Chapter 7

Halfspace Emptiness

The halfspace emptiness problem asks, given a set of points and a set of halfspaces, whether any halfspace contains a point. In this chapter, we will consider the following formulation of the problem: Given a set of points and hyperplanes, is every point above every hyperplane? Using linear programming [48, 109, 113, 136], we can decide in linear time whether the union of a set of halfspaces is $\mathbb{R}^d$. If it is, then *every* input point must lie in a halfspace; if not, then by an appropriate projective transformation, we can ensure that the halfspaces miss the point $(0, 0, \ldots, 0, \infty)$. If we use the duality transformation $(a_1, a_2, \ldots, a_d) \longleftrightarrow \sum_{i=1}^{d-1} a_i x_i = a_d + x_d$, then a point p is above a hyperplane h if and only if the dual point h^* is above the dual point p^*. Thus, in this formulation, the halfspace emptiness problem is self-dual.

The best known algorithms for this problem were developed for its online version: Given a set of n points, preprocess it to answer halfspace emptiness (or reporting) queries. In two and three dimensions, we can easily build a linear-size data structure, in $O(n \log n)$ time, that allows halfspace emptiness queries to be answered in logarithmic time [5, 39, 57]. In higher dimensions, a randomized algorithm due to Clarkson [47] answers halfspace emptiness queries in time $O(\log n)$ after $O(n^{\lfloor d/2 \rfloor + \varepsilon})$ preprocessing time. Matoušek [107] describes two halfspace emptiness data structures, one answering queries in time $O(n^{1-1/\lfloor d/2 \rfloor} \operatorname{polylog} n)$ time after $O(n \log n)$ preprocessing time, and the other answering queries in time $O(n^{1-1/\lfloor d/2 \rfloor} 2^{O(\log^* n)})$ after $O(n^{1+\varepsilon})$ preprocessing time. Combining Clarkson's and Matoušek's data structures, for a fixed parameter $n \le s \le n^{\lfloor d/2 \rfloor}$, one can answer queries in time $O((n \log n)/s^{1/\lfloor d/2 \rfloor})$ after $O(s \operatorname{polylog} n)$ preprocessing time [107, 3, 29]. For extensions and applications of halfspace range reporting, see [3, 4, 27, 29, 110, 108].

Given n points and m halfspaces, we can solve the offline halfspace emptiness problem in time

$$O\left(n \log m + (nm)^{\lfloor d/2 \rfloor/(\lfloor d/2 \rfloor+1)} \operatorname{polylog}(n+m) + m \log n\right),$$

using either Clarkson's data structure or one of Matoušek's data structures, depending on the relative growth rates of n and m. In two and three dimensions, the time bound simplifies to $O(n \log m + m \log n)$. If $n > m$, we actually solve the problem in the dual, by building a data structure to report if any halfspace contains a query point.

The only lower bound previously known for this problem is $\Omega(n \log m + m \log n)$, in the algebraic decision tree or algebraic computation tree models, by reduction from the set intersection problem [138, 16]. Thus, the two- and three-dimensional algorithms are optimal, but there is still a large gap in dimensions four and higher.

In this chapter, we develop a lower bound of $\Omega(n \log m + n^{2/3}m^{2/3} + m \log n)$ on the complexity of the halfspace emptiness problem in $\mathbb{R}^5$, matching known upper bounds up to polylogarithmic factors. We obtain marginally better bounds in dimensions 9 and higher. Using similar techniques, we also prove slightly better bounds for Hopcroft's problem in dimensions four and higher. Our lower bounds apply to *polyhedral partitioning algorithms*, a restriction of the class of partitioning algorithms introduced in the previous chapter. Informally, a polyhedral partitioning algorithm covers space with a constant number of constant-complexity polyhedra, determines which points and halfspaces intersect which polyhedra, and recursively solves the resulting subproblems.

The basic approach is the same as in the previous chapter. We first define *polyhedral covers*, and develop lower bounds on their combinatorial complexity. The main result of this chapter (Theorem 7.9) states that the running time of a polyhedral partitioning algorithm is bounded below by the polyhedral cover size of its input. The $\Omega(n^{4/3})$ lower bound then follows from the construction of a set of points and hyperplanes in $\mathbb{R}^5$, with all the points above all the hyperplanes, whose every polyhedral cover is that large.

7.1 Projective Polyhedra

Our lower bound argument relies heavily on certain properties of convex polytopes and polyhedra. Many of these properties are more easily proved, and have fewer special cases, if we state and prove them in projective space rather than affine Euclidean space. In particular, developing these properties in projective space allows us to more easily deal with unbounded and degenerate polyhedra and duality transformations. Everything we describe in this section can be formalized algebraically in the language of polyhedral cones and linear subspaces one dimension higher; we will give a much less formal, purely geometric treatment. For more technical details, we refer the reader to Chapters 1 and 2 of Ziegler's lecture notes [161].

The projective space $\mathbb{RP}^d$ can be defined as the set of lines through the origin in $\mathbb{R}^{d+1}$. Every k-dimensional linear subspace of $\mathbb{R}^{d+1}$ induces a $(k-1)$-dimensional *flat* f in $\mathbb{RP}^d$, and its orthogonal complement induces the *dual flat* f^*.

A *projective polyhedron* is a single closed cell, not necessarily of full dimension, in the arrangement of a finite number of hyperplanes in $\mathbb{RP}^d$. A *projective polytope* is a simply-connected projective polyhedron, or equivalently, a projective polyhedron that is disjoint from some hyperplane (not necessarily in its defining arrangement). Every projective polyhedron is (the closure of) the image of a convex polyhedron under some projective transformation, and every projective polytope is the image of a convex polytope. Every flat is also a projective polyhedron.

The *projective span* (or projective hull) of any subset $X \subseteq \mathbb{RP}^d$, denoted $\mathrm{span}(X)$, is the projective subspace of minimal dimension that contains it. The *relative interior* of a projective polyhedron is its interior in the subspace topology of its projective hull. A hyperplane *supports* a polyhedron if it intersects the polyhedron but not its relative interior. A flat has no supporting hyperplanes.

A *proper face* of a projective polyhedron is the intersection of the polyhedron and one or more of its supporting hyperplanes. Every proper face of a polyhedron is a lower-dimensional polyhedron. A *face* of a polyhedron is either a proper face or the polyhedron itself. We write $\Phi \leq \Pi$ to denote that a polyhedron Φ is a face of another polyhedron Π. The *dimension* of a face is the dimension of its projective hull. The dimension of the empty set is taken to be -1. The faces of a polyhedron form a lattice under inclusion. Every projective polyhedron has a face lattice isomorphic to that of a convex polytope, possibly of lower dimension.

The *apex* of a polyhedron Π is the intersection of all its supporting hyperplanes, or equivalently, its unique face of minimum dimension. The apex is empty if and only if the polyhedron is a polytope but not a single point; the apex is the whole polyhedron if and only if the polyhedron is a flat.

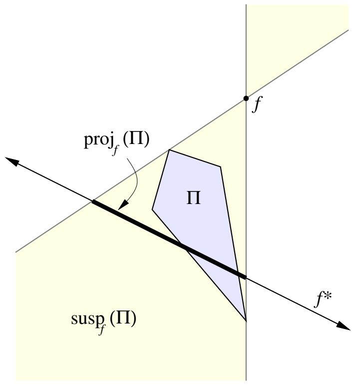

Figure 7.1. The suspension and projection of a polygon by a point.

The *dual* of a polyhedron Π, denoted Π^*, is defined to be the set of points whose dual hyperplanes intersect Π in one of its faces:

$$\Pi^* \triangleq \{p \mid (p^* \cap \Pi) \leq \Pi\}.$$

In other words, $h^* \in \Pi^*$ if and only if h either contains Π, supports Π, or completely misses Π. This definition generalizes both the polar of a convex polytope and the projective dual of a flat. We easily verify that Π^* is a projective polyhedron whose face lattice is the inverse of the face lattice of Π. In particular, Π and Π^* have the same number of faces. See [161, pp. 59–64] and [140, pp. 143–150] for similar definitions and more technical details.

For any subset $X \subseteq \mathbb{RP}^d$ and any flat f, the *suspension of* X *by* f, denoted $\mathrm{susp}_f(X)$, is formed by replacing each point in X by the span of that point and f:

$$\mathrm{susp}_f(X) \triangleq \bigcup_{p \in X} \mathrm{span}(p \cup f).$$

The suspension of a subset of projective space corresponds to an infinite cylinder over a subset of an affine space, at least when the apex of suspension is "at infinity".[1] The *projection of* X *by* f, denoted $\mathrm{proj}_f(X)$, is the intersection the suspension and the dual flat f^*:

$$\mathrm{proj}_f(X) \triangleq f^* \cap \mathrm{susp}_f(X),$$

In particular, $\mathrm{susp}_f(X)$ is the set of all points in $\mathbb{RP}^d$ whose projection by f is in $\mathrm{proj}_f(X)$. The projection of a subset of projective space corresponds to the orthogonal projection of a subset of affine space onto an affine subspace. See Figure 7.1.

[1] Ziegler [161, p. 33] calls this the *elimination* of X.

7.2 Polyhedral Separation

Let P be a set of points, let H be a set of hyperplanes, and let Π be a projective polyhedron in $\mathbb{RP}^d$. We say that Π *separates* P and H if Π contains P and the dual polyhedron Π^* contains the dual points H^*; that is, any hyperplane in H either contains Π, supports Π, or misses Π entirely. Both P and H may intersect the relative boundary of Π. We say that P and H are r-*separable* if there is a polyhedron with at most r faces that separates them.

The proofs of Theorems 6.6 and 6.11 implicitly relied on the following trivial observation: if we perturb a configuration of points and hyperplanes just enough to remove any incidences, and the resulting configuration is monochromatic, then the original configuration must have been loosely monochromatic. The following technical lemma establishes the corresponding, but no longer trivial, property of r-separable configurations. Informally, if a configuration is not r-separable, then arbitrarily small perturbations cannot make it r-separable. First-time readers are encouraged to skip the proof.

Technical Lemma 7.1. *Let H be a set of m hyperplanes in $\mathbb{RP}^d$. For all r, the set of point configurations $P \in (\mathbb{RP}^d)^n$ such that P and H are r-separable is topologically closed.*

Proof: There are two cases to consider: either the hyperplanes in H do not have a common intersection, or they intersect in a common flat. The proof of the second case relies on the first.

Case 1 ($\bigcap H = \emptyset$):

Any polyhedron that separates P and H must be completely contained in a closed d-cell $\mathcal{C}$ of the arrangement of H. Thus, it suffices to show, for each cell $\mathcal{C}$, that the set of n-point configurations contained in $\mathcal{C}$ and r-separable from H is topologically closed. Our approach is to show that this set is actually a compact semialgebraic set.

Fix a cell $\mathcal{C}$. Since every hyperplane in H passes through the apex of $\mathcal{C}$, both $\mathcal{C}$ and any polyhedra it contains must be polytopes. By choosing an appropriate hyperplane "at infinity" that misses $\mathcal{C}$, we can treat $\mathcal{C}$ and any polytopes it contains as *convex* polytopes in $\mathbb{R}^d$.

Let $A = \{a_1, a_2, \ldots, a_\nu\}$ and $B = \{b_1, b_2, \ldots, b_\nu\}$ be two indexed sets of points in $\mathbb{R}^d$, for some integer ν. We say that A *is simpler than* B, written $A \sqsubseteq B$, if for any subset of B contained in a facet of conv(B), the corresponding subset of A is contained in a facet of conv(A).[2] Equivalently, $A \sqsubseteq B$ if and only if for $d+1$ points in B, d of whose vertices lie on a facet of conv(B), the corresponding simplex in A either has the same orientation or is degenerate. Simpler point sets have less complex convex hulls — if $A \sqsubseteq B$, then conv(A) has no more vertices, facets, or faces than conv(B). If both $A \sqsubseteq B$ and $B \sqsubseteq A$, then the convex hulls of A and B are combinatorially equivalent.

If B is fixed, then the relation $A \sqsubseteq B$ can be encoded as the conjunction of at most $O(\nu^{\lfloor d/2 \rfloor + 1})$ algebraic inequalities of the form

$$\begin{vmatrix} a_{i_0 0} & a_{i_0 1} & \cdots & a_{i_0 d} \\ a_{i_1 0} & a_{i_1 1} & \cdots & a_{i_1 d} \\ \vdots & \vdots & \ddots & \vdots \\ a_{i_d 0} & a_{i_d 1} & \cdots & a_{i_d d} \end{vmatrix} \lozenge 0,$$

where $(a_{i_j 0}, a_{i_j 1}, \ldots, a_{i_j d})$ are the homogeneous coordinates of the point $a_{i_j} \in A$, and $\lozenge$ is either $\geq$, $=$, or $\leq$. In every such inequality, the corresponding points $b_{i_1}, b_{i_0}, \ldots, b_{i_d}$ (but not necessarily b_{i_0}) all lie on a facet of conv(B). For every d-tuple of points in B contained in a facet of conv(B), there are $\nu - d$ such inequalities, one for every other point. (If we replace the loose inequalities $\leq, \geq$ with strict inequalities $<, >$, the resulting expression encodes the combinatorial equivalence of conv(A) and conv(B).)

We can encode the statement "P is contained in $\mathcal{C}$ and is r-separable from H" as the following elementary formula:

$$\bigvee_{\nu=1}^{r} \bigvee_{\substack{B \in (\mathbb{R}^d)^\nu \\ \text{conv}(B) \text{ has at most } r \text{ faces}}} \left\{ \begin{array}{l} \exists\, a_1, a_2, \ldots, a_\nu \in \mathcal{C}: \\ \exists\, \lambda_1, \lambda_2, \ldots, \lambda_n \in [0,1]^\nu: \\ \qquad (A \sqsubseteq B) \;\wedge\; \displaystyle\bigwedge_{i=1}^{n} \left(\sum_{j=1}^{\nu} a_j \lambda_{ij} = p_i \;\wedge\; \sum_{j=1}^{\nu} \lambda_{ij} = 1 \right) \end{array} \right\} \tag{7.1}$$

Equivalently, in English:

> For some integer ν, and for some set B of ν points whose convex hull has at most r faces, there exists a set A of ν points in $\mathcal{C}$, such that A is simpler than B (so conv(A) has at most r faces) and every point in P is a convex combination of points in A (with barycentric coordinates Λ).

[2] Every set of points is simpler than itself. It would be more correct, but also unwieldier, to say "A is at least as simple as B".

Since there are only a finite number of combinatorial equivalence classes of convex polytopes with ν vertices [89], the formula is finite, and therefore defines a semi-algebraic set. It remains only to show that this set is closed.

For any fixed ν and B, the set of configurations $P \times A \times \Lambda \in (\mathbb{R}^d)^n \times \mathcal{C}^\nu \times ([0,1]^\nu)^n$ that satisfy the subexpression

$$(A \sqsubseteq B) \;\wedge\; \bigwedge_{i=1}^{n} \left(\sum_{j=1}^{\nu} a_j \lambda_{ij} = p_i \;\wedge\; \sum_{j=1}^{\nu} \lambda_{ij} = 1 \right)$$

is the intersection of the closed convex polytope $\mathcal{C}^{n+\nu} \times [0,1]^{\nu n}$, at most $(\nu - d)r$ closed algebraic halfspaces, νn quadratic surfaces, and νn hyperplanes, and is therefore closed and bounded. It follows that the set of point configurations P that satisfy the subexpression of (7.1) in braces is the projection of a compact set, and is therefore also compact. Finally, the set of configurations P satisfying the entire formula (7.1) is the union of a finite number of compact sets, and therefore must be compact.

This completes the proof of Case 1.

Case 2 ($\bigcap H \neq \emptyset$):

The previous argument will not work in this case, because the cells in the arrangement of H are not simply connected, and thus are not polytopes.

Let P be an arbitrary set of n points in $\mathbb{RP}^d$, such that P and H are *not* r-separable. To prove the lemma, it suffices to show that there is an open neighborhood $\mathcal{U} \in (\mathbb{RP})^d$ with $P \in \mathcal{U}$, such that, for all $Q \in \mathcal{U}$, Q and H are not r-separable.

Let $f = \bigcap H$, and let f^* be its dual flat. Without loss of generality, suppose the points $p_1, p_2, \ldots, p_m \in P$ are disjoint from f, and the points $p_{m+1}, \ldots, p_n \in P$ are contained in f. Denote these two subsets of P by $P \setminus f$ and $P \cap f$, respectively. Either subset may be empty. Note that $\mathrm{proj}_f(P) = \mathrm{proj}_f(P \setminus f)$, since by definition $\mathrm{proj}_f(f)$ is empty.

If any polyhedron Π separates P and H, then its projection $\mathrm{proj}_f(\Pi)$ separates the projected points $\mathrm{proj}_f(P)$ and the lower dimensional hyperplanes $H \cap f^*$. Conversely, if any polyhedron $\Pi \subseteq f^*$ separates $\mathrm{proj}_f(P)$ and $H \cap f^*$ then its suspension $\mathrm{susp}_f(\Pi)$ separates P and H. Thus, P and H are r-separable if and only if $\mathrm{proj}_f(P)$ and $H \cap f^*$ are r-separable.

Since P and H are *not* r-separable, neither are $\mathrm{proj}_f(P)$ and $H \cap f^*$. The lower-dimensional hyperplanes $H \cap f^*$ do not have a common intersection. Thus, Case 1 implies that the set of configurations $P' \in (f^*)^m$ such that P' and $H \cap f^*$ are r-separable is closed. It follows that there is an open set $\mathcal{U}' \subseteq (f^*)^m$, with $\mathrm{proj}_f(P) \in \mathcal{U}'$, such that for all $Q' \in \mathcal{U}'$, Q' and $H \cap f^*$ are not r-separable.

Let $\mathcal{U}'' \subseteq (\mathbb{RP}^d)^m$ be the set of m-point configurations P'' such that $\mathrm{proj}_f(P'') \in \mathcal{U}'$. Clearly, $\mathcal{U}''$ is an open neighborhood of $P \setminus f$, and no configuration in $Q'' \in \mathcal{U}''$ is r-separable from H.

Finally, if Q'' and H are not r-separable, then no superset of Q'' is r-separable from H. Let $\mathcal{U} = \mathcal{U}'' \times (\mathbb{RP}^d)^{n-m}$. Then $\mathcal{U}$ is an open subset of $(\mathbb{RP}^d)^n$ containing P. Since every configuration $Q \in \mathcal{U}$ has a subset Q'' that is not r-separable from H, we conclude that no $Q \in \mathcal{U}$ is r-separable from H, as claimed.

This completes the proof of Case 2, and thus the entire technical lemma. $\square$

The method we used to encode the condition "$\mathrm{conv}(A)$ has at most r faces" may seem somewhat convoluted. If we replace $A \sqsubseteq B$ with "$\mathrm{conv}(A)$ is combinatorially equivalent to $\mathrm{conv}(B)$", we get exactly the same semi-algebraic set, without needing to define the partial order $\sqsubseteq$. Unfortunately, testing whether two convex polytopes are combinatorially equivalent requires *strict* inequalities, whose corresponding semi-algebraic sets are *open*.

7.3 Polyhedral Covers

A r-*polyhedral cover* of a set P of points and a set H of hyperplanes is an indexed set of subset pairs $\{(P_i, H_i)\}$, where $P_i \subseteq P$ and $H_i \subseteq H$ for all i, such that

(1) For each index i, P_i and H_i are r-separable.

(2) For every point $p \in P$ and hyperplane $h \in H$, there is some index i such that $p \in P_i$ and $h \in H_i$.

We emphasize that the subsets P_i are not necessarily disjoint, nor are the subsets H_i. We refer to each subset pair (P_i, H_i) in an r-polyhedral cover as a r-*polyhedral minor*. The *size* of a polyhedral cover is the sum of the sizes of the subsets P_i and H_i.

Let $\pi_r(P, H)$ denote the size of the smallest r-polyhedral cover of P and H. Let $\pi^\circ_{d,r}(n, m)$ denote the maximum of $\pi_r(P, H)$ over all sets P of n points and H of m hyperplanes in $\mathbb{RP}^d$ with no incidences. When the subscript r is omitted, we take it to be a constant. Finally, recall that $I(P, H)$ denotes the number of point-hyperplane incidences between P and H.

Lemma 7.2. *Let* P *be a set of* n *points and* H *a set of* m *hyperplanes, such that no subset of* s *hyperplanes contains* t *points in its intersection. If* P *and* H *are* r*-separable, then* $I(P,H) \leq r(s+t)(n+m)$.

Proof: Let Π be a polyhedron with r faces that separates P and H. For any point $p \in P$ and hyperplane $h \in H$ such that p lies on h, there is some face $f \leq \Pi$ that contains p and is contained in h. For each face f of Π, let P_f denote the points in P that are contained in f, and let H_f denote the hyperplanes in H that contain f.

Since no set of s hyperplanes can all contain the same t points, it follows that for all f, either $|P_f| < t$ or $|H_f| < s$. Thus, we can bound $I(P,H)$ as follows.

$$I(P,H) \leq \sum_{f \leq \Pi} I(P_f, H_f) = \sum_{f \leq \Pi} (|P_f| \cdot |H_f|) \leq (s+t) \sum_{f \leq \Pi} (|P_f| + |H_f|)$$

Since Π has r faces, the last sum counts each point and hyperplane at most r times. □

The next lemma shows that sufficiently small perturbations of a configuration cannot decrease its polyhedral cover size.

Lemma 7.3. *Let* P *be a set of* n *points and* H *a set of* m *hyperplanes in* $\mathbb{RP}^d$. *For all point configurations* $Q \in (\mathbb{RP}^d)^n$ *sufficiently close to* P, $\pi_r(Q,H) \geq \pi_r(P,H)$.

Proof: Let P' be a subset of P, and for any other set Q of $|P|$ points, let Q' be the corresponding subset of Q. Let H' be a subset of H. Lemma 7.1 implies that there is an open set $\mathcal{U}(P',H') \subseteq (\mathbb{RP}^d)^n$ such that if $Q \in \mathcal{U}(P,H)$ and Q' and H' are r-separable, then P' and H' are r-separable.

Let $\mathcal{U}$ be the intersection of these $2^n 2^m$ open sets:

$$\mathcal{U} = \bigcap_{P' \subseteq P} \bigcap_{H' \subseteq H} \mathcal{U}(P',H').$$

For all $Q \in \mathcal{U}$, every r-polyhedral minor of Q and H corresponds to a r-polyhedral minor of P and H. Thus, for any r-polyhedral *cover* of Q and H, there is a corresponding r-polyhedral cover of P and H with exactly the same size. □

Theorem 7.4. $\pi_2^\circ(m,n) = \Omega(n + n^{2/3}m^{2/3} + m)$.

Proof: Let P be a set of n points and H a set of m lines in the plane with $I(P,H) = \Omega(n + n^{2/3}m^{2/3} + m)$, as described by Lemma 6.4.

Consider subsets $P_i \subseteq P$ and $H_i \subseteq H$ such that P_i and H_i are r-separable. Since two distinct lines in the plane intersect in a single point, Lemma 7.2 implies that $I(P_i, H_i) \leq 4r(|P_i| + |H_i|)$. It follows that any collection of r-polyhedral minors that includes every incidence between P and H must have size at least $I(P,H)/4r$. Thus, $\pi_{2,r}(P,H) = \Omega(n + n^{2/3}m^{2/3} + m)$ for any constant r.

Finally, Lemma 7.3 implies that we can perturb P slightly, removing all the incidences, without decreasing the polyhedral cover size. □

A similar argument derives the following lower bound from Lemma 6.13. As usual, the best lower bounds from each dimension have been combined into a single expression.

Theorem 7.5. $\pi^{\circ}_d(n,m) = \Omega\left(\sum_{i=1}^{d}\left(n^{1-2/i(i+1)}m^{2/(i+1)} + n^{2/(i+1)}m^{1-2/i(i+1)}\right)\right)$.

Following the terminology in the previous chapter, we call say that a point-hyperplane configuration in $\mathbb{R}^d$ is *monochromatic* if every point lies above every hyperplane. Monochromatic configurations have no incidences. Let $\hat{\pi}_{d,r}(m,n)$ denote the maximum of $\pi_r(P,H)$ over all monochromatic configurations of n points and of m hyperplanes in $\mathbb{R}^d \subset \mathbb{RP}^d$.

Lemma 6.15 and the arguments in Theorem 7.4 immediately imply the following lower bound.

Theorem 7.6. $\hat{\pi}_5(n,m) = \Omega(n + n^{2/3}m^{2/3} + m)$.

We can improve this bound very slightly in higher dimensions. Define the family of functions $\sigma_d : \mathbb{R}^{d+1} \to \mathbb{R}^{\binom{d+2}{2}}$ as follows.

$$\sigma_d(x_0, x_1, \ldots, x_d) \triangleq (x_0^2, x_1^2, \ldots, x_d^2, \sqrt{2}\,x_0x_1, \sqrt{2}\,x_0x_2, \ldots, \sqrt{2}\,x_{d-1}x_d)$$

For any two vectors $u, v \in \mathbb{R}^{d+1}$, we have $\langle\sigma_d(u), \sigma_d(v)\rangle = \langle u, v\rangle^2$, where $\langle\cdot,\cdot\rangle$ is the standard vector inner product. In a more geometric setting, σ_d maps points and hyperplanes in $\mathbb{R}^d$, represented in homogeneous coordinates, to points and hyperplanes in $\mathbb{R}^D$, also in homogeneous coordinates, where $D = \binom{d+2}{2} - 1 = d(d+3)/2$. If the point p is incident to the hyperplane h, then $\sigma_d(p)$ is also incident to $\sigma_d(h)$; otherwise, $\sigma_d(p)$ is above $\sigma_d(h)$.

Lemma 6.13 now immediately implies:

Theorem 7.7. *For all* $D \geq d(d+3)/2$,

$$\hat{\pi}_D(n,m) = \Omega\left(\sum_{i=1}^{d}\left(n^{1-2/i(i+1)}m^{2/(i+1)} + n^{2/(i+1)}m^{1-2/i(i+1)}\right)\right).$$

It is clear that $\hat{\pi}_3(n,m) = \Theta(n+m)$, since both the convex hull of any set of points and the upper envelope of any set of planes have linear size triangulations. We conjecture that $\hat{\pi}_d(n,n) = \Omega(n^{\lfloor d/2\rfloor/(\lfloor d/2\rfloor+1)})$ for all d, but are unable to prove this when $d = 4$ or $d \geq 6$.

7.4 Polyhedral Partitioning Algorithms

A *polyhedral partition graph* is a partition graph in which every query region is a projective polyhedron with at most r faces, for some fixed constant r. A typical value for r might be 2^{d+1} (every query region is a simplex) or $3^d + 1$ (every query region is a combinatorial cube). We still do not require the query regions to be disjoint. A *polyhedral partitioning algorithm* is a partitioning algorithm whose partition graph is polyhedral.

Given sets P of points and H of hyperplanes in $\mathbb{R}^d$ as input, a polyhedral partitioning algorithm for the halfspace emptiness problem constructs a polyhedral partition graph and uses it to drive the following divide-and-conquer process, which is slightly different from that used to solve Hopcroft's problem. As before, the algorithm starts at the root and proceeds through the graph in topological order, and at every node except the root, points and hyperplanes are passed in along incoming edges from preceding nodes. For each node v, let $P_v \subseteq P$ denote the points and $H_v \subseteq H$ the hyperplanes that reach v; at the root, we have $P_{root} = P$ and $H_{root} = H$. If v is a primal node, then for every query region $\Pi \in \mathcal{R}_v$, the points in P_v that are contained in Π and the hyperplanes in H_v *whose lower halfspaces* intersect Π traverse the corresponding outgoing edge. If v is a dual node, then for every $\Pi \in \mathcal{R}_v$, the points $p \in P_v$ whose dual hyperplanes p^* intersect *or lie above* Π and the hyperplanes $h \in H_v$ whose dual points h^* are contained in Π traverse the corresponding outgoing edge.[3]

[3]Alternately, we could let the points whose dual hyperplanes intersect Π and the hyperplane whose dual points intersect or lie below Π traverse the edge. Using this alternate formulation has no effect on our results. In fact, we can allow our partition graphs to have *four* types of non-leaf nodes — primal or dual; point/halfspace or ray/hyperplane — without changing our results, or even significantly altering their proofs.

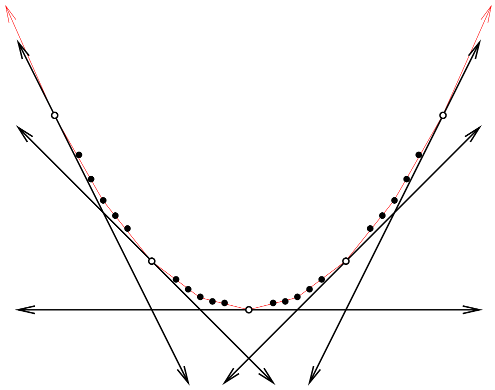

Figure 7.2. Worst-case configuration for halfspace emptiness. Tangent points are shown in white.

To solve the halfspace emptiness problem, a partitioning algorithm reports that all the points are above the hyperplanes if and only if no leaf in the partition graph is reached by both a point and a hyperplane. Clearly, if such an algorithm reports that every point is above every hyperplane, it must be correct.

Theorem 7.8. *Any polyhedral partitioning algorithm that solves the halfspace emptiness problem in $\mathbb{R}^d$, for any $d \geq 2$, must take time $\Omega(n \log m + m \log n)$ in the worst case.*

Proof: It suffices to consider the following configuration, where n is a multiple of m. P consists of n points on the unit parabola $x_d = x_1^2/2$ in $\mathbb{R}^d$, and H consists of m hyperplanes tangent to the parabola and orthogonal to the (x_1, x_d) plane, placed so that n/m points lie between adjacent points of tangency. All the points in P are above all the hyperplanes in H. The dual points H^* also lie on the parabola $x_d = x_1^2/2$, and the dual hyperplanes P^* are also tangent to that parabola.

Following the proof of Theorem 6.18, for any point, we call the hyperplane whose tangent point is closest in the positive x_1-direction the point's *partner*. Every hyperplane is the partner of n/m points. A node v *splits* a point-hyperplane pair if both the point and the hyperplane reach v, and none of the outgoing edges of v is traversed by both the point and the hyperplane. A hyperplane h is *active at level* k if no node in the first k levels splits h from any of its partners.

Suppose v is a primal node. For each hyperplane h that v splits from one of its partner points p, mark some query polyhedron $\Pi \in \mathcal{R}_v$ that contains p but misses h. Since Π has at most r faces, the intersection of Π and the parabola consists of at most r arcs, so

Π can be marked at most r times. Since there are at most Δ polyhedra in $\mathcal{R}_v$, at most $r\Delta$ hyperplanes become inactive at v. Similarly, if v is a dual node, then v splits at most $r\Delta$ points from their partners.

Thus, the number of hyperplanes that are inactive at level k is less than $r\Delta^{k+2}$. In particular, at level $\lfloor \log_\Delta(m/r) \rfloor - 3$, at least $m(1 - 1/\Delta)$ hyperplanes are still active. It follows that at least $n(1 - 1/\Delta)$ points each traverse at least $\lfloor \log_\Delta(m/r) \rfloor - 3$ edges. We conclude that the total running time of the algorithm is at least

$$n(1 - 1/\Delta)(\lfloor \log_\Delta(m/r) \rfloor - 3) = \Omega(n \log m).$$

Symmetric arguments establish a lower bound of $\Omega(m \log n)$ when $n < m$. □

The restriction to polyhedral partitioning algorithms is necessary for the lower bound to hold, since the problem can be solved in *linear* time in the generic partitioning algorithm model. The partition graph consists of a single primal node with two query regions: the convex hull of the points and its complement. If every point is above every hyperplane, then no hyperplane intersects the convex hull of the points.

This lower bound is tight, up to constant factors, in two and three dimensions.

Theorem 7.9. *Let $\mathcal{A}$ be a polyhedral partitioning algorithm that solves the halfspace emptiness problem, and let P be a set of points and H a set of hyperplanes, such that every point is above every hyperplane. Then $T_{\mathcal{A}}(P, H) = \Omega(\pi(P, H))$.*

Proof: From the proof of Theorem 6.19, we immediately have the following inequality:

$$\begin{aligned}\Delta \cdot T_{\mathcal{A}}(P, H) \geq & \sum_{\text{primal edge } e} \Big(\#\text{ points traversing } e + \#\text{ hyperplanes missing } e\Big) + \\ & \sum_{\text{dual edge } e} \Big(\#\text{ hyperplanes traversing } e + \#\text{ points missing } e\Big)\end{aligned}$$

For each primal edge e, let P_e be the set of points that traverse e, and let H_e be the set of hyperplanes that miss e. The edge e is associated with a query polyhedron Π. Every point in P_e is contained in Π, and every hyperplane in H_e is disjoint from Π. Since Π has at most r faces, P_e and H_e are r-separable.

Similarly, for each dual edge e, let H_e be the hyperplanes that traverse it, and P_e the points that miss it. The associated query polyhedron Π separates the dual points H_e^* and the dual hyperplanes P_e^*. By the definition of dual polyhedra, Π^* separates P_e and H_e.

For every point $p \in P$ and hyperplane $h \in H$, there is node that splits them (since otherwise the algorithm would return the wrong answer) and thus some edge e with $p \in P_e$ and $h \in H_e$. It follows that the collection of subset pairs $\{(P_e, H_e)\}$ is an r-polyhedral cover of P and H whose size is at least $\Delta \cdot T_{\mathcal{A}}(P, H)$ and, by definition, at most $\pi_r(P, H)$. □

We emphasize that every point must be above every hyperplane for this lower bound to hold. If some point lies below a hyperplane, then the trivial partitioning algorithm, whose partition graph consists of a single leaf, correctly "detects" the pair.

Corollary 7.10. *The worst-case running time of any polyhedral partitioning algorithm that solves the halfspace emptiness problem in* $\mathbb{R}^D$ *is* $\Omega(n \log m + n^{2/3}m^{2/3} + m \log n)$ *for all* $D \geq 5$ *and*

$$\Omega\left(n \log m + \sum_{i=2}^{D}\left(n^{1-2/i(i+1)}m^{2/(i+1)} + n^{2/(i+1)}m^{1-2/i(i+1)}\right) + m \log n\right)$$

for all $D \geq d(d+3)/2$.

Proof: Theorems 7.8 and 7.9 together imply that the worst case running time is $\Omega(n \log m + \hat{\pi}_d(n, m) + n \log m)$. The lower bounds then follow immediately from Theorem 7.6 and 7.7. □

Partitioning algorithms for the halfspace emptiness problem can (and do [47, 107]) apply a version of the "containment shortcut" described in Section 6.3.4. If some query region lies entirely in a hyperplane's lower halfspace, then the hyperplane need not traverse the corresponding edge. Instead, if any point lies in that region, we immediately halt and report that some point is below a hyperplane. Although this shortcut decreases the running time of the algorithm, we easily verify that Theorem 7.9 still applies in the faster model.

Our techniques also allow us to slightly improve earlier lower bounds for Hopcroft's problem in higher dimensions, matching our lower bounds for the counting problem in Corollary 6.22.

Theorem 7.11. *Let* $\mathcal{A}$ *be a polyhedral partitioning algorithm that solves Hopcroft's problem, and let* P *be a set of points and* H *a set of hyperplanes such that* $I(P, H) = 0$. *Then* $T_{\mathcal{A}}(P, H) = \Omega(\pi(P, H))$.

Combining this theorem with Theorem 7.5, we conclude:

Corollary 7.12. *The worst-case running time of any polyhedral partitioning algorithm that solves Hopcroft's problem in* $\mathbb{R}^d$ *is*

$$\Omega\left(n \log m + \sum_{i=2}^{d}\left(n^{1-2/i(i+1)} m^{2/(i+1)} + n^{2/(i+1)} m^{1-2/i(i+1)}\right) + m \log n\right).$$

7.5 Conclusions and Open Problems

We have proven a lower bound of $\Omega(n^{4/3})$ on the complexity of the offline halfspace emptiness problem in five dimensions. Our lower bounds apply to a broad class of geometric divide-and-conquer algorithms that recursively partition their input by a division of space into constant-complexity polyhedra.

The most obvious open problem is to improve our results. The correct complexity in d dimensions is almost certainly $\Theta(n^{2-2/\lfloor d/2 \rfloor})$, but we achieve this bound only when $d = 5$. In particular, the four dimensional case is wide open. It is not even known whether the four-dimensional halfspace emptiness problem is harder, or easier, than Hopcroft's problem in the plane [75].

The inner product doubling maps σ_d can be used to reduce Hopcroft's problem in $\mathbb{R}^d$ to halfspace emptiness in $\mathbb{R}^{d(d-3)/2}$ in linear time. Is there an efficient reduction from Hopcroft's problem to halfspace emptiness that only increases the dimension by a constant factor (preferably two)?

Our lower bounds are ultimately based on the construction of point-hyperplane configurations whose incidence graphs have several edges but no large complete bipartite subgraphs. Better such configurations would immediately lead to better lower bounds. Lower bounds in the Fredman/Yao semigroup arithmetic model have a similar basis. For example, Chazelle's lower bounds for offline simplex range searching [38] is based on a similar configuration of points and slabs. (See also [44].) Can we derive better polyhedral cover size bounds for points and hyperplanes from these configurations?

Another open problem is to prove tight lower bounds for *online* halfspace range query problems. Brönnimann, Chazelle, and Pach [23] have proven time-space tradeoffs for halfspace counting data structures in the Fredman/Yao semigroup model. Specifically, they prove that any data structure that uses space $n \le s \le n^d$ has worst-case query time

$$\Omega\left(\frac{(n/\log n)^{1-\frac{d-1}{d(d+1)}}}{s^{1/d}}\right).$$

Results of Matoušek [111] imply the upper bound $O((n/s^{1/d}) \operatorname{polylog} n)$, which is almost certainly optimal (except possibly for the polylog factor), so the lower bounds have significant room for improvement. Chazelle and Rosenberg [44] have developed quasi-optimal tradeoffs for simplex reporting data structures in Tarjan's pointer machine model, but no lower bounds are known for *halfspace* reporting. No lower bounds are known for online halfspace emptiness queries in any model of computation. One possible approach, suggested by Pankaj Agarwal (personal communication), is to model range query data structures with partition graphs and to prove tradeoffs between the total size of the graph (space) and the size of the subgraph induced by a query range (time).

A problem closely related to halfspace range searching is linear programming. The best known data structures of linear programming queries are based on data structures for halfspace emptiness [110] and halfspace reporting queries [27]. However, no nontrivial lower bounds are known for linear programming queries in any model of computation. One application of particular interest is deciding, given a set of points, whether every point is a vertex of the set's convex hull. Bounds for this problem closely match the best known bounds for halfspace emptiness [29], but the best known lower bound is $\Omega(n \log n)$. It seems unlikely that a lower bounds can be derived for this problem in the partitioning algorithm model, since the extremity of a point depends on several other points arbitrarily far away. Perhaps the techniques developed in Part I are more applicable.

Finally, extending our lower bounds into more traditional models of computation, such as algebraic decision trees or algebraic computation trees, is an important and extremely difficult open problem. A lower bound bigger than $\Omega(n \log m + m \log n)$ for *any* offline range searching problem in these models would be a major breakthrough.

> *Here, however, a word of warning may be in order: do* not *try to visualize* n*-dimensional objects for* $n \geq 4$*. Such an effort is not only doomed to failure—it may be dangerous to your mental health. (If you do succeed, then you are in trouble.) To speak of* n*-dimensional geometry with* $n \geq 4$ *simply means to speak of a certain part of algebra.*
>
> — Vašek Chvátal, *Linear Programming*, 1983

> *This is wrong, and even Chvátal acknowledges the fact that the correspondence between intuitive geometric terms and algebraic machinery can be used in both ways.*
>
> — Günter Ziegler, *Lectures on Polytopes*, 1995