Helly-type Theorems and Generalized Linear Programming

Nina Amenta *

Abstract

Recent combinatorial algorithms for linear programming can also be applied to certain non-linear problems. We call these Generalized Linear Programming, or GLP, problems. We connect this class to a collection of results from combinatorial geometry called Helly-type theorems. We show that there is a Helly-type theorem about the constraint set of every GLP problem. Given a family H of sets with a Helly-type theorem, we give a paradigm for finding whether the intersection of H is empty, by formulating the question as a GLP problem. This leads to many applications, including linear expected time algorithms for finding line transversals and mini-max hyperplane fitting. Our applications include GLP problems with the surprising property that the constraints are non-convex or even disconnected.

1 Introduction

The fixed dimensional linear programming algorithms developed in computational geometry have a very combinatorial flavor. Both Clarkson [?] and Seidel [?] gave linear programming algorithms which, because they relied only on the combinatorial rather than the geometric properties of linear halfspaces, could be applied to various non-linear problems. The deterministic fixed dimensional algorithms of Dyer [?] and Megiddo [?] were generalized to specific non-linear problems albeit with considerable effort [?], [?]. In an important step, Sharir and Welzl [?] gave a randomized algorithm, and defined an abstract framework under which their algorithm could be applied. They gave a number of non-linear problems that fell into this framework. A re-analysis [?] of this algorithm, in response to the subexponential randomized simplex algorithm given by Kalai [?], showed that it too could solve arbitrary-dimensional linear programming in expected subexponential time, although this result does not immediately apply to any of the non-linear problems. The abstract framework of [?] applies to [?] as well, which, with an additional assumption, can be derandomized to give deterministic linear time algorithms for fixed dimensional problems [?]. We will call the class of problems described by the abstract framework *Generalized Linear Programming*, or GLP, problems.

We may begin to wonder about the relationship between GLP and the vast body of previous work in combinatorial geometry and mathematical programming. Does the framework merely restate some known characterization of "easy" non-linear problems? For instance, both Clarkson [?] and Seidel [?] noted that their algorithms can be applied to the *convex programming problem*: minimize a convex objective function over the intersection of a family of convex sets. Most of the known examples of GLP problems can be formulated as convex programs. Can they all?

*University of California, Berkeley, and The Geometry Center, Minneapolis, MN. The author was supported at Berkeley by a U.C. President's Dissertation Year Fellowship. Some of this work was done while visiting the Freie Universität, Berlin.

In this paper we forge a connection to one relevant area of combinatorial geometry, the class of results known as Helly-type theorems. Helly-type theorems have the same combinatorial structure as the following archetypical theorem:

Helly's Theorem *Let K be a family of at least $d+1$ convex sets in E^d, and assume K is finite or that every member of K is compact. If every $d+1$ members of K have a point in common, then there is a point common to all the members of K.*

This is one of the fundamental properties of convexity. There are many similar theorems with the same logical structure, for objects other than convex sets, for properties other than intersection, or for special cases in which $d+1$ is replaced by a different constant k. Combinatorial geometers collect Helly-type theorems, in much the same way that computer scientists collect NP-complete problems [?], [?], [?].

There is a natural computational problem associated with any Helly-type theorem. For instance, we might ask if Helly's theorem can help us get an algorithm that takes a family of convex sets and returns a point in their intersection, if one exists. An algorithm for this problem in [?] uses $O(n^{d+1})$ calls to a subroutine that finds a point in the intersection of $d+1$ convex sets. Finding the *minimum* point in the intersection, with respect to a convex objective function, is convex programming. For fixed dimension d, the GLP algorithms solve convex programming using only expected $O(n)$ calls to a slightly stronger subroutine that finds the minimum in the intersection of $d+1$ sets. Can we apply GLP in an analogous way to all Helly theorems? What should play the role of the objective function?

2 Definitions and Background

Let C be a family of objects, and $\mathcal{P}$ a predicate on subsets of C. A *Helly-type theorem* for C is something of the form:

There is a constant k such that for all finite $H \subseteq C$, $\mathcal{P}(H)$, if and only if, for every $B \subseteq H$ with $|B| \leq k$, $\mathcal{P}(B)$.

The constant k is called the *Helly number* of C with respect to the predicate $\mathcal{P}$. We are interested in a particular subclass of Helly-type theorems, in which the objects in C are sets, and $\mathcal{P}(G)$, $G \subseteq C$, is that the sets in G intersect in a common point. We write $\bigcap G$ for $\{x \in X \mid x \in h, \forall h \in G\}$, and we say that a family of sets *intersects* when $\bigcap G \neq \emptyset$, that is, when $\mathcal{P}(G)$ is true.

A *set system* is a pair (X, C), where X is a set and C is a family of subsets of X. We say (X, C) is a *Helly system* if there is a k such that C has Helly number k with respect to the intersection predicate $\mathcal{P}$. The natural computational problem associated with a Helly system is, given a subset $H \subseteq C$, return a point $x \in \bigcap H$, or show that $\bigcap H = \emptyset$.

Most Helly-type theorems can be re-stated in terms of the intersection predicate. For example, consider the theorem that family H of points in E^d is contained in a unit ball if and only if every $d+1$ points are contained in a unit ball. This is equivalent to saying that the family of unit balls centered at points of H intersects if and only if every subfamily of $d+1$ balls centered at points of H intersects.

We can state most Helly-type theorems as follows. There is a set X (E^d, in the example), a family $C \subseteq 2^X$ (the points of E^d), and a set Y (the set of all unit balls). The Helly-type theorem is given in terms of a predicate $\mathcal{P}$ on subsets of C (that the points are contained in a unit ball), where $\mathcal{P}$ is defined in terms of another predicate $\mathcal{Q}$ on pairs in $C \times Y$ (that the unit ball $y \in Y$ contains a point $p \in E^d$). For $H \subseteq C$, $\mathcal{P}(H)$ iff $\exists\ (y \in Y)\ \forall\ (h \in H)\mathcal{Q}(h, y)$. To state the theorem in terms of the intersection predicate, let $c_h = \{y \in Y \mid \mathcal{Q}(h, y)\}$, for $h \in H$ (the set of all unit balls containing point h, which, when represented by their centers, forms a unit ball around h), and $C' = \{c_h \mid h \in H\}$. Then (Y, C') is a Helly system.

Now we review the abstract framework for generalized linear programming from [?]. A *generalized linear programming* (or GLP) problem is a family H of *constraints* and an objective function w from subfamilies of H to a totally ordered set Λ. The pair (H, w) must obey the following conditions:

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

2. Locality: 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 $w(G + h) > w(G)$

The set Λ must contain a special maximal element Ω; for $G \subseteq H$, if $w(G) = \Omega$, we say G is *infeasible*; otherwise we call G *feasible*. A *basis* is a subfamily B such that for every $h \in B$, $w(B - h) < w(B)$. Notice that every $G \subseteq H$ contains a basis $B \subseteq G$ with $w(G) = w(B)$; B is called a *basis for* G. The *combinatorial dimension* of a GLP problem is the maximum size d of any feasible basis; an infeasible basis may have size $d + 1$.

A GLP algorithm takes a GLP problem (H, w) and returns a basis B for H. The randomized GLP algorithm of Matoušek, Sharir and Welzl [?] uses two primitive operations. A *basis computation* takes a family G of at most $d + 1$ constraints and finds a basis for G. A *violation test* takes a basis B and a constraint h, and returns true if B is a basis of $B + h$. Let t_v be the time required for a violation test and t_b be the time required for a basis computation. Their algorithm runs in expected time linear in n and subexponential in d, assuming that both t_v and t_b are polynomial in d.

A GLP problem is *fixed dimensional* if d is constant. With the dependence on d hidden in the big-Oh notation, the algorithm of [?] takes expected $O(t_v n + t_b \lg n)$ time. Thus to get an algorithm that runs in expected linear time, t_v must be $O(1)$, although we can afford to spend up to $O(n/\lg n)$ time per basis computation.

3 Statement of Results

Our starting point is the following simple theorem.

Theorem 3.1 *Let (H, w) be a GLP problem with combinatorial dimension k. H has the property $w(H) \leq \lambda$ if and only if every $B \subseteq H$ with $|B| \leq k + 1$ has the property $w(B) \leq \lambda$*

This means that there is a Helly-type theorem corresponding to the constraint set of every fixed dimensional GLP problem.

Our Main Theorem goes in the other direction; we give a paradigm for solving the natural computational problem associated with a Helly-type theorem as a fixed dimensional GLP problem. What this means is that we construct an objective function w for a set system (X, H) with Helly number k, so that (H, w) is a GLP problem. Intuitively, the objective function is defined by the following process. We shrink the sets in H until they fail to intersect, then gradually grow them back to their original size. The value of w on a subfamily $G \subseteq H$ is the time at which G first intersects. For a Helly-type theorem about geometric objects, we can usually accomplish this by scaling the objects themselves.

A problem (H, w) constructed using our paradigm meets the Monotonicity Condition and has combinatorial dimension k, but does not necessarily meet the Locality Condition. We observe, however, that when the the constraints in every subfamily $G \subseteq H$ first intersect in a single unique point, then (H, w) *does* meet the Locality Condition and is a fixed dimensional GLP problem. We call this the Unique Minimum Condition.

As applications of our paradigm, we show that the following are fixed dimensional GLP problems, giving expected linear time algorithms:

1. Finding a line transversal of translates of a convex (but possibly complicated) object in the plane,
2. Finding a line transversal in E^d for some special classes of objects,
3. Finding the closest hyperplane to a family of points under the *weighted* L^∞ *metric*, in which every coefficient of every point is equipped with a weight,
4. Finding the closest line to a family of points under the weighted L^∞ metric,
5. Finding the closest hyperplane to a family of points under a convex polytopal metric,
6. Finding, for a convex object C of constant complexity and a family K of convex objects of constant complexity, the largest homothet of C contained in the intersection of K, or the smallest homothet of C containing K, or the smallest homothet of C intersecting every member of K.

A second way to apply the paradigm is to assume that there is a single set $\Lambda \in H$ that grows with time. The value of w on a subfamily G is then the first time at which $G + \Lambda$ intersects. This second approach gives a GLP (H, w) of combinatorial dimension $k - 1$ rather than k, again assuming that the problem meets the Unique Minimum Condition. As an application, we give a fixed-dimensional GLP for finding a point in the intersection of a family of sets, each of which is the union of two convex sets, given that the intersection never has more than two connected components. Notice that this is an example of a GLP that is certainly *not* a convex programming problem; the feasible region fails not only to be convex, but also to be connected.

For linear programming, one way of ensuring the Unique Minimum Condition is using a lexicographic objective function. We consider lexicographic objective functions for general GLP problems. Although they are useful in specific applications, we find that they do not provide a general theoretical solution to the problem of establishing the Unique Minimum Condition, since there are GLP problems for which adopting a lexicographic objective function causes the combinatorial dimension to double.

In fact, we show that there is no completely general construction of an objective function for a Helly system, by giving the following example.

Theorem 3.2 *For all $n > 1$, there is a family H of $2n$ sets with Helly number two such that for any function w that assigns the value Ω to infeasible subfamilies, and such that (H, w) meets the Monotonicity and Locality Conditions, the combinatorial dimension of (H, w) is n.*

Besides the applications presented here, we expect this work will be useful in identifying new GLP problems. It is often difficult and sometimes impossible to reduce a GLP problem to linear or convex form; these results give an alternative approach to getting an efficient algorithm. It also makes it easier to implement programs for these problems, since a single implementation of a GLP algorithm can be equipped with specialized subroutines to solve any one of them.

4 Helly-type Theorems from GLP Problems

We begin with the proof that there is Helly-type theorem corresponding to the constraint family of every fixed dimensional GLP problem.

Theorem ?? *Let (H, w) be a GLP problem with combinatorial dimension k. H has the property $w(H) \leq \lambda$ if and only if every $B \subseteq H$ with $|B| \leq k+1$ has the property $w(B) \leq \lambda$.*

Proof: Let $w(H) \leq \lambda$. By the Monotonicity Condition, every $B \subseteq H$ must have $w(B) \leq w(H) \leq \lambda$. Going in the other direction, H must contain a basis B with $w(B) = w(H)$, with $|B| \leq k+1$. So if every subfamily B with $|B| \leq k+1$ has $w(B) \leq \lambda$, then $w(H) = w(B) \leq \lambda$.
□
It is interesting that this proof does not use the Locality Condition. Notice that when λ is the special symbol Ω, this means that every infeasible family of constraints contains an infeasible subfamily of size $\leq k+1$. We can use this idea to prove new Helly-type theorems.

Application 4.1 *Line transversal of boxes in E^d.*

A *line transversal* of a family of objects is a line that intersects every object. Let a *positive* line transversal be one directed into the positive orthant of E^d. An argument in [?] implies that finding a positive line transversal of a family of axis-aligned boxes in E^3 is a GLP problem with combinatorial dimension four (this problem was later reduced to linear programming [?]). This gives the following Helly-type theorem: there is a positive line transversal of a family of axis-aligned boxes in E^3 if and only if there is a positive line transversal for every subfamily B of boxes such that $|B| \leq 5$. Or, in terms of the intersection property, the system (X, C) has Helly number 5, where X is the set positive lines in E^3, and every $c \in C$ is the set of positive lines intersecting some axis-aligned box in the input family.

We can also show that certain problems are *not* GLP using Theorem 3.1. This is useful in the same way as a lower bound, in that it rules out a particular line of attack.

Application 4.2 *Line transversal of convex sets in E^3.*

In [?], the authors give an example of a linearly ordered family A of n disjoint convex compact sets in E^3 that has no line transversal, although every $(n-2)$-element subfamily of A has a line transversal consistent with the ordering. This tells us that we cannot hope to apply GLP in this situation, even given a linear ordering.

5 Paradigm for constructing GLP

Now we turn to the question of constructing a GLP problem to solve the natural computational problem concerning a Helly system (X, C): given a family $H \subseteq C$, find a point in $\bigcap H$ or show that none exists.

Recall that the range Λ of a GLP objective function w can be any totally ordered set. We define an objective function using a well-behaved function $w' : X \rightarrow \Lambda$ on the points $x \in X$, which we shall call a *ground-set objective function.*

Definition 5.1 *For a $G \subseteq C$, let $w(G) = \min\{w'(x) \mid x \in \bigcap G\}$, and $w(G) = \Omega$ when $\bigcap G = \emptyset$. We call $w : 2^H \rightarrow \Lambda$ the objective function* induced *by w' on (X, C).*

For example, when formulating linear programming in the GLP framework, the value of w on subfamily G of constraints is the minimum value that a linear ground-set objective function w' achieves on the points that are feasible with respect to G.

There is a problem with this definition, however, since there may be no minimum point even when $\bigcap G \neq \emptyset$, for instance when a linear program is unbounded below.

Definition 5.2 *An objective function w induced by a ground-set objective function w' is* well defined *when $w(G)$ is defined for every $G \subseteq H$.*

In linear programming, we get a well defined objective function by putting a "bounding box" around the problem, or symbolically compactifying the space by representing points at infinity.

Observation 5.1 *A GLP problem (H, w) satisfies the Monotonicity Condition whenever w is a well-defined objective function induced by a ground-set objective function w'.*

This follows from the fact that adding a constraint only eliminates feasible points, so the value of the minimum remaining feasible point can only go up. Certain functions w' also produce a function w that meets the Locality Condition.

Definition 5.3 *Let w be a well-defined objective function induced by a ground-set objective function w'. If $|\{x \in \bigcap G \mid w'(x) = w(G)\}| = 1$, for all $G \subseteq H$, then we say w' satisfies the* Unique Minimum Condition *on (X, H).*

This definition says that every subfamily not only has a minimum, but that this minimum is achieved by a unique point.

Observation 5.2 *If w' meets the Unique Minimum Condition on (X, H), then (H, w) satisfies the Locality Condition.*

This is because if $w(G) = w(F)$, for $F \subseteq G$, is achieved only at a single point x, then $w(G+h) > w(G)$ only if $x \notin h$, in which case $w(F+h) > w(F)$. There is one easy way to satisfy the Unique Minimum Condition.

Observation 5.3 *If $w'(x) \neq w'(y)$ for any two distinct points $x, y \in X$, then w' satisfies the Unique Minimum Condition on (X, H), for any $H \subseteq 2^X$.*

In general, however, the GLP we get when we use an arbitrary function w' to ensure the Unique Minimum Condition will *not* be fixed dimensional.

We now define an extension of a Helly system to a structure that admits an objective function in a natural way.

Definition 5.4 *A* nested family *is an indexed family of subsets* $\overline{h} = \{h_\lambda \mid \lambda \in \Lambda\}$*, such that* $h_\alpha \subseteq h_\beta$*, for all* $\alpha, \beta \in \Lambda$ *with* $\alpha < \beta$*.*

Let $\overline{H}$ be a family of nested families. We write $H_\lambda = \{h_\lambda \mid \overline{h} \in \overline{H}$ and $\lambda \in \Lambda\}$. For example, Λ might be $[0,1]$, and each $\overline{h} \in \overline{H}$ might be the set of balls of radius at most one centered at a particular point, indexed by the radius. Then H_λ would be a family of balls of radius exactly λ.

Definition 5.5 *A set system* $(X \times \Lambda, \overline{H})$ *is a* parameterized Helly system *with Helly number* k*, when*

1. *every* $\overline{h} \in \overline{H}$ *is a nested family, and*
2. (X, H_λ) *is a Helly system, with Helly number* k*,* $\forall \lambda$*.*

Figure 1 is a schematic diagram of a parameterized Helly system. The whole stack represents $X \times \Lambda$,

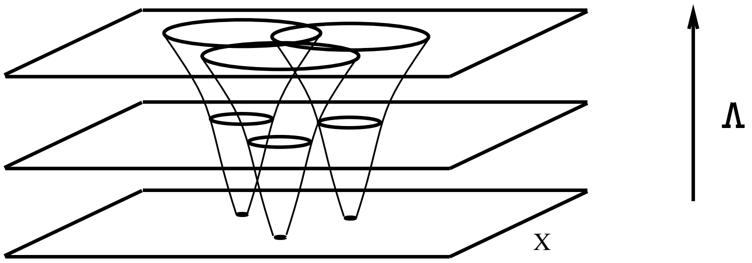

Figure 1: A parameterized Helly system

and each of the fluted things represents a set $\overline{h} \in \overline{H}$. Each $\overline{h}$ is a subset of $X \times \Lambda$. Because all the $\overline{h}$ are monotone with respect to Λ, the cross-section at λ (represented by one of the planes) is equivalent to the Helly system (X, H_λ). Notice that if $\overline{G} \subseteq \overline{H}$ does not intersect at some value λ_2, then $\overline{G}$ also fails to intersect at all $\lambda_1 < \lambda_2$, and if $\overline{G} \subseteq \overline{H}$ intersects at λ_1, then $\overline{G}$ also intersects at all $\lambda_2 > \lambda_1$.

A parameterized Helly system has a natural objective function w.

Definition 5.6 *For a point* $(x, \lambda) \in X \times \Lambda$*, let* $w'(x, \lambda) = \lambda$*. The objective function induced by* w' *is the* natural objective function *of a parameterized Helly system* $(\Lambda \times X, \overline{H})$*.*

So the function w' is just projection into the λ coordinate, and for $\overline{G} \subseteq \overline{H}$, $w(\overline{G}) = \lambda^* = min\{\lambda \mid \bigcap G_\lambda \neq \emptyset\}$, or $w(\overline{G}) = \Omega$ if $\overline{G}$ does not intersect at any value of λ.

Theorem 5.1 (Main Theorem) *Let* $(X \times \Lambda, \overline{H})$ *be a parameterized Helly system with Helly number* k *and natural objective function* w*. If* w *is well defined and meets the Unique Minimum Condition then* $(\overline{H}, w)$ *is a GLP problem of combinatorial dimension* k*.*

Proof: Observations 5.1 and 5.2 indicate that $(\overline{H}, w)$ meets the Monotonicity and Locality Conditions.

To prove that $(\overline{H}, w)$ has combinatorial dimension k, we have to show that the size of any basis is at most k. Consider any basis $\overline{B} \subseteq \overline{H}$. The definition of a basis says that for any $\overline{h} \in \overline{B}$, $w(\overline{B} - \overline{h}) < w(\overline{B})$. Let $\lambda^{max} = max\{w(\overline{B} - \overline{h}) \mid \overline{h} \in \overline{B}\}$. The basis $\overline{B}$ does not intersect at λ^{max}, but for any $\overline{h} \in \overline{B}$, $w(\overline{B} - \overline{h}) \leq \lambda^{max}$, which means that $\overline{B} - \overline{h}$ intersects at λ^{max}. Since $(X, H_{\lambda^{max}})$ has Helly number k, $\overline{B}$ must contain a subfamily $\overline{A}$ with $|\overline{A}| \leq k$, such that $\overline{A}$ does not intersect at λ^{max}. Every $\overline{h} \in \overline{B}$ must be in $\overline{A}$, since otherwise it would be the case that $\overline{A} \subseteq (\overline{B} - \overline{h})$ for some $\overline{h}$. This cannot be, because $\overline{A}$ does not intersect at λ^{max} while every $(\overline{B} - \overline{h})$ does. Therefore $\overline{B} = \overline{A}$ and $|\overline{B}| \leq k$.
□

Application 5.1 *Line transversal of translates in the plane.*

Let T be a family of disjoint translates of a single convex object O in E^2. Tverberg [?] showed that if every family $B \subseteq H$ with $|B| \leq 5$ admits a line transversal, then H also admits a line transversal. Egyed and Wenger [?] gave a deterministic linear time algorithm to find a line transversal. Showing that the problem can be formulated as GLP gives a simpler, although randomized, linear time algorithm.

We assume that the family of translates is in general position (we will define general position in a moment); if not, we use a standard perturbation argument. The set X is the set of lines in the plane, and we let $\Lambda = [0, 1]$. We abuse notation so that t refers both to a translate $t \in T$ and to the set of lines intersecting t. So a subfamily $G \subseteq T$ intersects when there is a line that intersects every translate in G. We pick a distinguished point q in the interior of the archetypical object O. For a particular translate t, let λt be the homothet of O that results from scaling translate t by a factor of λ, keeping the point corresponding to q fixed in the plane. Let the family $\overline{t} = \{\lambda t \mid 0 \leq \lambda \leq 1\}$, and $\overline{T} = \{\overline{t} \mid t \in T\}$.

Every line that intersects the homothet $\lambda_1 t$ also intersects $\lambda_2 t$ for any $\lambda_2 > \lambda_1$. So each $\overline{t}$ is a nested family of lines. The t_λ are always disjoint, every $(X, \lambda T)$ is a Helly system with Helly number 5, and $(X \times \Lambda, \overline{T})$ is a parameterized Helly system.

The natural objective function $w(\overline{G})$, where $\overline{G} \subseteq \overline{T}$, is the minimum λ such that G_λ intersects. In the case where $\overline{G}$ consists of a single translate, we define $w(\overline{G}) = 0$. Notice that for certain degenerate placements of the translates (see figure ??) it is possible for there to be two or even three distinct line transversals at $\lambda^* = w(\overline{G})$.

The general position assumption is that the line transversal at λ^* is always unique.

$(\overline{T}, w)$ is a GLP problem with combinatorial dimension 5. Either the GLP algorithm finds a line transversal at some value of $\lambda \leq 1$, or no line transversal of the input exists.

When O is a polygon with a constant number of sides, neither this algorithm nor Egyed and Wenger's is very interesting, since we can find a line transversal via a constant number of fixed dimensional linear programming problems. Both algorithms are useful only for more complicated polygons, in which the number of sides depends on n, or for non-polygonal objects.

Recall that the algorithm in [?] runs in $O(t_v n + t_b \lg n)$ time, where t_v is the time required for a violation test and t_b is the time required for a basis computation. In this application, a violation test

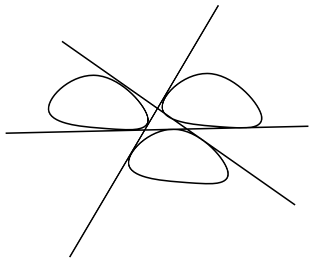

Figure 2: Degenerate input

determines whether the current minimum line x intersects a new homothet λt. For any x, there is a diametral pair of points on O such that x intersects a homothet λt if and only if x passes between the corresponding points on λt. We require the basis computation subroutine to find such a diameteral pair whenever it finds a new minimum line x, so t_v is $O(1)$. The running time is then limited by t_b; when the complexity of O is such that t_b is $O(n/\lg n)$, we get an expected linear time algorithm.

Notice that here the dimension of the space X of lines in the plane is two. If there were an affine structure on X such that the constraints T were convex subsets of X, then the Helly number of the system (X, T) would be three. But examples show that the bound of five is in fact tight, which means that this is a GLP that is *not* a convex program. This is also a natural example of a GLP problem in which the minimal object does not "touch" every constraint in the basis.

This paradigm may be profitably applied to many other Helly-type theorems.

Application 5.2 *Homothets spanning convex sets.*

We use Theorem 2.1 from [?],

Theorem (Vincensini and Klee) *Let K be a finite family of at least $d+1$ convex sets in E^d, and let C be a convex set in R^d. Then there is a translate of C that [intersects/ is contained in/contains] all members of K if and only if there is such a translate for every every $d+1$ members of K.*

We apply the paradigm by either growing or shrinking the convex body C, to get an algorithm that takes as input a finite family K of at least $d+1$ convex sets in E^d and a convex set C and returns either the smallest homothet of C that contains $\bigcup K$, the largest homothet of C contained in $\bigcap K$, or the smallest homothet of C that intersects every member of K. These problems can be seen a generalization of Megiddo's problem of finding the smallest ball containing a set of balls [?]. All of them are convex programming problems. The combinatorial dimension in each case is $d+1$, and the running time again depends on the complexity of the objects. When C and all the elements of K are of constant complexity we get an expected linear time algorithm. In other cases, preprocessing can often be used to reduce the obvious running times.

For instance, there is an $O(k + n\lg^2 n)$ algorithm in [?] for finding the largest homothet of a k-vertex convex polygon P inside an n-vertex convex polygon Q in E^2. The problem arises as a restriction of a pattern matching problem; we can think of Q as input from a low-level vision system and P as a model of an object we'd like to match. This is a restriction of the problem because we allow P to scale and translate, but not to rotate.

In this simple case, the GLP formulation, after preprocessing, is just linear programming. We use the halfspaces supporting Q as constraints. Let us identify translates of P with translates of an arbitrary distinguished point $p \in P$. For a fixed scale factor λ, the possible translates that put P on the correct side of a halfspace h put p in a halfspace h' parallel to h and offset by the distance from p to some vertex v of P. We can assign the correct vertex v to each halfspace h in $O(n+k)$ time by merging the lists of face normals for P and Q. As we increase the scale factor λ, the distance from p to v changes linearly, and h' traces out a three-dimensional halfspace in $\lambda \times E^2$. The linear program maximizes λ over the intersection of the h' in $O(n)$ time, so the total time is $O(n+k)$.

Application 5.3 *Special cases of line transversals in E^d.*

In general, finding line transversals is significantly more difficult in dimension $d > 2$ than it is in the plane, but there are a few special cases in which Helly-type theorems help us get a linear time algorithm. Theorems 5.6 and 5.7 in [?], due to Grünbaum, concern, respectively, a family of $d-1$ dimensional convex sets, all of which lie in a family of parallel hyperplanes, and a family of closed balls such that the distance between any two is greater than the sum of their diameters. In both these cases, if there is a line through every $2d-1$ objects then there is a line through all of them. Again, the paradigm can be applied to give a linear time algorithm to find a line transversal. When the parallel convex sets are polytopes, we can formulate this GLP as a linear program as well. In the case of the closed balls, the Helly-type theorem breaks down when the balls get too close together. Therefore the constraints (a constraint is the set of lines that intersects a ball) again cannot be convex under any affine structure we may impose on the space of lines, and this is another example of a non-convex GLP problem.

6 Optimization problems

Notice that, for example, the line transversal algorithm for translates in Application 5.1 finds a line that minimizes the maximum distance from the family of fixed points, under the quasi-metric whose unit ball is the object O. This mini-max property is useful in and of itself. We can use the Main Theorem with known GLP problems much as we would apply parametric search, to find a minimum parameter value at which the problem is feasible. While parametric search usually adds an additional logarithmic factor to the running time, here the expected time remains linear.

Application 6.1 *Weighted L^∞ linear interpolation.*

The input to this problem is a family of n points in R^d, with an axis-aligned rectangle T_p centered at each point p. Note that each T_p may have different dimensions. We define the distance from a hyperplane y to p to be the smallest nonnegative real value λ such that T_p intersects y when scaled by λ around p. We call this the *weighted L^∞ metric*. The linear interpolation problem is to find the hyperplane y that minimizes the maximum distance to any point.

This problem arises when each coordinate of each point is given a weight, producing box-shaped error regions. This occurs, for example, when the point coordinates are calculated and error is bounded using interval arithmetic, or when complicated error regions are approximated by bounding boxes. The general-dimensional version of the problem was considered in [?], [?], and in [?], where it was shown to be NP-hard. Showing that the problem can be formulated as GLP gives a expected linear time algorithm for the fixed dimensional case.

Define a *positive* hyperplane to be one that is oriented so that its normal vector is directed into the positive orthant of E^d. There is a diameteral pair of vertices v^+, v^- on each box such that, at any fixed value of λ, a positive hyperplane y intersects the box if and only if v^+ and v^- lie in its positive and negative halfspaces, respectively. Finding a positive hyperplane transversal of the boxes at a fixed value of λ is thus the geometric dual form of a d-dimensional linear program with $2n$ constraints. If a fixed positive hyperplane y goes through a box at λ_1, it also does so at any $\lambda_2 > \lambda_1$. So for each constraint point v in the linear program, the set of hyperplanes for which v lies in the correct halfspace form a nested family parameterized by λ. The natural objective function returns the the minimum value λ^* at which a positive hyperplane transversal of the boxes exists.

For degenerate input, the boxes might admit a transversal by a lower-dimensional affine subspace at λ^*, and there would be an infinite set of positive hyperplane transversals. The simplest example is when three rectangles in E^2 meet in a point. We again use a perturbation argument to ensure that the boxes are in general position.

We conclude that finding the closest positive hyperplane in the weighted L^∞ metric is a GLP problem of combinatorial dimension $d+1$. Even when the boxes all have the same dimensions, the problem cannot be formulated as a linear program. For a given family of boxes, we define a separate GLP problem for each orthant of E^d. The solution to the whole problem will be hyperplane that achieves the minimum λ of any of the 2^d GLP problems.

Application 6.2 *Linear interpolation with a polytopal metric.*

Consider the problem of finding a hyperplane transversal of a family of polytopes whose facets are drawn from a set U that is the union of a constant number of families of parallel hyperplanes. In [?] there is a similar reduction of this problem to a fixed number of linear programming problems. Applying our paradigm gives an expected linear time algorithm for fitting a hyperplane to a point family under any quasi-metric whose unit ball is a polytope with a constant number of sides, or, more generally, in which each point has an error metric whose unit ball is a polytope with facets drawn from U.

Application 6.3 *Line fitting in the weighted L^∞ metric.*

Megiddo has shown [?] that the problem of finding a line transversal for a family of axis-aligned boxes in E^d can be formulated as a collection of linear programs in dimension $2d-2$. We can again apply the paradigm to find the closest line to a family of points under the weighted L^∞ metric defined above. This technique can be extended to take two input sets of points and weights $(P_1, T_1), (P_2, T_2)$, and find a pair of parallel lines l_1, l_2 such that the maximum of the weighted L^∞

distance from (P_1, T_1) to l_1 and from (P_2, T_2) to l_2 is minimized.

7 Nesting by intersection

To use our paradigm, we had to assume that each constraint could be parameterized to form a nested family. So far we have been able to do this in a natural way by shrinking or growing geometric objects, but it is not clear that this technique is applicable to all problems. In this section, we present another way of parameterizing a family of constraints, essentially due to Hoffman [?]. He assumes the existence of one nested family of constraints and uses it to build an objective function. In this way he relates the Helly number k to something he calls the *binding constraint number*; in our terms, he shows that the size of every basis is no greater than $k-1$. Convex programming is an example of a GLP that can be defined in this way.

Theorem 7.1 *Let (X, C) be a Helly system with Helly number k, with $\Lambda \subset C$ such that Λ is a nested family and $\emptyset, X \in \Lambda$. (X, C) can be extended to a parameterized Helly system $(X \times \Lambda, \overline{C})$ with natural objective function w. When, on a finite subfamily $\overline{H} \subseteq \overline{C}$, w is well defined and meets the Unique Minimum Condition, then $(\overline{H}, w)$ is a GLP problem of combinatorial dimension $k-1$.*

Proof: The range of our objective function will be the nested family Λ itself. So $\lambda \in \Lambda$ is a subset of X, rather than, for instance, a real number.

We extend every element $h \in C$ of our original family of constraints to a nested family $\overline{h}$ as follows. For $\lambda \in \Lambda$, we let $h_\lambda = h \cap \lambda$. Again we write $\overline{C} = \{\overline{h} \mid h \in C\}$, and $C_\lambda = \{h_\lambda \mid h \in C\}$.

The natural objective function w and the corresponding ground-set objective function w' are defined as usual for a parameterized Helly system. Assuming that w' meets the Unique Minimum Condition on $(X \times \Lambda, \overline{H})$ the Main Theorem tells us that $(\overline{H}, w)$ is a GLP problem of combinatorial dimension no greater than k. It remains therefore to be shown that the combinatorial dimension is in fact no greater than $k-1$, that is, that any feasible basis $\overline{B} \subseteq \overline{H}$ has $|\overline{B}| \leq k-1$.

We use another version of the argument in the proof of the Main Theorem. Let $\lambda^* = w(\overline{B})$. There is an element $\overline{h}^{max} \in \overline{B}$, such that $\lambda^{max} = w(\overline{B} - \overline{h}^{max}) \geq w(\overline{B} - \overline{h})$ for all $\overline{h} \in \overline{B}$. Remember that $\lambda^{max} \in C$.

Now let $B = \{h \in C \mid \overline{h} \in \overline{B}\}$. $A = B + \lambda^{max}$ fails to intersect, because $B_{\lambda^{max}}$ fails to intersect. Since $\overline{B}$ is a basis, $B - h + \lambda^{max}$ does intersect, for every $h \in B$, and we know that B intersects because B_{λ^*} intersects.

So all of A's proper subsets do intersect, while A does not. Since $A \subseteq C$ and (X, C) is a Helly system with Helly number k, $|A| \leq k$. This means $|B| = |\overline{B}| \leq k-1$.
□

The tricky part here is using the nested family Λ itself as the range of the objective function. We have to do it this way because Λ may have any order type; it may not be isomorphic to some convenient totally ordered set like the real line.

Recall that the Main Theorem used the fact that the Helly number of each (X, H_λ) was k to show that the combinatorial dimension was bounded by k. Here, although the combinatorial dimension is at most $k-1$, the Helly number of each (X, H_λ) can still be as great as k. For example, in the figure below, the constraints and λ are all convex. The combinatorial dimension of the problem is two, while the Helly number is three.

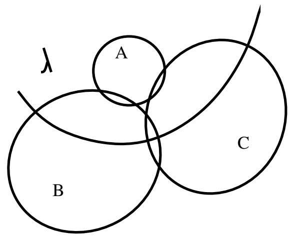

Figure 3: Degenerate input

Of course, we have no guarantee that a Helly system contains a nested family Λ that we can use to construct a function w that satisfies the Unique Minimum Assumption. We can, however, find such sets in specific applications. As an example, we use nesting by intersection to define a GLP problem in which the constraints are not even connected, let alone convex. This is interesting in that it implies we can sometimes use a single GLP to solve mathematical programming problems in which the feasible region becomes disconnected, so long as the number of connected components remains fixed.

Application 7.1 *Pairs of convex sets.*

Let C be the set of all convex sets in E^d, defined so as to include the empty set. Let Z be the family of all sets consisting of at most two disjoint convex sets, $Z = \{(C_1 \cup C_2) \mid C_1, C_2 \in C \text{ and } C_1 \cap C_2 = \emptyset\}$. A sets formed by taking the intersection of a subfamily of Z does not necessarily belong to Z. But consider a subfamily $Z' \subseteq Z$ with the special property that $\bigcap G \in Z$ for any $G \subseteq Z'$. (E^d, Z') is a Helly system with Helly number $2d+2$ [?]. As a concrete example, let each $z \in Z'$ be a pair of closed balls of radius one, separated by a distance of 1 (kind of like dumbbells).

Notice that we can adjoin the set Λ of closed balls centered at the origin to any family of pairs Z', so that $Z' \cup \Lambda$ still has the special property. Each ball can be considered a pair with the empty set, and the intersection of a ball with any pair of convex sets will produce ≤ 2 convex sets. Let S_r be the ball of radius r centered at the origin. For any $h \in Z'$, let $\overline{h} = \{h_r \mid h_r = h \cap S_r\}$. So for any $G \subseteq Z'$, the natural objective function $w(G)$ is the smallest r such that $\bigcap G \cap S_r \neq \emptyset$. Notice that for any G, if $r = w(G)$, $|\bigcap G \cap S_r| \in \{1, 2\}$. The appropriate general position assumption, which we can simulate by perturbation one member of each pair of sets, is that $|\bigcap G \cap S_r| = 1$. So (Z', w) is a GLP problem, with combinatorial dimension $2d+1$.

8 Lexicographic objective functions

So far our theoretical arguments have just assumed the Unique Minimum Condition. Even for linear programming, this requires a perturbation argument. The objective function

$$minimize\ x_0$$

fails to ensure a unique minimum, since the minimal value of x_0 might be achieved over a face of any dimension in the feasible polytope. One conceptual approach to this problem is to perturb each of the constraint hyperplanes, but this is messy in practice. Another approach is to perturb the objective function, or equivalently rotate the polytope slightly so that no facet is parallel to the hyperplane $x_0 = 0$. This can also be messy in practice. Instead, we usually use a *lexicographic* objective function

$$minimize\ \langle x_0, x_1, \ldots, x_{d-1}\rangle$$

where the minimum point is the one that minimizes x_{d-1}, over all that minimize x_{d-2}, etc. On any finite instance of LP, this linear lexicographic function is equivalent to the perturbed linear function

$$(1) \qquad x_0 + \epsilon x_1 + \epsilon^2 x_2 + \ldots + \epsilon^{d-1} x_{d-1}$$

for ϵ infinitesimally small.

In this section we use this idea to build a GLP objective function for other Helly systems. It is tempting to hope that we can construct a lexicographic objective function that ensures the Unique Minimum Condition for any Helly problem, but in fact we find that lexicographic objective functions interact strangely with the combinatorial dimension.

Here is one way of looking a problem with a lexicographic objective function. Say v_p is a lexicographic objective function whose range has p parameters. The most significant parameter is determined by a simple one-parameter objective function w_p, so that (H, w_p) is a problem that meets the Monotonicity Condition but not necessarily the Locality Condition. At each value λ of w_p, H_λ is the constraint set of another problem whose lexicographic objective function v_{p-1} has $p-1$ parameters. By analogy with linear programming, if every subproblem (H_λ, v_{p-1}) has combinatorial dimension d, we expect the whole problem (H, v_p) to have combinatorial dimension $d+1$. This is in fact the case when v_p is equivalent to a perturbation of a single parameter objective function. But this equivalence relies on geometric arguments and is not inherent in the GLP framework. In fact, when every (H_λ, v_{p-1}) has combinatorial dimension at most d, the combinatorial dimension of (H, v_p) might be as great as $d + k$, where k is combinatorial dimension of (H, w_p).

Despite this theoretical oddity, in most applications of interest, lexicographic objective functions *are* equivalent to perturbations of single parameter objective functions, and so can be used without increasing the combinatorial dimension. In particular, lexicographic objective functions can be used to remove the general position assumption for the the optimization problems of Section 6 and the pairs of convex sets in Application 7.1.

Let $(X \times \Lambda, \overline{H})$ be a parameterized Helly system with Helly number k and natural objective function w. For all $\lambda \in \Lambda$, we assume a function $v_\lambda : 2^{H_\lambda} \rightarrow \Lambda'$, where Λ' is a totally ordered set containing a maximum element Ω, such that (H_λ, v_λ) is a GLP problem of combinatorial dimension at most d. The functions v_λ may themselves be lexicographic. We define a lexicographic objective function v in terms of w and the functions v_λ. Let $v : 2^{\overline{H}} \rightarrow \Lambda \times \Lambda'$ be $v(\overline{G}) = (\lambda^*, \kappa^*)$, where $\overline{G} \subseteq \overline{H}$, $\lambda^* = w(\overline{G})$ and $\kappa^* = v_{\lambda^*}(G_{\lambda^*})$. We impose a lexicographic order on $\Lambda \times \Lambda'$, with $(\lambda, \kappa) > (\lambda', \kappa')$ if $\lambda > \lambda'$, or if $\lambda = \lambda'$ and $\kappa > \kappa'$.

Theorem 8.1 *If $(X \times \Lambda, \overline{H})$ is a parameterized Helly system with Helly number k and natural objective function w such that*
1. *w is well defined, and*
2. *for all λ, there is a function $v_\lambda : 2^{H_\lambda} \rightarrow \Lambda'$, where Λ' is a totally ordered set, such that (H_λ, v_λ) is a GLP problem of combinatorial dimension d,*

then there is a function v such that $(\overline{H}, v)$ is a GLP problem of combinatorial dimension $\leq k + d$.

Proof: Let v be defined as above. We know that v is well defined because w and all of the v_λ are well defined. As usual, Observation 5.1 tells us that $(\overline{H}, v)$ obeys the Monotonicity Condition.

Now consider $\overline{F} \subseteq \overline{G} \subseteq \overline{H}$, with $v(\overline{G}) = v(\overline{F}) = (\lambda^*, \kappa^*)$ and $v(\overline{G} + \overline{h}) = (\lambda, \kappa) > v(\overline{G})$. Either $\lambda > \lambda^*$, or $\kappa > \kappa^*$. If $\lambda > \lambda^*$, then $v_{\lambda^*}(G_{\lambda^*} + h_{\lambda^*}) = \Omega$, by the definition of w, and $v_{\lambda^*}(G_{\lambda^*} + h_{\lambda^*}) > v_{\lambda^*}(G_{\lambda^*})$. Otherwise, $\kappa > \kappa^*$, that is, $v_{\lambda^*}(G_{\lambda^*} + h_{\lambda^*}) > v_{\lambda^*}(G_{\lambda^*})$. In either case, by the Locality Condition on v_{λ^*}, $v_{\lambda^*}(F_{\lambda^*} + h_{\lambda^*}) > v_{\lambda^*}(F_{\lambda^*})$, and $v(\overline{F} + \overline{h}) > v(\overline{F})$. So the lexicographic function v also satisfies the Locality Condition.

Finally we consider the combinatorial dimension. Let $\overline{B}$ be a basis for any $\overline{G} \subseteq \overline{H}$. Then $v(\overline{B} - \overline{h}) = (\lambda, \kappa) < v(\overline{B}) = (\lambda^*, k^*)$, for any $\overline{h} \in \overline{B}$. Let the subset $\overline{B}_1 = \{\overline{h} \in \overline{B} \mid v(\overline{B} - \overline{h}) = (\lambda, \kappa)$ and $\lambda < \lambda^*\}$. Since the combinatorial dimension of $(\overline{B}_{\lambda^*}, v_{\lambda^*})$ is d, $\overline{B} - \overline{B}_1$ contains at most d constraints.

As we did in the proof of the Main Theorem, let Let $\lambda^{max} = max\{\lambda \mid v(\overline{B} - \overline{h}) = (\lambda, \kappa), \overline{h} \in \overline{B}_1\}$. Again, $\overline{B}$ fails to intersect at λ^{max} and hence must contain a set $\overline{A}$ of size $\leq k$ that also fails to intersect. And again, every $\overline{h} \in \overline{B}_1$ must also be in $\overline{A}$, since $\overline{B} - \overline{h}$ intersects at λ^{max} and $\overline{A}$ does not, so $\overline{A} \not\subseteq \overline{B} - \overline{h}$. So $|\overline{B}_1| \leq |\overline{A}| \leq k$, and $|\overline{B}| \leq k + d$.
□

Notice that since the parameterized system $(X \times \Lambda, \overline{H})$ has Helly number k, so does the Helly system (X, H_{λ^*}), and Theorem 3.1 tells us that $k <= d + 1$. So the combinatorial dimension of a problem for which every (H_λ, v_λ) has combinatorial dimension d is at most $2d + 1$.

Certainly this bound on the combinatorial dimension is not always tight. For d-dimensional linear programming, for instance, this theorem gives an upper bound of $2d - 1$ on the combinatorial dimension, since each H_λ is the constraint set of a $(d-1)$-dimensional linear program, and $(E^d, \overline{H})$ is a parameterized Helly system with Helly number d. None the less, the theorem gives the best general bound, as shown by the following example.

Theorem 8.2 *There is a parameterized Helly system $(X \times \Lambda, \overline{H})$ with Helly number $d + 1$, and a family of functions v_λ such that every (G_λ, v_λ) is a GLP problem with combinatorial dimension d, such that (H, v) has combinatorial dimension $2d + 1$, where v is the lexicographic objective function.*

Proof: We define a problem in which every H_λ is the constraint set of a d-dimensional linear programming problem. Every nested family $\overline{h} \in \overline{H}$ is a subset of $E^d \times \mathcal{R}$ of the form $a \cdot x \geq b_\lambda$, where a is a vector, $\cdot$ is dot-product, $x \in X$ and b_λ is a scalar quantity that varies with λ as follows.

Let b and c be constants. For $\lambda < c$, $b_\lambda = b$, and for $\lambda \geq c$, $b_\lambda = -\infty$. Essentially, a constraint remains in force below c, and recedes to $-\infty$ above. The function v_λ can be any linear function. Notice that $(E^d \times \mathcal{R}, \overline{H})$ is a parameterized Helly system with Helly number $d+1$, and that every (G_λ, v_λ) has combinatorial dimension d.

We can easily construct an instance of this problem for which a basis $\overline{B} \subseteq \overline{H}$ has size $|\overline{B}| = 2d+1$. Figure 2 is an example for the case $d = 1$.

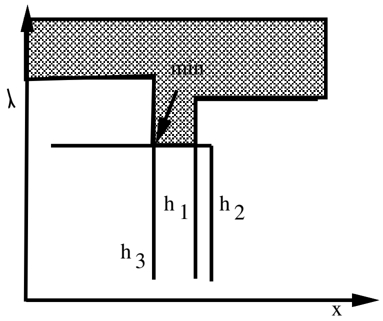

Figure 4: Size of basis is three

The constraints h_1 and h_2 determine the minimum value of λ, while constraint h_3 determines the minimum value of x.

To construct an example in general dimension, let u be any linear objective function on E^d. Construct a tiny simplex by intersecting a family A' of positive halfspaces with $|A'| = d+1$. Let A be the corresponding family of negative halfspaces; A is the constraint set of an infeasible d-dimensional linear program, with A a basis for (A, u). Now let B be the basis of a feasible d-dimensional linear programming problem, with the same objective function u, such that $|B| = d$ and the simplex s is strictly contained in the interior of $\bigcap B$. We parameterize A and B by λ, as above, assigning some arbitrary value c_a to every $\overline{h} \in \overline{A}$, and some other value $c_b > c_a$ to every $\overline{h} \in \overline{B}$. Let the constraint set of the whole problem be $\overline{H} = \overline{A} \cup \overline{B}$. For every $\lambda \in S$, let the function v_λ be u. Then $v(\overline{H}) = (\lambda^*, \kappa^*)$, with $\lambda^* = c_a$, and $\kappa^* = u(B)$. For any $\overline{h} \in \overline{A}$, $w(\overline{H} - \overline{h}) = -\infty$, and for any $\overline{h} \in \overline{B}$, $v_{\lambda^*}(H_{l^*} - h_{l^*}) < k^*$. So $\overline{H}$ is a basis for $\overline{H}$ of size $2d+1$.
□

Finally, we formalize the statement that if there is a single-parameter objective function w that is equivalent to a lexicographic objective function v, then the combinatorial dimension of (H, v) is no greater than that of (H, w).

Observation 8.1 *Let (H, v) be a GLP problem, with v a lexicographic objective function. If there is a function w such that (H, w) is a GLP problem of combinatorial dimension k, and $v(F) = v(G)$ implies $w(F) = w(G)$, for $F \subseteq G \subseteq H$, then the combinatorial dimension of (H, v) is no greater than k.*

Notice that we are not required to actually produce such a function w; the fact that it exists ensures that the combinatorial dimension is bounded by k even when using the lexicographic function v. So, for example, given an instance of d-dimensional linear programming, we don't actually have to find an ϵ small enough to define the perturbed linear function in Expression ??

9 Helly systems with unbounded combinatorial dimension

We have given a paradigm for constructing a GLP objective function for a Helly system, given the Unique Minimum Assumption. We now show that some such additional assumption is necessary, by exhibiting a set system with a fixed Helly number that *cannot* be turned into a fixed dimensional GLP problem.

Theorem ?? *For all $n > 1$, there is a family H of $2n$ sets with Helly number two such that for any function w that assigns the value Ω to infeasible subfamilies, and such that (H, w) meets the Monotonicity and Locality Conditions, the combinatorial dimension of (H, w) is n.*

Proof: Let the universe X consist of the 2^n points at the vertices of an n dimensional hypercube, and let the constraint family H be the $2n$ subsets each of which lies in a facet of the hypercube. Notice that if a subfamily $G \subseteq H$ includes any pair of opposite facets, then G fails to intersect, and otherwise G does intersect. So the Helly number of (X, H) is two.

Any objective function w must assign $w(G) = \Omega$ to the infeasible families G that contain a pair of opposite facets. Meanwhile any feasible G that does not contain a pair of opposite facets will have $w(G) = \lambda \in \Lambda$, with $\lambda < \Omega$. Let $\lambda^{max} = \max\{\lambda \in \Lambda | \lambda < \Omega \text{ and } \lambda = w(G) \text{ for some } G \subseteq H\}$, and consider a subfamily G with $w(G) = \lambda^{max}$. If $|G| < n$, then there exists a pair (h^+, h^-) of facets, such that G contains neither h^+ nor h^-. This means that $G + h^+$ is also feasible. By the Monotonicity Condition, $w(G + h^+) \geq w(G)$; and since $w(G)$ is maximal, we can conclude that $w(G + h^+) = w(G) = \lambda^{max}$. This argument shows that there must be a subfamily G of size n with $w(G) = \lambda^{max}$.

Now we show that there is no basis B for such a subfamily G such that $B \neq G$. Assume, for the sake of contradiction, and without loss of generality, that there is an element $h^+ \in G$ such that $h^+ \notin B$. Then $B + h^-$ is still feasible, so $w(B + h^-) = w(B) = \lambda^{max}$. But $w(G + h^-) = \Omega$. Since $B \subset G$ and $w(B) = w(G)$, this means that w fails to satisfy the Locality Condition, contrary to the assumption of the theorem. So w must be a function which makes $B = G$, and (H, w) must have combinatorial dimension n.
$\square$

This theorem says that the class of problems whose constraint sets have a fixed Helly number with respect to the intersection property is strictly greater than the class of fixed dimensional GLP problems.

10 Concluding Remarks

Our parameterized Helly system paradigm should be useful in producing computational versions of other interesting Helly theorems such as those using spherical convexity and those concerning separating surfaces.

There are other similar collections of theorems, such as Gallai-type theorems and Hadwiger-type theorems, and Helly-type theorems in which the fact that all subfamilies of size at most k have some property $\mathcal{P}$ implies that the whole family has some other property $\mathcal{Q}$. It would be interesting to find algorithmic applications of these.

The major open problem is to characterize the Helly systems (X, H) for which there is an objective function w that gives a fixed dimensional GLP problem (H, w).

11 Acknowledgments

The author thanks Raimund Seidel for suggesting the problem in Section 9, Raimund Seidel, Günter Rote and an anonymous referee for many useful comments on earlier versions of this paper, and Emo Welzl and Alon Efrat for pointing out references.

References

[A92] Nina Amenta. Finding a line transversal of axial objects in three dimensions, *Proceeding of the 3rd Annual ACM-SIAM Symposium on Discrete Algorithms*, (1992) pages 66-71.

[AD92] David Avis and Mike Doskas. Algorithms for high dimensional stabbing problems, *Discrete Applied Mathematics*, vol. 27, (1990) pages 39-48.

[AH91] David Avis and Micheal E. Houle. Computational aspects of Helly's theorem and its relatives, *Proceedings of the Third Canadian Conference on Computational Geometry*, (1991) pages 11-14.

[AGPW] Boris Aronov, Jacob E. Goodman, Richard Pollack and Rephael Wenger. There is no Hadwiger number for line transversals in higher dimensions. *unpublished manuscript*, cited in [?], theorem 2.9.

[CM93] Bernard Chazelle and Jiří Matoušek. On linear-time deterministic algorithms for optimization problems in fixed dimension, *Proceeding of the 4th Annual ACM-SIAM Symposium on Discrete Algorithms*, (1993) pages 281-90.

[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.

[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.

[D92] James Demmel. The componentwise distance to the nearest singular matrix, *SIAM Journal of Matrix Analysis and Applications*, vol. 13, no. 1, (1992) pages 10-19.

[D84] Martin Dyer. On a multidimensional search technique and its application to the Euclidean one-center problem, *SIAM Journal on Computing* 15, (1986) pages 725-38.

[D92] Martin Dyer. A class of convex programs with applications to computational geometry, *Proceedings of the 8th Annual Symposium on Computational Geometry*, (1992) pages 9-15.

[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.

[EW89] P. Egyed and Rephael Wenger. Stabbing pairwise disjoint translates in linear time, *Proceedings of the 5th Annual Symposium on Computational Geometry*, (1989) pages 364-9

[GPW93] Jacob E. Goodman, Richard Pollack and Rephael Wenger. Geometric Transversal Theory, in New Trends in Discrete and Computational Geometry, Springer Verlag, to appear.

[GM61] Branko Grünbaum and Theodore S. Motzkin. On components in some families of sets, *Proceedings of the American Mathematical Society*, vol. 12, (1961) pages 607-13.

[H79] Alan J. Hoffman. Binding constraints and Helly numbers, *Annals of the New York Academy of Sciences* **319**, (1979) pages 284-8.

[K92] Gil Kalai. A subexponential randomized simplex algorithm, *24th Annual ACM Symposium on the Theory of Computation* (1992), pages 475-482.

[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.

[M83] Nimrod Megiddo. Linear programming in linear time when the dimension is fixed, *SIAM Journal on Computing* 12, (1983) pages 759-76.

[M89] Nimrod Megiddo. On the ball spanned by balls, *Discrete and Computational Geometry* 4, (1989) pages 605-10.

[M91] Nimrod Megiddo. *personal communication*, (1991)

[PR92] Svatopluk Poljak and Jiří Rohn. Radius of nonsingularity, *Mathematics of Systems, Signals and Control*, to appear.

[R89] Jiří Rohn. Linear interval equations, *Linear Algebra and its Applications*, vol 126, (1989) pages 39-78.

[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.

[T91] Sivan Toledo. Extremal Polygon Containment Problems, *Proceedings of the 7th Annual Symposium on Computational Geometry*, (1991) pages 176-185.

[T89] Helge Tverberg. Proof of Grünbaum's conjecture on common transversals for translates, *Discrete and Computational Geometry*, vol. 4, (1989) pages 191-203.