $d \geq 3$ and d odd, such that $n_i = \Theta(n)$ and $n_j = \Theta(n)$, respectively, we cannot compute $CH_{d+1}(\mathscr{P})$ faster than the worst-case optimal algorithm by Chazelle [6], which, in our setting, runs in $O(n^{\lfloor \frac{d+1}{2} \rfloor})$ time. If this is not the case, however, it might pay off to use an output-sensitive algorithm for constructing the convex hull of the point set P formed by the vertices of the $\mathcal{P}_i$'s. In Table 1 we summarize the various convex hull algorithms that are applicable in our case, and we report on their asymptotic complexity, both in the generic setting, as well as the case where we have a constant number m of disjoint d-polytopes in $\mathbb{E}^{d+1}$. In the first four rows of the table we focus on 3-polytopes in $\mathbb{E}^4$. In rows 5 to 7 of the table we consider both deterministic and randomized algorithms that can be used for any $d \geq 3$ odd, whereas in the last two rows of the table we have improved bounds for $d \geq 5$ odd, for the algorithms reported on the first two rows of the table.

Below we distinguish between 3-polytopes in $\mathbb{E}^4$ and d-polytopes in $\mathbb{E}^{d+1}$, where $d \geq 5$ and d odd. In the rest of the section f will denote the number of facets of the output convex hull computed, whereas F will denote the quantity $F = \sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor \frac{d}{2} \rfloor}$. Moreover, for simplicity of notation, we will use α to denote the quantity $\lfloor \frac{d+1}{2} \rfloor$; the case $d = 3$ is, thus, equivalent to $\alpha = 2$, whereas the case $d \geq 5$ and d odd is equivalent to $\alpha \geq 3$. Finally, in our analysis below we will assume, without loss of generality, that $n_1 \geq n_2 \geq \cdots \geq n_m$; under this assumption, and given that m is constant, we have that $n_1 = \Theta(n)$, while for F we have: $F = \Theta(n_2 n_1^{\lfloor \frac{d}{2} \rfloor}) = \Theta(n_2 n^{\lfloor \frac{d}{2} \rfloor}) = \Theta(n_2 n^{\alpha-1})$.

Three-dimensional polytopes in $\mathbb{E}^4$. One of the earliest algorithms is Seidel's shelling algorithm [9] that runs in $O(n^2 + f \log n)$ time. The preprocessing step of Seidel's algorithm was later on improved by Matoušek and Schwarzkopf [11], resulting in an $O(n^{2-2/(\alpha+1)+\epsilon} + f \log n)$ time algorithm, for any fixed $\epsilon > 0$, which for $\alpha = 2$ gives an $O(n^{4/3+\epsilon} + f \log n)$ time algorithm. Chan, Snoeyink and Yap [16] describe a divide-and-conquer algorithm for constructing four-dimensional convex hulls in $O((n+f) \log^2 f)$ time. Finally, Chan [15] improved the gift-wrapping algorithm of

Algorithm/Reference	$\alpha = \lfloor \frac{d+1}{2} \rfloor$	Complexity type	Time (general)	Time (our case)
Seidel [9]	2	Worst-case	$O(n^2 + f \log n)$	$O(n^2 + F \log n)$
Matoušek and Schwarzkopf [11]	2	Worst-case	$O(n^{4/3+\epsilon} + f \log n)$	$O(n^{4/3+\epsilon} + F \log n)$
Chan, Snoeyink and Yap [16]	2	Worst-case	$O((n + f) \log^2 f)$	$O(F \log^2 n)$
Modified gift-wrapping (this paper)	2	Worst-case	N/A	$O(F \log n) \{\star\}$
Chazelle [6]	≥ 2	Worst-case	$O(n^\alpha)$	$O(n^\alpha)$
Chan [15]	≥ 2	Worst-case	$O(n \log f + (nf)^{\alpha/(\alpha+1)} \log^{O(1)} n)$	$O((nF)^{\alpha/(\alpha+1)} \log^{O(1)} n)$
Clarkson and Shor [28]	≥ 2	Expected	$O(n^\alpha)$	$O(F) \{\star\}$
Seidel [9]	≥ 3	Worst-case	$O(n^2 + f \log n)$	$O(F \log n)$
Matoušek and Schwarzkopf [11]	≥ 3	Worst-case	$O(n^{2-2/(\alpha+1)+\epsilon} + f \log n)$	$O(F \log n)$

Table 1: The various algorithms that can be applied in our setting. The dimension d is always at least 3 and odd, and n always denotes the number of points for which the convex hull is computed. The last two columns display the time complexities of the various algorithms considered in the generic case (no restrictions on the points' configuration), and the case of disjoint convex polytopes (our case), respectively. The algorithms and complexities in the first three rows are specific to four dimensions ($d = 3$ or, equivalently, $\alpha = 2$), the algorithms and complexities in the middle four rows are applicable for any $d \geq 3$ odd ($\alpha \geq 2$), whereas the algorithms and complexities in the lower two rows are valid for $d \geq 5$ only (i.e., for $\alpha \geq 3$). For output-sensitive algorithms, f is the number of facets of the output convex hull, whereas in the last column F denotes the quantity $F = \sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor \frac{d}{2} \rfloor} = \sum_{1 \leq i \neq j \leq m} n_i n_j^{\alpha-1}$. Note that $\alpha \geq 2$, $F = \Omega(n^{\lfloor \frac{d}{2} \rfloor}) = \Omega(n^{\alpha-1})$, and $F = O(n^{\lfloor \frac{d+1}{2} \rfloor}) = O(n^\alpha)$. The complexities marked as $\{\star\}$ in the last column either refer to an algorithm presented in this paper or to an analysis performed in this paper.

Chand and Kapur [7], yielding an $O(n \log f + (nf)^{1-1/(\alpha+1)} \log^{O(1)} n)$ time algorithm, which, for $\alpha = 2$ has time complexity $O(n \log f + (nf)^{2/3} \log^{O(1)} n)$ time. In the disjoint 3-polytopes setting, we have $f = O(F) = O(n_2 n)$, which yields the complexities shown in rows 1–3 and 5–6 of the last column in Table 1. In particular, for Chan's algorithm [15], we have the following bound for its time complexity:

$$
\begin{aligned}
O(n \log F + (nF)^{1-1/(2+1)} \log^{O(1)} n) &= O(n \log (n_2 n) + (n^2 n_2)^{2/3} \log^{O(1)} n) \\
&= O(n \log n + n^{4/3} n_2^{2/3} \log^{O(1)} n) \\
&= O(n^{4/3} n_2^{2/3} \log^{O(1)} n).
\end{aligned}
$$

Among the output-sensitive algorithms discussed above, it is clear that the algorithm by Matoušek and Schwarzkopf has better time complexity than by that Seidel. The algorithm by Chan, as well as that by Chan, Snoeyink and Yap, can yield better asymptotic complexities than Matoušek and Schwarzkopf's algorithm for a certain range for the size of n_2 (e.g., for $n_2 = O(1)$, Matoušek and Schwarzkopf's algorithm has complexity $O(n^{4/3+\epsilon})$, Chan's algorithm has complexity $O(n^{4/3} \log^{O(1)} n)$, while the algorithm of Chan, Snoeyink and Yap has complexity $O(n \log^2 n)$). However, it is always the case that the asymptotic complexity of Chan, Snoeyink and Yap's algorithm is better than that of Chan's algorithm.

In what follows we describe a custom modification of Chand and Kapur's gift-wrapping algorithm that applies ideas similar to those in Chan's optimal output-sensitive algorithm for 3-dimensional convex hulls [14]. The algorithm has worst-case time complexity $O(F \log n)$, hence outperforming all algorithms discussed above, except possibly the worst-case optimal algorithm by Chazelle. Consider each point set P_i separately, and compute the polytope $\mathcal{P}_i$, as well as its Dobkin-Kirkpatrick hierarchy [29]. Then, perform the standard gift-wrapping algorithm on the 4-dimensional set P as follows. First compute an initial facet f_0 of $CH_4(P)$. Until all facets of $CH_4(P)$ have been computed, perform, as usual, the gift-wrapping steps of the algorithm: at each gift-wrapping step, consider a facet f of $CH_4(P)$ that has already been computed. Let t_j, $j = 1, 2, 3, 4$, be the four triangles of f, and for each triangle t_j determine a point $q \in P$ such that $f' = CH_3(t_j \cup \{q\})$ has the maximum possible angle with f. The maximum-angle query is done by considering each polytope $\mathcal{P}_i$ separately: for each $\mathcal{P}_i$ we determine the point q_i such that $CH_3(t_j \cup \{q_i\})$ has the maximum possible angle with f. Then, among all q_i's, $1 \leq i \leq m$, choose q to be the point that produces a tetrahedron that maximizes the angle with f. Unless f' has already been detected, add f' to the list of computed facets. Computing $\mathcal{P}_i$ takes $O(n_i \log n_i)$ time, which gives a total of $O(n \log n)$ for computing all m polytopes. The Dobkin-Kirkpatrick hierarchy can be computed in linear time in the size of the polytope, i.e., all such hierarchies can be computed in $O(n)$ total time. At each gift-wrapping step we consider four triangles, while for each triangle we consider each polytope $\mathcal{P}_i$, $1 \leq i \leq m$. For each such polytope we perform a maximum-angle query, which can be done in $O(\log n_i)$ time using the polytope's Dobkin-Kirkpatrick hierarchy. Since m is constant, the cost of computing all q_i's is $\sum_{i=1}^{m} O(\log n_i) = O(\log n)$, while determining q among the q_i's takes $O(m) = O(1)$ time. As a result, each gift-wrapping step of the algorithm takes $O(\log n)$ time. To compute f_0 we need three gift-wrapping steps, i.e., the starting facet for the gift-wrapping algorithm can be computed in $O(\log n)$ time also. The number of gift-wrapping steps is proportional to the number of facets of $CH_4(P)$. Since this is in $O(F)$, we conclude that the time complexity of the gift-wrapping algorithm described above is $O(F \log n)$.

d-dimensional polytopes in $\mathbb{E}^{d+1}$, where $d \geq 5$ odd. In dimension $d \geq 5$ and d odd we have $F = \Omega(n^2)$. The complexity of the applicable worst-case algorithms, whether output-sensitive or not, is shown in rows 5–6 and 8–9 of Table 1. Notice that for $d \geq 5$ odd, or equivalently $\alpha \geq 3$, the running time of Seidel's [9] or Matoušek and Schwarzkopf's [11] algorithm is $O(F \log n)$. Regarding

Chan's output-sensitive algorithm [15], its time complexity becomes:

$$O(n \log F + (nF)^{1-1/(\alpha+1)} \log^{O(1)} n) = O(n \log F + (nF)^{\alpha/(\alpha+1)} \log^{O(1)} n)$$
$$= O(n \log n + (n_2 n^\alpha)^{\alpha/(\alpha+1)} \log^{O(1)} n)$$
$$= O(n_2^{\alpha/(\alpha+1)} n^{\alpha^2/(\alpha+1)} \log^{O(1)} n).$$

Since $\alpha \geq 3$ and $n \geq n_2$, we have $n_2^{\alpha/(\alpha+1)} n^{\alpha^2/(\alpha+1)} \geq n_2 n^{\alpha-1}$, which implies that Chan's algorithm does not yield a better asymptotic complexity that the algorithm by Seidel or that by Matoušek and Schwarzkopf. This remains true even if the improvement ideas of the previous paragraph are applied.

Expected complexity. Another option is to apply Clarkson and Shor's randomized incremental algorithm for computing d-dimensional convex hulls [28] (refer also to row 7 of Table 1). The algorithm in [28] runs in $O(n) \cdot \sum_{r=1}^{n} \frac{C_r}{r^2}$ expected time, where C_r denotes the expected combinatorial complexity of the convex hull of a random subset, of size r, of the input set of points.

Let N_i be a random variable indicating the number of points from P_i contained in a random subset of P of size r. Clearly $\mathrm{Ex}[N_i] = (r/n)n_i$. It is now tempting to apply Theorem 3 to those expectations to claim that

$$C_r = O\left(\sum_{1 \leq i \neq j \leq m} \mathrm{Ex}[N_i]\mathrm{Ex}[N_j]^{\alpha-1} \right) = O\left(\left(\frac{r}{n}\right)^\alpha n_i n_j^{\alpha-1} \right).$$

Although the resulting upper bound is correct, the argument is fallacious since actually

$$C_r = O\left(\mathrm{Ex}\left[\sum_{1 \leq i \neq j \leq m} N_i N_j^{\alpha-1} \right] \right) = O\left(\sum_{1 \leq i \neq j \leq m} \mathrm{Ex}\left[N_i \cdot N_j^{\alpha-1} \right] \right),$$

and in general $\mathrm{Ex}[N_i N_j^{\alpha-1}] \neq O\left(\mathrm{Ex}[N_i]\mathrm{Ex}[N_j]^{\alpha-1} \right)$. However, in the case at hand this asymptotic bound actually holds.

To see this note that determining $\mathrm{Ex}[N_i N_j^{\alpha-1}]$ is asymptotically the same as determining $E_{ij} = \mathrm{Ex}[\binom{N_i}{1}\binom{N_j}{\beta}]$, where $\beta = \alpha - 1$. Let $p(r_i, r_j)$ be the probability that a random r-subset of P contains exactly r_i points from P_i and r_j points from P_j and the remaining $r' = r - r_i - r_j$ points from the other P_k's. We have

$$E_{ij} = \mathrm{Ex}\left[\binom{N_i}{1}\binom{N_j}{\beta} \right] = \sum_{\substack{r_i+r_j+r'=r \\ r_i,r_j,r' \geq 0}} p(r_i, r_j) \binom{r_i}{1}\binom{r_j}{\beta}.$$

We have $p(r_i, r_j) = \binom{n_i}{r_i}\binom{n_j}{r_j}\binom{n'}{r'}/\binom{n}{r}$, where $n' = n - n_i - n_j$ and $r' = r - r_i - r_j$. Thus we have

$$E_{ij} = \sum_{\substack{r_i+r_j+r'=r \\ r_i,r_j,r' \geq 0}} \frac{\binom{n_i}{r_i}\binom{n_j}{r_j}\binom{n'}{r'}}{\binom{n}{r}} \binom{r_i}{1}\binom{r_j}{\beta}.$$

Applying the binomial identity $\binom{A}{B}\binom{B}{C} = \binom{A}{C}\binom{A-C}{B-C}$ three times, namely $\binom{n_i}{r_i}\binom{r_i}{1} = \binom{n_i}{1}\binom{n_i-1}{r_i-1}$, $\binom{n_j}{r_j}\binom{r_j}{\beta} = \binom{n_j}{\beta}\binom{n_j-\beta}{r_j-\beta}$, $\binom{n}{r}\binom{r}{\beta+1} = \binom{n}{\beta+1}\binom{n-\beta-1}{r-\beta-1}$, this sum turns into

$$\frac{\binom{r}{\beta+1}}{\binom{n}{\beta+1}} \binom{n_i}{1}\binom{n_j}{\beta} \sum_{\substack{r_i+r_j+r'=r \\ r_i,r_j,r' \geq 0}} \frac{\binom{n_i-1}{r_i-1}\binom{n_j-\beta}{r_j-\beta}\binom{n'}{r'}}{\binom{n-\beta-1}{r-\beta-1}}.$$

The last sum evaluates to 1 since it essentially counts all ways of choosing subsets of size $r - \beta - 1$ from a set of size $n - \beta - 1$ that is partitioned into sets of size $n_i - 1$, $n_j - \beta$, and n'. Thus we get

$$E_{ij} = \frac{\binom{r}{\beta+1}}{\binom{n}{\beta+1}} \binom{n_i}{1} \binom{n_j}{\beta} = O\left(\left(\frac{r}{n}\right)^\alpha n_i n_j^{\alpha-1}\right). \tag{3}$$

From Theorem 3 we now get that the expected complexity of the convex hull of a random subset of P of size r is

$$C_r = O\left(\operatorname{Ex}\left[\sum_{1 \le i \ne j \le m} N_i N_j^{\alpha-1}\right]\right) = O\left(\sum_{1 \le i \ne j \le m} \operatorname{Ex}[N_i N_j^{\alpha-1}]\right),$$

which by (3) is

$$O\left(\left(\frac{r}{n}\right)^\alpha \sum_{1 \le i \ne j \le m} n_i n_j^{\alpha-1}\right) = O\left(\left(\frac{r}{n}\right)^\alpha F\right).$$

The complexity of Clarkson and Shor's algorithm thus becomes:

$$O\left(n \cdot \sum_{r=1}^{n} \frac{1}{r^2} \left(\frac{r}{n}\right)^\alpha F\right) = O\left(n^{1-\alpha} F \cdot \sum_{r=1}^{n} r^{\alpha-2}\right) = O(n^{1-\alpha} F \cdot n^{(\alpha-2)+1}) = O(F).$$

Summarizing our analysis above of the various possible algorithms, we arrive at the following theorem.

Theorem 5. *Let $\mathscr{P} = \{\mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_m\}$ be a set of m disjoint d-polytopes in $\mathbb{E}^{d+1}$, where $d \ge 3$ and d odd. Let $n_i = f_0(\mathcal{P}_i)$, $1 \le i \le m$, and $n = \sum_{i=1}^{m} n_i$. We can compute the convex hull $CH_{d+1}(\mathscr{P})$ in $O(\min\{n^{\lfloor \frac{d+1}{2} \rfloor}, (\sum_{1 \le i \ne j \le m} n_i n_j^{\lfloor \frac{d}{2} \rfloor}) \log n\})$ worst-case time, and $O(\sum_{1 \le i \ne j \le m} n_i n_j^{\lfloor \frac{d}{2} \rfloor})$ expected time.*

4. Convex hulls of spheres with a constant number of distinct radii

In this section we derive tight upper and lower bounds on the worst-case complexity of the convex hull of a set of spheres in $\mathbb{E}^d$ having a constant number $m \ge 2$ of distinct radii.

4.1. Upper bounds

Let Σ be a set of n spheres $\sigma_k = (c_k, r_k)$, $1 \le k \le n$, in $\mathbb{E}^d$, and let $CH_d(\Sigma)$ be the convex hull of the spheres in Σ. We will assume that the spheres are in non-degenerate position in the sense no $d + 2$ of the vectors σ_k are affinely dependent unless they all agree in their last component (which specifies the radius). It will become clear later that this non-degeneracy condition implies that no hyperplane is tangent to more than d spheres from Σ. Algorithmically this non-degeneracy condition can be enforced by symbolically perturbing the centers of the spheres.

A *facet of circularity* ℓ of $CH_d(\Sigma)$, $0 \le \ell \le d - 1$, is a maximal connected portion of the boundary of $CH_d(\Sigma)$ consisting of points where the supporting hyperplanes are tangent to a given set of $(d - \ell)$ spheres of Σ. In the special case where all spheres have the same radius, $CH_d(\Sigma)$ is combinatorially equivalent to the convex hull $CH_d(C)$ of the set C of centers of the spheres in Σ, in the sense that each facet of circularity ℓ of $CH_d(\Sigma)$ corresponds to a unique $(d - \ell - 1)$-face of $CH_d(C)$, for $0 \le \ell \le d - 1$.

We consider here the case where the radii r_k can take m distinct values, i.e., $r_k \in \{\rho_1, \rho_2, \ldots, \rho_m\}$. Without loss of generality we may assume that $0 < \rho_1 < \rho_2 < \ldots < \rho_m$. We identify $\mathbb{E}^d$ with the hyperplane $H_0 = \{x_{d+1} = 0\}$ of $\mathbb{E}^{d+1}$ and we call the $(d+1)$-axis of $\mathbb{E}^{d+1}$ the *vertical axis*, while the expression *above* will refer to the $(d+1)$-coordinate. Let Π_i, $1 \le i \le m$, be the hyperplane $\{x_{d+1} = \rho_i\}$ in $\mathbb{E}^{d+1}$, and let P be the point set in $\mathbb{E}^{d+1}$ obtained by mapping each sphere σ_k to the point $p_k = (c_k, r_k) \in \mathbb{E}^{d+1}$. Let P_i denote the subset of P containing points that

belong to the hyperplane Π_i, and let n_i be the cardinality of P_i. We denote by $\mathcal{P}$ the convex hull of the points in P (i.e., $\mathcal{P} = CH_{d+1}(P)$). We further denote by $\mathcal{P}_i$ the convex hull of the points in P_i (i.e., $\mathcal{P}_i = CH_d(P_i)$); more precisely, we identify Π_i with $\mathbb{E}^d$, and then define $\mathcal{P}_i$ to be the convex hull of the points in P_i, seen as points in $\mathbb{E}^d$. We use $\mathscr{P}$ to denote the set of the $\mathcal{P}_i$'s. Note that by our non-degeneracy assumption all facets of the $(d+1)$-polytope $\mathcal{P}$ are simplices except possibly the "top" and "bottom" facets which correspond to $\mathcal{P}_1$ and $\mathcal{P}_m$.

Let λ_0 be the half lower hypercone in $\mathbb{E}^{d+1}$ with arbitrary apex, vertical axis, and angle between axis and directrices equal to $\frac{\pi}{4}$. By $\lambda(p)$ we denote the translated copy of λ_0 with apex at p; observe that the intersection of the hypercone $\lambda(p_k)$ with the hyperplane H_0 is identical to the sphere σ_k. Let Λ be the set of the lower half hypercones $\{\lambda(p_1), \lambda(p_2), \ldots, \lambda(p_n)\}$ in $\mathbb{E}^{d+1}$ associated with the spheres of Σ. The intersection of the convex hull $CH_{d+1}(\Lambda)$ with H_0 is equal to $CH_d(\Sigma)$.

Let us call a hyperplane H *tilted* iff its normal is at angle $\pi/4$ with the vertical axis. Note that H is tilted iff it is tangent to a translate λ of λ_0 along a generatrix of λ. Let O' be a point in the interior of $\mathcal{P}$. We then have the following:

Theorem 6 ([21, Theorem 1]). *Any hyperplane of $\mathbb{E}^d$ supporting $CH_d(\Sigma)$ is the intersection with H_0 of a unique hyperplane H of $\mathbb{E}^{d+1}$ satisfying the following three properties:*

1. *H supports $\mathcal{P}$,*
2. *H is tilted,*
3. *H is above O'.*

Conversely, let H be a hyperplane of $\mathbb{E}^{d+1}$ satisfying the above three properties. Its intersection with H_0 is a hyperplane of $\mathbb{E}^d$ supporting $CH_d(\Sigma)$.

Boissonnat *et al.* [21] then use polarity to obtain the dual polar of $\mathcal{P}$. Given a hyperplane $H \in \mathbb{E}^{d+1}$, we denote by $H^\star$ its dual point, and given a point $p \in \mathbb{E}^{d+1}$, we denote by $p^\star$ its dual hyperplane and by $p^{\star-}$ the halfspace bounded by $p^\star$ containing O'. Then, according to [21], the following proposition holds:

Proposition 7 ([21, Proposition]).

1. *The polytope $\mathcal{P}^\star = p_1^{\star-} \cap p_2^{\star-} \cap \cdots p_n^{\star-}$ of $\mathbb{E}^{d+1}$ is dual to the polytope $\mathcal{P}$, i.e., there is a bijection between the ℓ-faces of $\mathcal{P}$ and the $(d-\ell)$-faces of $\mathcal{P}^\star$ which reverses the relation of inclusion. Each hyperplane supporting $\mathcal{P}$ along an ℓ-face F has its polar point on the corresponding $(d-\ell)$-face $F^\star$ of $\mathcal{P}^\star$.*

2. *The polar set of the tilted hyperplanes is a hypercone K with apex at O', a vertical axis, and an angle between axis and directrices equal to $\pi/4$.*

3. *The polar set of the hyperplanes above O' is the half space $x'_{d+1} > 0$.*

A consequence of the above proposition is the following (again, following the arguments in [21]): the polar set of the hyperplanes that

1. support the convex hull of the points in $\mathcal{P}$,
2. are tilted, and
3. are above O'

is the set $S = \mathcal{P}^+ \cap K \cap \{x'_{d+1} > 0\}$, where $\mathcal{P}^+$ is the boundary of $\mathcal{P}^\star$. In other words, the points in S correspond one-to-one with the hyperplanes that support the set of spheres Σ. In particular, if F is a face of $\mathcal{P}^\star$ defined the duals of points $p_{i_1}, \ldots, p_{i_\ell}$ and $x \in F \cap K \cap \{x'_{d+1} > 0\}$, then x corresponds to a hyperplane that supports Σ in spheres $\sigma_{i_1}, \ldots, \sigma_{i_\ell}$, and connected components of such x's correspond to faces of the convex hull of Σ.

Note that $F \cap K$ can have many connected components (e.g. think of intersecting a polygon and a circular or parabolic curve). However, the intersection of a simplex and cone K can consist of only a constant number of components. Thus if $\mathcal{P}^+$ is triangulated into N simplices, then the number of connected components in S can be at most $O(N)$, and hence the number of faces of the convex hull $CH_d(\Sigma)$ is $O(N)$. In the following we show that N is sufficiently small if we use the bottom-vertex triangulation of $\mathcal{P}^\star$.

Lemma 8. *The bottom-vertex triangulation of $\mathcal{P}^\star$ contains at most $N = O(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor \frac{d}{2} \rfloor})$ simplices.*

Proof. For a polytope (or polytopal complex) Q let us denote with $\overline{f}(Q)$ the total number of all its faces (of all dimensions). For a simple polytope Q it is easy to show that the number of simplices in the bottom-vertex triangulation is $O(\overline{f}(Q))$. However this fails to be the case for non-simple polytopes, and note that $\mathcal{P}^\star$ need not be simple, since $\mathcal{P}$ need not be simplicial: the "top" and "bottom" facets corresponding to $\mathcal{P}_1$ and $\mathcal{P}_m$ need not be simplices.

For a polytope Q let $\mathring{Q}$ denote its *barycentric subdivision*, which is a triangulation of Q that is defined as follows: if Q has dimension 0, i.e. it is a point, then $\mathring{Q} = Q$. If $d > 0$ pick a point c in the relative interior of Q, and for each facet G of Q and each simplex Δ in its barycentric subdivision $\mathring{G}$, include the simplex spanned by Δ and c in $\mathring{Q}$. It is well known that for a d-polytope Q the d-simplices in $\mathring{Q}$ correspond one-to-one with increasing maximal chains in the face lattice of Q. Since the face lattice of Q and its dual $Q^\star$ are the same except for inclusion reversion it follows that $\mathring{Q}$ and $\mathring{Q^\star}$ have the same number of d-faces and actually $\overline{f}(\mathring{Q}) = \overline{f}(\mathring{Q^\star})$ holds.

Let $\widehat{Q}$ be a bottom-vertex triangulation of Q. From the definitions it is clear that we have $\overline{f}(\widehat{Q}) \leq \overline{f}(\mathring{Q})$. For our lemma we therefore get

$$\overline{f}(\widehat{\mathcal{P}^\star}) \leq \overline{f}(\mathring{\mathcal{P}^\star}) = \overline{f}(\mathring{\mathcal{P}}),$$

and it remains to bound $\overline{f}(\mathring{\mathcal{P}})$.

For this purpose note first – taking d as constant – that for a d-simplex S we have $\overline{f}(\mathring{S}) = O(1)$. Next note that for any polytope Q we have $\overline{f}(\mathring{Q}) \leq 2 \cdot \sum_{G \text{ facet of } Q} \overline{f}(\mathring{G})$. This implies that for a simplicial polytope Q we have $\overline{f}(\mathring{Q}) = O(\overline{f}(Q))$, and this also implies that in our case at hand $\overline{f}(\mathring{\mathcal{P}}) = O(\overline{f}(\mathcal{P}))$, since at most 2 facets of $\mathcal{P}$ are not simplices, while their boundary complexes are simplicial by our non-degeneracy assumption. But by Theorem 3 we have $\overline{f}(\mathcal{P}) = O(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor \frac{d}{2} \rfloor})$. $\qquad\qquad\square$

Summarizing we can state:

Theorem 9. *Let a set Σ of spheres in $\mathbb{E}^d$, consisting of n_i spheres of radius ρ_i, $1 \leq i \leq m$, with $m \geq 2$ constant. The worst-case complexity of the convex hull $CH_d(\Sigma)$ is $O(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor \frac{d}{2} \rfloor})$.*

4.2. Balanced polytopes

In this subsection we describe a family of even-dimensional polytopes, called *balanced polytopes*, that play a crucial role in our lower bound construction for the sphere convex hull problem (see next subsection). A balanced polytope $\mathcal{P}$ in $\mathbb{E}^d$, $d = 2\delta$, with n vertices, has the following property: there exists a subset $\mathcal{B}$ of the facets of $\mathcal{P}$, such that:

1. the facets in $\mathcal{B}$ are simplicial,
2. the cardinality of $\mathcal{B}$ is $\Theta(n^\delta) = \Theta(n^{\lfloor \frac{d}{2} \rfloor})$, and
3. there exists a $(d-1)$-sphere Σ, such that for every facet F in $\mathcal{B}$, Σ intersects the interior of F, but none of the ridges of $\mathcal{P}$ that belong to the boundary of F.

We will call a facet in $\mathcal{B}$ a *balanced facet* of $\mathcal{P}$, while $\mathcal{B}$ will be called, naturally, the *set of balanced facets* of $\mathcal{P}$. As we will see in the next subsection, our lower bound construction is based the existence of the set of balanced facets, and we will exploit their properties.

For any even dimension $d = 2\delta$, consider δ unit circles $C_1, \ldots, C_\delta$, with their centers at the origin and C_j lying in the plane spanned by the $x_{2j-1} - x_{2j}$ axes. We are going to place points on each C_j as follows. Let $\nu = \lfloor \frac{n}{\delta} \rfloor$, and $\nu' = n - \nu\delta = n \pmod{\delta}$. For each j, $1 \leq j \leq \delta - 1$, we place ν points on C_j so that they form a regular ν-gon. On C_δ, we place $\nu + \nu'$ points, where

the first ν points are placed so that they create a regular ν-gon, while the remaining ν' points are placed arbitrarily on C_δ between the ν-th and the first point of C_δ. More precisely, for each C_j, $1 \leq j \leq \delta$, the k-th vertex, $0 \leq k \leq \nu - 1$, is $(\cos(t_k), \sin(t_k))$, $t_k = \frac{2k\pi}{\nu}$, embedded in the x_{2j-1} and x_{2j} axes.

Let $\mathcal{P}$ be the convex hull of all the vertices of all circles C_j, $1 \leq j \leq \delta$, and notice that the vertices of $\mathcal{P}$ lie on the unit sphere $\mathbb{S}^{d-1}$ centered at the origin of $\mathbb{E}^d$. Call $\mathcal{B}$ the set of vertex subsets of $\mathcal{P}$ created by taking two points from each C_j, $1 \leq j \leq \delta$, where the indices of these two points are consecutive and at most ν (in other words we consider the first $\nu - 1$ edges per C_j). It is easy to verify that each vertex subset in $\mathcal{B}$ defines a simplicial facet for $\mathcal{P}$. Hence, we can identify the vertex subsets in $\mathcal{B}$ with the associated facets of $\mathcal{P}$. Moreover, the number of the vertex subsets (or facets) in $\mathcal{B}$ is $(\nu - 1)^\delta$, which means that $\mathcal{B}$ contains $\Theta(n^\delta)$ facets. We will show below that the facets in $\mathcal{B}$ are balanced facets.

Let F be a facet of $\mathcal{P}$ in $\mathcal{B}$, and recall that each pair of points in F coming from the same circle C_j have parameter values t_k and t_{k+1}, for some k, where $t_k = \frac{2k\pi}{\nu}$. Call θ the difference between t_{k+1} and t_k, i.e., $\theta = t_{k+1} - t_k = \frac{2\pi}{\nu}$. We may assume, without loss of generality, that the j-th pair of points of F come from the circle C_j and that the corresponding parameter values are $t_{j,1}$ and $t_{j,2}$, where $t_{j,2} - t_{j,1} = \theta$. Call b the barycenter of F, i.e.,

$$b = \frac{1}{d} \left[\sum_{j=1}^{\delta} (\cos t_{j,1} + \cos t_{j,2}) \boldsymbol{e}_{2j-1} + \sum_{j=1}^{\delta} (\sin t_{j,1} + \sin t_{j,2}) \boldsymbol{e}_{2j} \right].$$

It is now fairly easy to verify that for any vertex v of F, we have $\|b - v\|_2^2 = 1 - \frac{1}{d}(1 + \cos\theta)$. Moreover, $\|b\|_2^2 = \frac{1}{d}(1 + \cos\theta)$. Hence, b is equidistant from each vertex of F, which implies that b is the circumcenter of the unique, since F is a $(d-1)$-simplex, circumscribing $(d-1)$-sphere of the vertex set of F. Moreover, b is forcibly the point of F closest to the origin. To see this, first note that b is by construction (as the barycenter) an interior point of F (the important point here is that b is a point in the closure F and not in the complement of the closure of F with respect to its affine hull). Second, observe that b is also the point of the supporting hyperplane H_F of F closest to the origin: recall that the points in F lie on the unit sphere $\mathbb{S}^{d-1}$ in $\mathbb{E}^d$, and hence also on the intersection S of H_F with $\mathbb{S}^{d-1}$; the center of S, which is b, is by construction the point closest to the origin.

Summarizing the analysis above, we deduce that the distance of F from the origin is $(\frac{1}{d}(1 + \cos\theta))^{1/2}$, and this distance is realized with a point in the interior of F. Furthermore, notice that the distance of F from the origin is, in fact, independent from the choice of F in $\mathcal{B}$. In other words, the $(d-1)$-sphere Σ centered at the origin with radius $(\frac{1}{d}(1 + \cos\theta))^{1/2}$ touches every facet $F \in \mathcal{B}$ at an interior point of F and lies in the same halfspace, with respect to the supporting hyperplane H_F of F, as $\mathcal{P}$. Consider, now, a sphere Σ', centered also at the origin, with radius $(\frac{1}{d}(1 + \cos\theta))^{1/2} + \varepsilon'$, where $\varepsilon' > 0$. If we choose ε' small enough, Σ' intersects the interior of every facet F in $\mathcal{B}$, but none of the ridges on the boundary of F. In other words, every F in $\mathcal{B}$ is a balanced facet of $\mathcal{P}$, and $\mathcal{B}$ is the set of balanced facets of $\mathcal{P}$ satisfying the three properties mentioned at the beginning of this subsection.

4.3. Lower bound construction with two distinct radii

We will now exploit the construction of balanced polytopes of the previous subsection, in order to construct a set Σ of $\Theta(n_1 + n_2)$ spheres in $\mathbb{E}^d$, with $d \geq 3$ and d odd, where n_i spheres have radius ρ_i, $i = 1, 2$, and such that the complexity of $CH_d(\Sigma)$ is $\Omega(n_1 n_2^{\lfloor \frac{d}{2} \rfloor} + n_2 n_1^{\lfloor \frac{d}{2} \rfloor})$.

In what follows we assume that the ambient space is $\mathbb{E}^d$, where $d \geq 3$ is odd, and let $\delta = \lfloor \frac{d-1}{2} \rfloor = \lfloor \frac{d}{2} \rfloor$. Let H_1 and H_2 be the hyperplanes $\{x_d = z_1\}$ and $\{x_d = z_2\}$, where $z_1, z_2 \in \mathbb{R}$, and $z_2 > z_1 + 2(n_2 + 1)$; the quantity n_2 will be defined below. Consider a set Σ_i, $i = 1, 2$, of n_1 points, treated as spheres of $\mathbb{E}^d$ of zero radius, on the $(d-2)$-dimensional unit sphere $\mathbb{S}^{d-2}$ embedded in H_i and centered at the origin of H_i (please refer to Fig. 1(left), as well as Fig. 2 for the view of the construction from the positive x_d-axis). In other words, the points of Σ_i lie on the $(d-2)$-dimensional unit sphere of $\mathbb{E}^d$, centered at $(0, 0, \ldots, 0, z_i)$. The n_1 points in Σ_i are chosen as in

Subsection 4.2, and call $\mathcal{Q}_i$ their convex hull. By construction, $\mathcal{Q}_i$ is a balanced $(d-1)$-polytope, and call $\mathcal{B}_i$ the set of balanced facets of $\mathcal{Q}_i$. Recall that $\mathcal{B}_i$ has cardinality $(\lfloor \frac{n_1}{\delta} \rfloor - 1)^\delta$, and that each vertex subset in $\mathcal{Q}_i$ corresponds to a simplicial facet of $\mathcal{Q}_i$, thus yielding $\Theta(n_1^\delta) = \Theta(n_1^{\lfloor \frac{d}{2} \rfloor})$ facets for $\mathcal{Q}_i$. Finally, observe that $\mathcal{Q}_2$ is a translated copy of $\mathcal{Q}_1$ along the x_d-axis and vice versa.

The convex hull of the $2n_1$ points of $\Sigma_1 \cup \Sigma_2$ is a prism Δ. The prism Δ consists of $\Theta(n_1^{\lfloor \frac{d}{2} \rfloor})$ facets not lying on H_1 or H_2, called the *vertical facets*, the $(d-1)$-face of $CH_{d-1}(\Sigma_1)$, called the *bottom facet*, and the $(d-1)$-face of $CH_{d-1}(\Sigma_2)$, called the *top facet*. For each vertical facet F of Δ, we denote by $\vec{v}_F$ the unit normal vector of F pointing outside Δ, and by F^+ (resp., F^-) the positive (resp., negative) open halfspace delimited by the supporting hyperplane of F. Regarding the ridges of Δ, those that are intersections of vertical facets of Δ will be referred to as *vertical ridges*. Notice that the vertical ridges of Δ are perpendicular to H_1 and H_2. For each ridge of Δ in $\mathcal{B}_1$, there is a unique corresponding ridge in $\mathcal{B}_2$ (they are translated copies of each other), and together they form a vertical facet for Δ. We are going to denote by $\mathcal{B}_\Delta$ the set of vertical facets of Δ with ridges in $\mathcal{B}_1$ and $\mathcal{B}_2$, and we are going to call the vertical facets of Δ in $\mathcal{B}_\Delta$ the *balanced vertical facets* of Δ. Since $\mathcal{B}_1$ and $\mathcal{B}_2$ have cardinality $\Theta(n_1^{\lfloor \frac{d}{2} \rfloor})$, the same bound holds for the cardinality of $\mathcal{B}_\Delta$.

Let Y be the oriented hyperplane $\{x_1 = 0\}$ with unit normal vector $\vec{v} = (1, 0, \ldots, 0)$. Let also $\overline{Y^+}$ be the closed positive halfspace of $\mathbb{E}^d$ delimited by Y. Y contains the x_d-axis, and is perpendicular to the hyperplanes H_1 and H_2. Recall that the points of Σ_i have been chosen to lie on unit circles C_j, lying on the plane spanned by the $x_{2j-1} - x_{2j}$ axes, $1 \le j \le \delta$. Due to the way that the point set Σ_i has been constructed, at least $\lfloor \frac{1}{2} \lfloor \frac{n_1}{\delta} \rfloor \rfloor$ of the points in C_1 are contained in $\overline{Y^+}$. This further implies that at least $(\lfloor \frac{1}{2} \lfloor \frac{n_1}{\delta} \rfloor \rfloor - 1) \cdot (\lfloor \frac{n_1}{\delta} \rfloor - 1)^{\delta-1}$ balanced facets of $\mathcal{Q}_i$ are contained in $\overline{Y^+}$. We thus conclude that the number of balanced facets of $\mathcal{Q}_i$ in $\overline{Y^+}$ is $\Theta(n_1^\delta) = \Theta(n_1^{\lfloor \frac{d}{2} \rfloor})$; the same bound clearly holds for the number of balanced vertical facets of Δ in $\overline{Y^-}$.

Define now a set $\Sigma_3 = \{\sigma_0, \sigma_1, \ldots, \sigma_{n_2+1}\}$ of $n_2 + 2$ spheres in $\mathbb{E}^d$, where $\sigma_k = (c_k, \rho)$, and $c_k = (0, \ldots, 0, 2k+1)$, $0 \le k \le n_2 + 1$. In other words, the sphere σ_k is centered on the x_d-axis, with the d-th coordinate of its center c_k being $(2k+1)$, while its radius is ρ. The radius ρ of σ_k is chosen so that its projection on H_1 or H_2 defines a $(d-2)$-ball that intersects the balanced facets of $\mathcal{Q}_1$ or $\mathcal{Q}_2$, respectively, but none of the ridges incident to these balanced facets. Following the analysis in Subsection 4.2, such a choice for ρ is indeed possible: set $\rho = (\frac{1}{d-1}(1 + \cos\theta))^{1/2} + \varepsilon'$, where $\theta = \frac{2\pi}{n_1}$, and $\varepsilon' > 0$ is chosen small enough. As a result of this choice for ρ, each sphere σ_k satisfies the following two properties:

(1) it does not intersect any of the ridges incident to the balanced vertical facets of Δ, and
(2) it intersects the interior of all balanced vertical facets of Δ.

Notice also that none of the spheres in Σ_3 intersects the hyperplanes H_1 and H_2 (recall that $z_2 > z_1 + 2(n_2+2)$), while the spheres in Σ_3 are pairwise disjoint; these two observations, however, are not critical for our construction.

We are now going to perturb the centers of the spheres in Σ_3 to get a new set of spheres Σ_3' (see Fig. 1(right), as well as Fig. 3 for the view of the construction from the positive x_d-axis). Define σ_k' to be the sphere with radius ρ and center $c_k' = c_k + (\sum_{\ell=0}^{k} \frac{\varepsilon}{2^\ell})\vec{v} = c_k + \varepsilon(2 - \frac{1}{2^k})\vec{v}$, where $0 < \varepsilon \ll 1$. The quantity ε is chosen so that the spheres in Σ_3' satisfy almost the same conditions as the spheres in Σ_3. In particular, we require that condition (1) is still satisfied, while we relax the requirement on condition (2): we now require that σ_k' intersects the interior of all balanced vertical facets of Δ that are contained in $\overline{Y^+}$. In addition to the two conditions above, we also require that for each k, $0 \le k \le n_2 + 1$, the $(d-2)$-dimensional sphere $\sigma_k \cap \sigma_k'$ is contained in F^- for all balanced vertical facets F of Δ that are contained in $\overline{Y^+}$.

We will now show that for each pair (σ_k', F), where $1 \le k \le n_2$ and F is a balanced vertical facet of Δ in $\overline{Y^+}$, the spherical cap $F^+ \cap \sigma_k'$ induces a facet of circularity $(d-1)$ in $CH_d(\Sigma)$. Let F_1 and F_2 be the ridges of Δ on the boundary of F contained in the top and bottom facet, respectively. Finally, let S_k be the supporting hyperplane of σ_k parallel to F; we consider S_k

14

Figure 1: The lower bound construction in the case of two radii. The points in Σ_1 (resp., Σ_2) are shown in black (resp., white). The hyperplane Y is shown in green, while the prism Δ is shown in black. The facet F in blue is one of the vertical facets of Δ in Y^+. The sphere sets Σ_3 (left) and Σ_3' (right) are shown in red. The red spherical caps on the left correspond to a unique supporting hyperplane of $CH_d(\Sigma_1 \cup \Sigma_2 \cup \Sigma_3)$. The red spherical caps on the right correspond to facets of $CH_d(\Sigma_1 \cup \Sigma_2 \cup \Sigma_3')$ of circularity $(d-1)$.

Figure 2: View from the positive x_d-axis of the construction in Fig. 1(left). The silhouettes of all spheres in Σ_3 coincide.

Figure 3: View from the positive x_d-axis of the construction in Fig. 1(right).

to be oriented as F (i.e., the unit normal vector of S_k is $\vec{\nu}_F$), and thus σ_k lies in the closure of the negative halfspace delimited by S_k. Notice that S_k is also a supporting hyperplane for Σ_3. Let S'_k be the hyperplane we get by translating S_k by the vector $\varepsilon(2 - \frac{1}{2^k})\vec{\nu}$. S'_k supports σ'_k, but fails to be a supporting hyperplane for Σ'_3. More precisely, S'_k intersects all spheres σ'_j with $j > k$, whereas all spheres σ'_j with $j < k$, are contained in the negative open halfspace delimited by S'_k. We can, however, perturb S'_k so that it supports Σ'_3: simply slide S'_k on sphere σ'_k towards F_1, while maintaining the property that it remains parallel to F_1 and F_2. We keep sliding S'_k until it has empty intersection with any sphere σ'_j with $j > k$. Notice that due to the way we have perturbed the centers of the spheres in Σ_3 to get Σ'_3, the new hyperplane S''_k we get via this transformation is a supporting hyperplane for Σ'_3. In fact, S''_k is a supporting hyperplane for the sphere set $\Sigma = \Sigma_1 \cup \Sigma_2 \cup \Sigma'_3$ (it touches $CH_d(\Sigma)$ at σ'_k only), which implies that S''_k corresponds to a unique facet of circularity $(d-1)$ on $CH_d(\Sigma)$.

The same construction can be done for all k with $1 \leq k \leq n_2$, and for all balanced vertical facets of Δ in $\overline{Y^+}$. Since we have $\Theta(n_1^{\lfloor \frac{d}{2} \rfloor})$ balanced vertical facets of Δ in $\overline{Y^+}$, we can construct $n_2 \cdot \Theta(n_1^{\lfloor \frac{d}{2} \rfloor})$ distinct supporting hyperplanes of $CH_d(\Sigma)$, corresponding to distinct facets of circularity $(d-1)$ on $CH_d(\Sigma)$. Hence the complexity of $CH_d(\Sigma)$ is $\Omega(n_2 n_1^{\lfloor \frac{d}{2} \rfloor})$. Without loss of generality, we may assume that $n_2 \leq n_1$, in which case we have $n_2 n_1^{\lfloor \frac{d}{2} \rfloor} \geq \frac{1}{2}(n_2 n_1^{\lfloor \frac{d}{2} \rfloor} + n_1 n_2^{\lfloor \frac{d}{2} \rfloor})$. Hence, we arrive at the following:

Theorem 10. *Fix some odd $d \geq 3$. There exists a set Σ of spheres in $\mathbb{E}^d$, consisting of n_i spheres of radius ρ_i, $i = 1, 2$, with $\rho_1 \neq \rho_2$, such that the complexity of the convex hull $CH_d(\Sigma)$ is $\Omega(n_1 n_2^{\lfloor \frac{d}{2} \rfloor} + n_2 n_1^{\lfloor \frac{d}{2} \rfloor})$.*

4.4. Lower bound construction with m distinct radii

We can easily generalize the lower bound construction of the previous subsection in the case where we have n_i spheres of radius ρ_i, $1 \leq i \leq m$, $m \geq 3$, and the radii ρ_i are considered to be mutually distinct.

As in the previous subsection, the ambient space is $\mathbb{E}^d$, where $d \geq 3$ is odd. Let $N_1 = \sum_{i=2}^m n_i$ and $N_2 = n_1$. We construct the set $\Sigma = \Sigma_1 \cup \Sigma_2 \cup \Sigma'_3$ as in the previous subsection where Σ_1 and Σ_2 contain each N_1 points and Σ'_3 contains $N_2 + 2$ spheres of some appropriate radius ρ (recall that in the construction of the previous subsection $\rho \approx (\frac{1}{d-1}(1 + \cos\theta))^{1/2} \geq \frac{1}{\sqrt{d-1}}$). We then replace n_i among the N_1 points of Σ_1 (resp., Σ_2) by spheres with the same center and radius equal to r^i, where $0 < r \ll \frac{1}{\sqrt{d-1}}$ and $2 \leq i \leq m$. We choose r small enough so that the following two conditions hold:

(1) the prism $\Delta_r = CH_d(\Sigma_1 \cup \Sigma_2)$ is combinatorially equivalent[2] to the prism Δ_0 (this is the prism we get for $r = 0$, which is the prism Δ of the previous subsection), and

(2) the two requirements for the spheres in Σ'_3 are still satisfied: each σ'_k does not intersect any of the ridges[3] of the balanced vertical facets of Δ_r, while each σ'_k intersects the interior of all balanced vertical facets of Δ_r in $\overline{Y^+}$.

As described in the previous subsection, the convex hull $CH_d(\Sigma)$ of the set $\Sigma = \Sigma_1 \cup \Sigma_2 \cup \Sigma'_3$ of $2N_1 + N_2 + 2$ spheres has $N_2 \cdot \Theta(N_1^{\lfloor \frac{d}{2} \rfloor})$ facets of circularity $(d-1)$, and hence its complexity is $\Omega(N_2 N_1^{\lfloor \frac{d}{2} \rfloor}) = \Omega(n_1 (\sum_{i=2}^m n_i)^{\lfloor \frac{d}{2} \rfloor})$. Without loss of generality we may assume that $n_2 \geq n_1 \geq n_i$ for all $3 \leq i \leq m$, in which case we have: $n_1 (\sum_{i=2}^m n_i)^{\lfloor \frac{d}{2} \rfloor} \geq n_1 n_2^{\lfloor \frac{d}{2} \rfloor} \geq \frac{1}{m(m-1)}(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor \frac{d}{2} \rfloor})$. Since m is constant, we conclude that the complexity of $CH_d(\Sigma)$ is $\Omega(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor \frac{d}{2} \rfloor})$.

---

[2]Combinatorial equivalence here means that each facet of circularity ℓ of Δ_r corresponds to a unique $(d-\ell-1)$-face of Δ_0.

[3]Since Δ_r is no longer a polytope, Δ_r does not have any ridges (resp., vertical facets), but rather facets of circularity $d-2$ (resp., $d-1$) that are associated with ridges (resp., vertical facets) of Δ_0. On the other hand, due to the combinatorial equivalence of Δ_r and Δ_0 the term "ridges" (resp., "vertical facets") can be safely used here.

Theorem 11. *Fix some odd $d \geq 3$. There exists a set Σ of spheres in $\mathbb{E}^d$, consisting of n_i spheres of radius ρ_i, $1 \leq i \leq m$, with $m \geq 3$ constant, such that the complexity of the convex hull $CH_d(\Sigma)$ is $\Omega(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor \frac{d}{2} \rfloor})$.*

This theorem immediately also implies a lower bound on the worst case complexity of the convex hull of disjoint d-polytopes in $\mathbb{E}^{d+1}$. We have shown above that the total number Z of faces of $CH_d(\Sigma)$ is $\Omega(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor \frac{d}{2} \rfloor})$, where $d \geq 3$ odd and $m \geq 2$. When proving Theorem 9 we showed that $Z = O(X)$, where X is the number of faces of the convex hull of m disjoint d-polytopes in $\mathbb{E}^{d+1}$. But $Z = O(X)$ is equivalent to $X = \Omega(Z)$. Thus the construction above, which yields a large number of faces for the convex hull of spheres, also yields a large number of faces for the corresponding convex hull of disjoint d-polytopes. This establishes our lower bound claim in Theorem 3:

Corollary 12. *Let $\mathscr{P} = \{\mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_m\}$ be a set of $m \geq 2$ disjoint d-polytopes in $\mathbb{E}^{d+1}$, with $d \geq 3$, d odd, where both d and m are constant. The worst-case complexity of $CH_{d+1}(\mathscr{P})$ is $\Omega(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor \frac{d}{2} \rfloor})$, where $n_i = f_0(\mathcal{P}_i)$, $1 \leq i \leq m$.*

5. Computing convex hulls of spheres

In this section we focus our attention on the computation of the convex hull $CH_d(\Sigma)$ of Σ. We use the same notation as in Section 4.1. Given a set Σ of n spheres in $\mathbb{E}^d$, we saw in Section 4.1 that the faces of $CH_d(\Sigma)$ can be gleaned from the intersection of the boundary of a $(d+1)$-polytope with a spherical cone. Using the notation of that section, we need to compute $\mathcal{P}^\star \cap K \cap \{x'_{d+1} > 0\}$. Boissonnat *et al.* [21] have used this property in order to propose an algorithm for computing $CH_d(\Sigma)$ in $O(n^{\lceil \frac{d}{2} \rceil} + n \log n)$ time for any $d \geq 2$. Below, we describe a slightly modified algorithm that takes into account the fact that the radii of the spheres in Σ can take on a constant number of $m \geq 2$ distinct values and that also explicates how to intersect a face of $\mathcal{P}^\star$ with K, which is a non-trivial operation since such an intersection may consist of many connected components.

Our algorithm consists of the following six steps, where we use the notation from Section 4.1:

1. For all i with $1 \leq i \leq m$: determine the set $P_i = P \cap \Pi_i$ and construct the convex hull $\mathcal{P}_i = CH_d(P_i)$.
2. Compute the polytope $\mathcal{P} = CH_{d+1}(P)$, and choose a point O' inside $\mathcal{P}$.
3. Compute the polar polytope $\mathcal{P}^\star$ of $\mathcal{P}$ with respect to O'.
4. Compute a bottom-vertex triangulation Δ of $\mathcal{P}^\star$.
5. For each simplex D in Δ compute the intersection $D \cap K \cap \{x'_{d+1} > 0\}$.
6. From all these intersections recover the incidence graph of the facets in $CH_d(\Sigma)$.

Determining all the sets P_i takes $\Theta(n)$ time, whereas constructing the polytope $\mathcal{P}_i$ takes $O(n_i^{\lfloor \frac{d}{2} \rfloor} + n_i \log n_i)$ time. We thus conclude that step 1 of the algorithm takes $O(n^{\lfloor \frac{d}{2} \rfloor} + n \log n)$ time. Let X be the number of faces of $\mathcal{P}$. Step 2 computes $\mathcal{P}$ and takes time at least $\Omega(X)$. Finding the point O' and computing $\mathcal{P}^\star$ from $\mathcal{P}$ can be done in $O(X)$ time. The bottom-vertex triangulation Δ can be computed in time $O(X)$ also, moreover the number of its simplices is $O(X)$. In Step 5 constant time needs to be afforded for each simplex, leading to $O(X)$ time overall for this step. Finally, Step 6 can be completed in time $O(X)$ also. Thus the time taken for Steps 1 and 2 dominate the running time of the entire algorithm and we get the following:

Theorem 13. *Let Σ be a set of n spheres in $\mathbb{E}^d$, having a constant number of m distinct radii $\rho_1, \rho_2, \ldots, \rho_m$, with $d \geq 3$, d odd. Let n_i be the number of spheres in Σ with radius ρ_i, $1 \leq i \leq m$. We can compute the convex hull $CH_d(\Sigma)$ in $O(n^{\lfloor \frac{d}{2} \rfloor} + n \log n + T_{d+1}(n_1, n_2 \ldots, n_m))$ time, where $T_{d+1}(n_1, n_2 \ldots, n_m)$ stands for the time to compute the convex hull of m disjoint d-polytopes $\mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_m$ in $\mathbb{E}^{d+1}$, with $n_i = f_0(\mathcal{P}_i)$, $1 \leq i \leq m$.*

As described in Section 3, for any $d \geq 3$ and d odd, $CH_{d+1}(\{\mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_m\})$ can be computed in $O(\min\{n^{\lfloor\frac{d+1}{2}\rfloor}, F \log n\})$ worst-case time and $O(F)$ expected time, where $F = \sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor\frac{d}{2}\rfloor}$. Hence, for any odd $d \geq 3$, we can compute the convex hull $CH_d(\Sigma)$ in $O(\min\{n^{\lfloor\frac{d+1}{2}\rfloor}, (\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor\frac{d}{2}\rfloor}) \log n\})$ worst-case time, and in $O(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor\frac{d}{2}\rfloor})$ expected time.

6. Summary and open problems

In this paper we have considered the problem of computing the worst-case complexity of the convex hull $CH_{d+1}(\mathscr{P})$ of a set $\mathscr{P} = \{\mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_m\}$ of m disjoint convex d-polytopes in $\mathbb{E}^{d+1}$, for any odd $d \geq 3$. Denoting by n_i the number of vertices of $\mathcal{P}_i$, we have shown that the worst-case complexity of $CH_{d+1}(\mathscr{P})$ is $O(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor\frac{d}{2}\rfloor})$. This result suggests that, in order to compute $CH_{d+1}(\mathscr{P})$, it might pay off to apply an output-sensitive convex hull algorithm to the set of vertices in $\mathscr{P}$. Indeed, we show that for any odd $d \geq 3$, we can compute $CH_{d+1}(\mathscr{P})$ in $O(\min\{n^{\lfloor\frac{d+1}{2}\rfloor}, (\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor\frac{d}{2}\rfloor}) \log n\})$ worst-case time. The above algorithms are nearly optimal for any odd $d \geq 3$; it remains an open problem to compute $CH_{d+1}(\mathscr{P})$ in worst-case optimal $O(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor\frac{d}{2}\rfloor} + n \log n)$ time. This optimal complexity can be achieved, however, by applying Clarkson and Shor's randomized incremental construction paradigm for convex hulls. Following this paradigm, we show that $CH_{d+1}(\mathscr{P})$ can be computed in $O(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor\frac{d}{2}\rfloor})$ expected time.

Capitalizing on our result on the complexity of convex hulls of disjoint convex polytopes, we have shown that the worst-case complexity of the convex hull $CH_d(\Sigma)$ of a set Σ of n spheres in $\mathbb{E}^d$, with a constant number m of distinct radii $\rho_1, \ldots, \rho_m$, is $O(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor\frac{d}{2}\rfloor})$, for any odd $d \geq 3$, where n_i is the number of spheres with radius ρ_i. By means of an appropriate construction, described in Subsections 4.2–4.4, we have shown that the upper bound above is asymptotically tight, implying that our upper bound for $CH_{d+1}(\mathscr{P})$ is also tight. By slightly, but crucially, modifying the algorithm of Boissonnat *et al.* [21], $CH_d(\Sigma)$ may be computed in $O(n^{\lfloor\frac{d}{2}\rfloor} + n \log n + T_{d+1}(n_1, \ldots, n_m))$ time, where $T_{d+1}(n_1, \ldots, n_m)$ stands for the time needed to compute the convex hull of m disjoint d-polytopes in $\mathbb{E}^{d+1}$, where the i-th polytope has n_i vertices (cf. Section 5). Applying our bounds for $T_{d+1}(n_1, \ldots, n_m)$ mentioned above, we can compute $CH_d(\Sigma)$ in $O(\min\{n^{\lfloor\frac{d+1}{2}\rfloor}, (\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor\frac{d}{2}\rfloor}) \log n\})$ worst-case time, and in $O(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor\frac{d}{2}\rfloor})$ expected time. As in the disjoint polytopes' case, it remains an open problem to compute $CH_d(\Sigma)$ in optimal $O(\sum_{1 \leq i \neq j \leq m} n_i n_j^{\lfloor\frac{d}{2}\rfloor} + n \log n)$ worst-case time.

Finally, Boissonnat and Karavelas [22] have shown that convex hulls of spheres in $\mathbb{E}^d$ and additively weighted Voronoi cells in $\mathbb{E}^d$ are combinatorially equivalent. This equivalence suggests that we should be able to refine the worst-case complexity of an additively weighted Voronoi cell in any odd dimension, when the number of distinct radii of the spheres involved is considered constant.

References

[1] M. I. Karavelas, E. Tzanaki, Convex hulls of spheres and convex hulls of convex polytopes lying on parallel hyperplanes, in: Proceedings of the 27th Annual ACM Symposium on Computational Geometry (SCG'11), Paris, France, 2011, pp. 397–406.

[2] L. R. Graham, An efficient algorithm for determining the convex hull of a finite planar set, Inform. Process. Lett. 1 (4) (1972) 132–133.

[3] F. P. Preparata, S. J. Hong, Convex hulls of finite sets of points in two and three dimensions, Commun. ACM 20 (1977) 87–93.

[4] A. M. Andrew, Another efficient algorithm for convex hulls in two dimensions, Inform. Process. Lett. 9 (5) (1979) 216–219.

[5] F. Preparata, An optimal real-time algorithm for planar convex hulls, Commun. ACM 22 (1979) 402–405.

[6] B. Chazelle, An optimal convex hull algorithm in any fixed dimension, Discrete Comput. Geom. 10 (1) (1993) 377–409.

[7] D. R. Chand, S. S. Kapur, An algorithm for convex polytopes, J. ACM 17 (1) (1970) 78–86.

[8] R. A. Jarvis, On the identification of the convex hull of a finite set of points in the plane, Inform. Process. Lett. 2 (1973) 18–21.

[9] R. Seidel, Constructing higher-dimensional convex hulls at logarithmic cost per face, in: Proc. 18th Annu. ACM Sympos. Theory Comput., 1986, pp. 404–413.

[10] D. G. Kirkpatrick, R. Seidel, The ultimate planar convex hull algorithm?, SIAM J. Comput. 15 (1986) 287–299.

[11] J. Matoušek, O. Schwarzkopf, Linear optimization queries, in: Proc. 8th Annu. ACM Sympos. Comput. Geom., 1992, pp. 16–25.

[12] B. Chazelle, J. Matoušek, Derandomizing an output-sensitive convex hull algorithm in three dimensions, Comput. Geom. Theory Appl. 5 (1995) 27–32.

[13] N. M. Amato, E. A. Ramos, On computing Voronoi diagrams by divide-prune-and-conquer, in: Proc. 12th Annu. ACM Sympos. Comput. Geom., 1996, pp. 166–175.

[14] T. M. Chan, Optimal output-sensitive convex hull algorithms in two and three dimensions, Discrete Comput. Geom. 16 (4) (1996) 361–368.

[15] T. M. Chan, Output-sensitive results on convex hulls, extreme points, and related problems, Discrete Comput. Geom. 16 (4) (1996) 369–387.

[16] T. M. Chan, J. Snoeyink, C.-K. Yap, Primal dividing and dual pruning: Output-sensitive construction of four-dimensional polytopes and three-dimensional Voronoi diagrams, Discrete Comput. Geom. 18 (4) (1997) 433–454.

[17] J. Erickson, New lower bounds for convex hull problems in odd dimensions, SIAM J. Comput. 28 (4) (1999) 1198–1214.

[18] D. Avis, D. Bremner, R. Seidel, How good are convex hull algorithms?, Comput. Geom. Theory Appl. 7 (5–6) (1997) 265–302.

[19] F. Aurenhammer, Power diagrams: properties, algorithms and applications, SIAM J. Comput. 16 (1) (1987) 78–96.

[20] D. Rappaport, A convex hull algorithm for discs, and applications, Comput. Geom. Theory Appl. 1 (3) (1992) 171–181.

[21] J.-D. Boissonnat, A. Cérézo, O. Devillers, J. Duquesne, M. Yvinec, An algorithm for constructing the convex hull of a set of spheres in dimension d, Comput. Geom. Theory Appl. 6 (2) (1996) 123–130.

[22] J.-D. Boissonnat, M. I. Karavelas, On the combinatorial complexity of Euclidean Voronoi cells and convex hulls of d-dimensional spheres, in: Proc. 14th ACM-SIAM Sympos. Discrete Algorithms (SODA), 2003, pp. 305–312.

[23] J.-D. Boissonnat, A. Cérézo, J. Duquesne, An output-sensitive algorithm for constructing the convex hull of a set of spheres, in: IFIP Conference on Algorithms and efficient computation, 1992.

[24] F. Nielsen, M. Yvinec, Output-sensitive convex hull algorithms of planar convex objects, Internat. J. Comput. Geom. Appl. 8 (1) (1998) 39–66.

[25] R. Stanley, The upper-bound conjecture and Cohen-Macaulay rings, Stud. Appl. Math. 54 (1975) 135–142.

[26] R. Stanley, The number of faces of simplicial polytopes and spheres, in: J. E. Goodman, J. Malkevitch, E. Lutwak (Eds.), Discrete Geometry and Convexity, Vol. 440 of Annals of the New York Academy of Sciences, New York Academy of Sciences, 1985, pp. 212–223.

[27] P. Agarwal, M. Sharir, Arrangements and Their Applications, in: J.-R. Sack, J. Urrutia (Eds.), Handbook of Computational Geoemtry, Elsevier, 2000, pp. 49–119.

[28] K. L. Clarkson, P. W. Shor, Applications of random sampling in computational geometry, II, Discrete Comput. Geom. 4 (1989) 387–421.

[29] D. P. Dobkin, D. G. Kirkpatrick, Determining the Separation of Preprocessed Polyhedra - A Unified Approach, in: Proceedings of the 17th International Colloquium on Automata, Languages and Programming, ICALP '90, Springer-Verlag, London, UK, UK, 1990, pp. 400–413.