Bounded boxes, Hausdorff distance, and a new proof of an interesting Helly-type theorem.

Nina Amenta *
The Geometry Center
1300 South Second Street
Minneapolis, MN. 55454

Abstract

In the first part of this paper, we reduce two geometric optimization problems to convex programming: finding the largest axis-aligned box in the intersection of a family of convex sets, and finding the translation and scaling that minimizes the Hausdorff distance between two polytopes. These reductions imply that important cases of these problems can be solved in expected linear time. In the second part of the paper, we use convex programming to give a new, short proof of an interesting Helly-type theorem, first conjectured by Grünbaum and Motzkin.

1 Introduction

Linear programming is a popular tool in computational geometry. A related problem is

Convex Programming
Input: *A finite family H of closed convex sets in E^d and a convex function f.*
Output: *The minimum of f over $\bigcap H$, the intersection of H.*

Remember that a function $f : E^d \rightarrow \mathcal{R}$ is convex when $f(\lambda a + (1-\lambda)b) \leq \lambda f(a) + (1-\lambda) f(b)$, for all $a, b \in E^d$ and $0 \leq \lambda \leq 1$.

*email: nina@geom.umn.edu. Most of this work was done at the University of California, Berkeley, where the author was supported by a U.C. Presidents Dissertation Year Fellowship and an AT&T Graduate Research Program for Women Grant.

The elements of the family H are called the *constraints*. Like a linear program, a convex program has a single global minimum which is determined by a subfamily of constraints. And under reasonable computational assumptions on the constraints and the convex function, fixed-dimensional convex programming can be solved in expected linear (in $|H|$) time by any of the randomized combinatorial linear programming algorithms of [C90], [S90], [MSW92]. These algorithms are combinatorial in the sense that they operate by searching the subfamilies of constraints for one that defines the minimum; a more classic example of a combinatorial LP algorithm the simplex algorithm.

We give two new reductions from geometric optimization problems of practical interest to convex programming. The first problem is finding the largest volume axis-aligned box in the intersection of a family K of convex sets. Like the ubiquitous bounding box, this *bounded box* can be used in heuristics to approximate a more complicated volume. When the dimension is fixed and the elements of K are linear halfspaces (or otherwise easy to compute with), this reduction implies an expected $O(n)$ time algorithm, $n = |K|$. The planar case of this problem was raised in the context of a heuristic for packing clothing pattern pieces [DMR93]. They give an $O(n\alpha(n)\log n)$ algorithm to find the maximum area axis-aligned rectangle in any simple polygon in E^2. Our reduction may be useful for higher dimensional packing heuristics as well, and also when objects are to be decomposed into collections of axis-aligned boxes (eg. in ray tracing, [AK89],pp 219-23).

Our second reduction concerns some particular cases of the much-studied problem of minimizing the Hausdorff distance between two objects under a group of transformations. The *one-directional Hausdorff distance* from a set A to a set B, $\vec{H}(A, B)$, is the maximum distance from any point in A to the nearest point in B. The *Hausdorff distance* between A and B, $H(A, B)$, is $\max\{\vec{H}(A, B), \vec{H}(B, A)\}$. The Hausdorff distance is used in pattern recognition and computer vision as a a measure of the difference in shape between the two sets; typically a family P of critical points is extracted from an image and compared with a stored template A. The problem is to choose a transformation from some family which, when applied to A, minimizes $H(P, A)$ or $\vec{H}(P, A)$.

For convex polytopes A, B we reduce the problem of finding the translation and scaling which minimizes $H(A, B)$ to a convex program. In the plane, this convex program can be solved in expected $O(n)$ time, where n is the total number of vertices. A similar convex program, motivated by the scenario above, finds the translation which minimizes the $\vec{H}(P, A)$, where P is a point set and A is a convex polytope. Again in the planar case, this convex program requires expected $min\{O(mn), O(n \lg n + m)\}$) time, where $n = |P|$ and A has m vertices.

There are many results on minimizing the Hausdorff distance between various objects under different groups of motions. Most of them solve more difficult problems and require more time. For planar point sets P, Q, an algorithm of [HKS91] finds the translation that minimizes $H(P, Q)$ in $O(mn(m + n) \lg(mn))$ time, where $n = |P|, m = |Q|$. Allowing rotation as well, an algorithm of [CGHKKK94] minimizes $H(P, Q)$ in $O(m^2n^2(m+n) \lg^2 mn)$ time. Algorithms for measuring the Hausdorff distance for fixed polygons, allowing no transformations at all, require $O(n \lg n)$ time for simple polygons [ABB91], and $O(n)$ for convex polygons [A83]. The algorithms implied by our reductions are also comparatively simple and implementable.

We turn from these practical issues to a theoretical question from combinatorial geometry. We give a new proof of the theorem of [Mo73]:

Theorem 1.1 *Let C_d^k be the family of all sets in R^d consisting of the disjoint union of at most k closed convex sets. Let $I_d^k \subset C_d^k$ be a subfamily with the special property that it is closed under intersection. Then the intersection of I_d^k is nonempty if and only if the intersection of J is nonempty, for all $J \subseteq I_d^k$ with $|J| \leq k(d + 1)$.*

This is called a *Helly-type theorem* because of its combinatorial structure: the whole family of sets has nonempty intersection if and only if all of its constant-size subfamilies do. There are many Helly-type theorems; it is appealing to think that there is some fundamental topological property underlying them all. This theorem is interesting because it suggests this fundamental property might be that the intersection of every subfamily is somehow homologically of constant complexity.

Grünbaum and Motzkin conjectured Theorem 0.1 [GM61], and proved the case $k = 2$, using a more general axiomatic structure in place of convexity. The case $k = 3$ was proved by Larman [L68]. Morris settled the conjecture in his thesis. His proof, however, is quite long (69 pages) and involved, and a better proof has been called for [E93].

Using convex programming, we give a short and insightful proof. Our approach is to introduce a function f, and then show that the problem of minimizing f over a family I_d^k belongs to the class *GLP* (for Generalized Linear Programming). Informally, GLP is the class of problems which can be solved by combinatorial LP algorithms. This is already interesting, as an example of a problem which is combinatorially similar to LP although geometrically the intersection of the constraints fails not only to be convex, but even to be connected. Theorem 0.1 follows by applying an easy theorem from [A93], that there is a Helly-type theorem about the constraint set of every GLP problem.

2 Setup

In this section we give some background on convex programming, GLP, and the combinatorial LP algorithms.

First we give the formal definition of GLP, using an abstract framework due to Sharir and Welzl [SW92]. A GLP problem is a pair (H, w), where H is a family of constraints, generally sets, and $w : 2^H \to \Lambda$ is a function which takes a subfamily of constraints to an element of a totally ordered set Λ. Λ has a special maximum element Ω. When

$w(G) < \Omega$, G is *feasible*, and otherwise G is *infeasible.*

In convex programming, for example, $w(G)$, for $G \subseteq H$, is the minimum of f over $\bigcap G$, and if $\bigcap G = \emptyset$, then $w(G) = \Omega$. In order to ensure that w is defined on every subfamily of constraints, we may have to enclose the problem in a compact "bounding box". We say that w is the objective function *induced* by f.

(H, w) is a GLP problem if

1. For all $F \subseteq G \subseteq H$: $w(F) \leq w(G)$

2. For all $F \subseteq G \subseteq H$ such that $w(F) = w(G)$ and for each $h \in H$:
$w(F + h) > w(F)$ if and only if $w(G + h) > w(G)$
(by $F + h$, we mean $F \cup \{h\}$)

A *basis* is a subfamily $G \subseteq H$ such that $w(G-h) < w(G)$, for all $h \in G$. The *combinatorial dimension* of a GLP is the maximum cardinality of any feasible basis.

Every convex program meets Condition 1, since adding more constraints to a problem can only increase the minimum. It is also well-known that a d-dimensional convex program has combinatorial dimension d; there is a formal proof in [A93'].

But notice that it is not the case that every convex program is a GLP problem, since it may fail to satisfy Condition 2. Condition 2 is always satisfied, however, when the minimum of f over the intersection of every subfamily $G \subseteq H$ is achieved by exactly one point. This observation implies that the following cases of convex programming are always GLP problems of combinatorial dimension d.

A function $f : E^d \rightarrow \mathcal{R}$ is *strictly convex* when $f(\lambda a + (1 - \lambda)b) < \lambda f(a) + (1 - \lambda)f(b)$, for all $a, b \in E^d$ and $0 < \lambda < 1$.

Strictly convex programming
Input: *A family H of compact convex subsets of E^d, and a strictly convex function f.*
Output: *The minimum of f over $\bigcap H$.*

The minimum is achieved at one unique point since any point x on the the line segment between any two feasible points y, z with $f(y) = f(z)$ has $f(x) < f(y) = f(z)$.

The *lexicographic function* $f : E^d \rightarrow \mathcal{R}^d$ takes a point to it's coordinates $x = < x_1, \ldots, x_d >$, which are totally ordered so that $x > y$ if $x_1 > y_1$, or if $x_1 = y_1$ and $x_2 > y_2$, and so on.

Lexicographic convex programming
Input: *A family H of compact convex subsets of E^d, and the lexicographic function f.*
Output: *The minimum of f over $\bigcap H$.*

The minimum is certainly achieved at a single point, since each point has a unique value. But this is not, strictly speaking, a convex programming problem, since f is not a function into $\mathcal{R}$. We observe, however, that f is related to the linear function

$$g_\epsilon(x) = x_1 + \epsilon x_2 + \epsilon^2 x_3 + \ldots + \epsilon^{d-1} x_d$$

for infinitessimally small ϵ. To make this relationship precise, we adopt the terminology of [M94], and say that an objective function v is a *refinement* of a function w when, for $F, G \subseteq H$, $w(G) > w(F)$ implies $v(G) > v(F)$.

Let w and v_ϵ be the objective functions induced by f and g_ϵ, respectively. For any finite family H of constraints, there is some ϵ small enough so that v_ϵ is a refinement of w. Note that there may be no ϵ small enough that $v_\epsilon(F) = v_\epsilon(G)$ whenever $w(F) = w(G)$, as illustrated in figure 1. Finally,

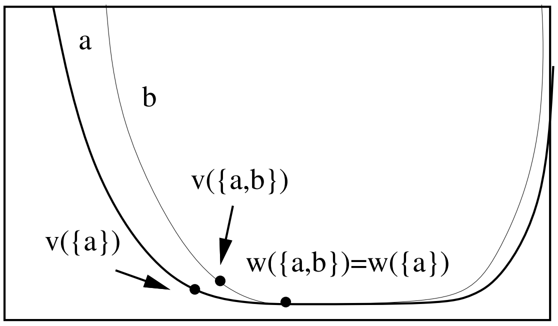

Figure 1: Minima are different under v_ϵ

we observe that if a function w has some refinement v such that (H, v) meets Condition 1 and has combinatorial dimension d, then so does (H, w). So Lexicographic Convex Programming is a GLP problem of combinatorial dimension d, since (H, v_ϵ) is a convex program and so meets Condition 1 and has combinatorial dimension d, and (H, w) meets Condition 2 as well.

The computational requirement under which any of the combinatorial LP algorithms can be applied to any GLP problem is that there is a subroutine available for the following problem:

Basis computation

Input: *A basis G and a constraint h.*
Output: *A basis $G' \subseteq G + h$ such that $w(G') = w(G + h)$.*

This operation corresponds to a pivot step in the simplex algorithm. In d-dimensional convex programming, a basis computation minimizes f over $\bigcap G$, where $|G| \leq d+1$. When a basis computation can be done in constant time, then any of [C90], [S90], [MSW92] require expected $O(n)$ time, where $n = |H|$. All of these algorithms have been implemented for LP, and the algorithm in [S90] has also been applied to the particular convex program of finding the smallest ball enclosing a family of points in E^d [W91].

3 Bounded boxes

In this section we prove

Theorem 3.1 *Finding the largest volume axis-aligned box in the intersection of a family K of n convex bodies in E^d is a strictly convex program in E^{2d} with $2^d n$ constraints.*

Proof sketch: We parameterize an axis-aligned box by a pair of vectors $\mathbf{x}, \mathbf{a} \in \mathcal{R}^{\mathbf{d}}$, where $x_1, \ldots, x_d$ are the coefficients of the lexicographically minimum vertex of the box, and $a_1, \ldots, a_d$ are positive offsets in each coordinate direction. This parameteriztion defines a space of boxes, in which we will construct a convex program.

For each convex body $C \in K$, we will define 2^d constraints, one for each box vertex. Note that a box is contained in a convex body C if and only if all of its vertices are. Let us label the vertices with 0-1 vectors in the natural way, so that $(0, \ldots, 0)$ is the lexicographic minimum corner of the box. The set of boxes for which vertex $\mathbf{u}$ is contained in C is $h = \{\mathbf{x}, \mathbf{a} \mid \mathbf{x} + (\mathbf{u} \otimes \mathbf{a}) \in \mathbf{C}\}$ (here $\otimes$ is coordinate-wise multiplication, and $+$ is translation). This is convex.

To prevent the largest volume box determined any subproblem from being unbounded, we require the bounded box to be contained in a very large bounding box which is guranteed to contain $\bigcap K$. This adds one more convex constraint to the problem in the space of boxes.

It remains to show that maximizing the volume of the box corresponds to minimizing some convex function over $\bigcap H$. The volume of a box (x, a), negated, is given by

$$g(a) = -\prod_{i=1}^{d} a_i$$

with all the a_i constrained to be positive. This is not a convex function, but

$$f(a) = -\log\left(\prod_{i=1}^{d} a_i\right) = -\sum_{i=1}^{d} \log(a_i)$$

is a strictly convex function. Minimizing f over $\bigcap H$ is a strictly convex programming problem in E^{2d}.
$\square$

When the elements of K are linear halfspaces, then the constraints are as well. When the elements of K are of constant complexity, so that a basis computation requires $O(1)$ time for fixed d, the largest volume axis-aligned box can be found in expected $O(n)$ time. Note that f and g are minimized at the same point, so the basis computation may be implemented using g.

4 Hausdorff distance

Now we consider the problem of finding the translation and scaling which minimize the Hausdorff distance between two convex polytopes A and B in E^d. The boundary of the set of points at distance λ from A (that is, the Minkowski sum of A with the closed disk of radius λ centered at the origin) is called the λ-*offset surface.* $\vec{H}(B, A) \leq \lambda$ when

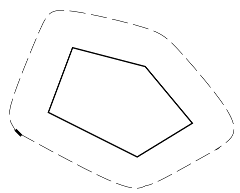

Figure 2: Offset surface

every vertex of B lies within the λ-offset surface of A.

We can think of the scaling and translation transformations as being applied to A alone. We define a $(d+2)$-dimensional *transformation space* in which the coordinates of each point represent a

d-dimensional translation vector τ to be applied to A, a scale factor σ to be applied to A, and an offset distance λ. For any point $b \in B$, let $H^b_{(\tau,\sigma,\lambda)}$ be the subset of transformation space such that b is within the λ-offset surface of the homothet of A scaled by σ and translated by τ. Similarly, let $H^a_{(\tau,\sigma,\lambda)}$ be the set of translations and scalings of A which put a point $a \in A$ inside the λ-offset surface of B.

Theorem 4.1 *The scaling and translation that minimizes the Hausdorff distance between two polytopes in E^d can be found by a lexicographic convex program in E^{d+2}.*

Proof: Let $V(A)$ and $V(B)$ be the vertex sets of A and B, respectively. Consider a vertex $b \in V(B)$. Fixing $\sigma = 1$ and $\lambda = 0$, we find that the set of translations of A which cover b, $H^b_{(\tau,1,0)}$, is itself a translate of the convex set $-A$. Allowing σ to vary, we find that $H^b_{(\tau,\sigma,0)}$ is a cone over $-A$, also convex, into the direction of the σ coordinate. Finally, we allow λ to vary as well. Notice that, as λ varies, the set of disks of radius λ centered at the origin forms a convex cone, C. $H^b_{(\tau,\sigma,\lambda)}$ is the Minkowski sum of $H^b_{(\tau,\sigma,0)}$ with C. This is convex because the Minkowski sum of convex bodies is convex. Each vertex of B produces one such convex constraint.

Now consider a vertex $a \in V(A)$. Again $H^a_{(\tau,1,0)}$ is convex, this time a translate of $-B$. As σ varies, A scales, and vertex a moves in some direction v, along line $a + \sigma v$. So $H^a_{(\tau,\sigma,0)}$ is a convex cylinder over $-B$, and $H^a_{(\tau,\sigma,\lambda)}$ is the Minkowski sum of the cylinder with C.

We use the objective function given by the lexicographic function f on the transformation space, where λ is the most significant coordinate, followed by σ and then the d coordinates of the translation τ. The minimum point in the intersection of the constraints $H^a_{(\tau,\sigma,\lambda)}$ and $H^b_{(\tau,\sigma,\lambda)}$, with respect to f, represents a scaling and translation which minimizes the Hausdorff distance between A and B. $\square$

In the plane, the most important case for existing applications, we can implement this convex program so that

Theorem 4.2 *The scaling and translation which minimizes the Hausdorff distance between two polygons in the plane can be found in expected $O(n)$ time, where $n = |V(A)| + |V(B)|$.*

Proof: We redefine each polygon as the intersection of pieces of constant complexity, and associate each piece with a subset of the vertices of the other polygon. Each piece is the infinite wedge formed by a vertex of one of the polygons and the rays supporting the adjacent sides, which we shall call an *angle.* If every vertex of A is within the λ-offset surface of every angle from B, and visa versa, $H(A, B)$ is no greater than λ. This gives a GLP with $O(n^2)$ constraints, pairing every vertex of A with every angle from B, and visa versa.

We get a linear number of constraints by noting that for every angle α from B, all of A is within the λ-offset surface of α if every vertex from a *critical subset* of $V(A)$ is. The faces of B divide the circle of normal directions into a family I_B of closed intervals, each interval corresponding to an angle. A vertex v of A is critical for an angle α if v is extremal in A for any direction in the interval induced by α. The face normals of A also divide the circle of normal directions into a family I_A of closed intervals, each corresponding to the set of directions in which a particular vertex of A is extremal. When the interval corresponding to a vertex v intersects the interval corresponding to an angle α, v is extremal for α. The critical subsets for every angle can be found by merging I_A and I_B in linear time. Each of the n intersections of an interval in I_A with an interval in I_B gives a vertex-angle pair which produces a constraint. We also construct the n constraints induced by angles of A and vertices of B.

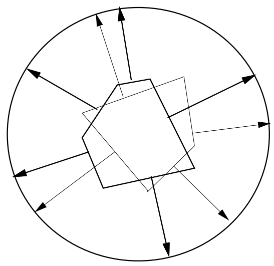

Figure 3: Intervals on the circle of normals

The lexicographic minimum point in the intersection of any four constraints can be found in constant time, so we get an expected $O(n)$ time algo-

rithm.
□

We now turn our attention to the problem of minimizing the one-directional Hausdorff distance $\vec{H}(P, A)$ from a set of P points to a convex polytope A. We can always scale A so that it is large enough to cover all the points, making $\vec{H}(P, A)$ zero. When the only transformation allowed is translation, however, we find

Theorem 4.3 *The translation that minimizes the one-directional Hausdorff distance from a set P of points to a convex polytope A in E^d can be found by a lexicographic convex program in E^{d+1}. For $d = 2$, we can solve the problem in expected $min\{O(mn), O(n \lg n + m)\}$ time, where $n = |P|$ and $m = |V(A)|$.*

Proof sketch: If the $O(mn)$ term is less than the $O(n \lg n + m)$ term (few vertices and many points), we construct a convex program in which each point in P produces n constraints of constant complexity, similar to those in the previous reduction. If not, we compute the convex hull of P and use a one-directional, translation-only modification of the previous construction.

Note that the Hausdorff distance used in this section can be derived from any metric on E^d, not just L^2.

5 A new proof of an interesting Helly-type theorem

In this section we will use an easy but powerful theorem from [A93], which gives us a simple technique for proving Helly-type theorems.

Theorem 5.1 *Let (H, w) be a GLP problem with combinatorial dimension d. The intersection of H is nonempty if and only if the intersection of every $G \subseteq H$ with $|G| \leq d + 1$ is nonempty.*

Let C_d^k be the family of all sets in R^d consisting of the disjoint union of at most k closed convex sets. A family I of sets is *intersectional* if, for every $H \subseteq I$, $\bigcap H \in I$. C_d^k is not intersectional. But consider some subfamily $I_d^k \subset C_d^k$ which *is* intersectional. This may "just happen" to be true, or I_k^d may be intersectional for some geometric reason. For example consider of any family J of sets like the one in figure 4 where each set is a pair of balls of diameter δ, separated by a distance of at least δ, kind of like dumbbells. The family formed by tak-

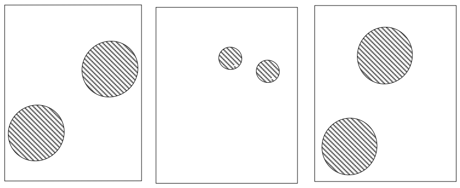

Figure 4: Generating family of I_d^k

ing the intersection of every subfamily of J forms an intersectional family I_d^k, since every intersection consists of at most two convex components.

Morris [Mo73] proved the following

Theorem 5.2 *Any intersectional family $I_d^k \subseteq C_d^k$ has Helly number $k(d + 1)$.*

We use Theorem 4.1 to give a new short and intuitive proof. Given any finite family $H \subseteq I_d^k$, we construct a GLP problem with H as the constraints. Let f be the lexicographic objective function on E^d, and define $w(G) = \min\{f(x) \mid x \in \bigcap G\}$, for all $G \subseteq H$. Since minima are identified with points, we will speak of $w(G)$ as a point and ignore the fine distinction between the point x itself and the value $f(x)$.

Theorem 5.3 *Finding $w(H)$ is a GLP problem of combinatorial dimension $k(d + 1) - 1$.*

Proof: It is easy to see that the problem satisfies Condition 1, since adding constraints to a subproblem can only increase the minimum. And is satisfies Condition 2 since every value of $w(G)$ is identified with a unique point, which is either in or out of h.

Recall that the combinatorial dimension is the largest cardinality of any basis B such that $\bigcap B$ is nonempty. We will count the constraints in any basis B by carefully removing selected constraints one by one, while building up a subfamily S of "sacred" constraints which may not be removed in later steps.

We will maintain two invariants. The first is that $w(B - h) < w(B)$ for all $h \in B - S$. The second invariant is that for all $h \in B - S$, the minimum point $w(B-h)$ lies in a different convex component of $\bigcap(B-h)$ from the point $w(B)$.

First we choose the subfamily S so that the invariants are true initially. Since $\bigcap B \neq \emptyset$, there is a minimum point $w(B)$ in some convex component of $\bigcap B$. Each $h \in B$ is the disjoint union of convex sets; for each h, the point $w(B)$ is contained in exactly one of them. Call this convex set C_h, and let $C = \{C_h \mid h \in B\}$. The pair (C, w) is a lexicographic convex programming problem, a GLP problem of combinatorial dimension d, with $w(C) = w(B)$. So C must contain a basis B_C with $|B_C| \leq d$. We set $S = \{h \in B \mid C_h \in B_C\}$.

How does this ensure the invariants? Since B is a basis, the first invariant holds for any subset S. The second invariant holds because all the constraints which contributed a convex component to B_C are in S, and for any $h \in B - S$, $w(B-h) < w(B) = w(B_C)$. That is, since the point $w(B)$ is the lowest point in $\bigcap B_C$, and $w(B-h)$ is lower than $w(B)$, the point $w(B-h)$ cannot be in $\bigcap B_C$, and hence must be in a different convex component of $\bigcap(B-h)$.

Now we turn our attention to selecting a constraint to remove from B. We use the fact that all the points $w(B-h)$ are distinct, for all $h \in B-S$. This is true because the point $w(B-h) \notin h$, so that for any other $h' \in B$, since $h \in (B-h')$, $w(B-h) \notin \bigcap(B-h')$. Since the $w(B-h)$ are distinct, there is some $h_{max} \in B-S$ such that $w(B-h_{max}) > w(B-h)$ for all other $h \in B-S$.

So consider removing h_{max} from B. Since $w(B-h) < w(B-h_{max})$, for any other $h \in B-S$, certainly $w(B-h-h_{max}) < w(B-h_{max})$. So the first invariant is maintained for $B-h_{max}$ and S. To re-establish the second invariant, we have to add more elements to S. We do this in the same way as before, by finding the at most d constraints which determine the minimum of the convex component containing $w(B-h_{max})$. We add these constraints to S, and set $B = B - h_{max}$.

We iterate this process, selecting constraints to remove from B and adding constraints to S, until $B-S$ is empty, that is, $B = S$. We now show that each removed constraint h accounts for at least one convex component $\mathcal{C}_{\langle}$ in $\bigcap S$. Removing h from B caused a new minimum point $w(B-h)$ to be created. This point was the minimum point in some convex component $\mathcal{C}_{\langle}$ of $\bigcap(B-h)$. We added the constraints determining $w(B-h)$ to S, so $w(B-h)$ has to remain the minimum point in its convex component, throughout the rest of the process. This ensures that, although $\mathcal{C}_{\langle}$ may later become part of some larger component, it will never become part of a larger component with a lower minimum point. Each subsequent component created by the removal of another constraint from B *will* in fact have a lower minimum point, so the component containing $w(B-h)$ must remain distinct from all later components. Thus every removed constraint h will account for at least one distinct component in $\bigcap S$.

Since I_k^d is an intersectional family, no subfamily of constraints can have more than k convex components in its intersection. Since $\bigcap B$ was initially nonempty, we started with at least one convex component, and at most d constraints in S. No more than $k-1$ constraints were removed, and each constraint removed added at most d constraints to S. So the total size of $|B| \leq (k-1)+|S| \leq k(d+1)-1$.

□

Theorems 4.1 and 4.3 together imply Theorem 4.2.

References

[A93] Nina Amenta, Helly Theorems and Generalized Linear Programming, *Proceedings of the 7th Annual Symposium on Computational Geometry*, (1993) pages 63-72.

[A93'] Nina Amenta, *Helly Theorems and Generalized Linear Programming* (PhD thesis), Geometry Center Technical Report, Number GCG61, (1993).

[ABB91] Helmut Alt, Bernd Behrends, and Johannes Blömer. Approximate matching of polygonal shapes, *Proceedings of the 7th Annual Symposium on Computational Geometry*, (1991), pages 186-93.

[AK89] James Arvo and David Kirk. A survey of ray tracing acceleration techniques, in *An Introduction to Ray Tracing* (Andrew Glassner, ed), Academic Press, San Diego, (1989), pages 201-262.

[A83] Mikhail J. Atallah. A linear time algorithm for the Hausdorff distance between convex polygons, *Information Processing Letters* **17** (1983), pages 207-9.

[CGHKKK94] L. Paul Chew, Michael T. Goodrich, Daniel P. Huttenlocher, Klara Kedem, Jon M. Kleinberg, and Dina Kravets. Geometric Pattern Matching under Euclidean Motion, *Proceedings of the Fifth Canadian Conference on Computational Geometry*, (1993) pages 151-156

[C90] Kenneth L. Clarkson. Las Vegas algorithms for linear and integer programming when the dimension is small, *manuscript*, 1990. An earlier version appeared in *Proceedings of the 29th Annual Symposium on Foundations of Computer Science*, pages 452-5, 1988.

[DMR93] Karen Daniels, Victor Milenkovic and Dan Roth. Finding the maximum area axis-parallel rectangle in a polygon, *Proceedings of the Fifth Canadian Conference on Computational Geometry*, (1993) pages 322-327.

[DGK63] Ludwig Danzer, Branko Grünbaum, and Victor Klee. Helly's Theorem and it's relatives, *Proceedings of the Symposium on Pure Mathematics*, Vol. 7, Convexity (1963) pages 101-180. American Mathematical Society, Providence, RI.

[E93] Jürgen Eckhoff. Helly, Radon and Carathody type theorems, Chapter 2.1 in *Handbook of Convex Geometry*, P.M. Gruber and J.M. Willis, eds., (1993) Elsevier Science Publishers B. V., Amsterdam.

[GM61] Branko Grünbaum and Theodore S. Motzkin. On components in some families of sets, *Proceedings of the American Mathematical Society*, vol. 12, (1961) pages 607-613.

[HKS91] Daniel P. Huttenlocher, Klara Kedam and Micha Sharir. The Upper Envelope of Voronoi Surfaces and its Applications, *Proceedings of the 7th Annual Symposium on Computational Geometry*, (1991) pages 194-203.

[L68] D.G. Larman. Helly type properties of unions of convex sets, *Mathematika* 15 (1968) pages 53-59.

[M94] Jiří Matoušek. On geometric optimization with few violated constraints, *These Proceedings*.

[MSW92] Jiří Matoušek, Micha Sharir and Emo Welzl. A subexponential bound for linear programming, *Proceedings of the 8th Annual Symposium on Computational Geometry* (1992) pages 1-8.

[Mo73] Howard Cary Morris. *Two Pigeon Hole Principles and Unions of Convexly Disjoint Sets*, PhD thesis, California Institute of Technology, Pasadena, California (1973).

[S90] Raimund Seidel. Linear programming and convex hulls made easy, *Proceedings of the 6th Annual Symposium on Computational Geometry*, (1990) pages 211-215.

[SW92] Micha Sharir and Emo Welzl. A combinatorial bound for linear programming and related problems, *Proceedings of the 9th Annual Symposium on Theoretical Aspects of Computer Science* (1992), Lecture Notes in Computer Science 577, pages 569-579.

[W91] Emo Welzl. *Smallest enclosing disks (balls and ellipsoids)*, Technical report number B 91-09 Fachbereich Mathematik, Freie Universität Berlin, Berlin, Germany (1991)