DMV

An Update on the Hirsch Conjecture

Edward D. Kim · Francisco Santos

Abstract The Hirsch conjecture was posed in 1957 in a question from Warren M. Hirsch to George Dantzig. It states that the graph of a d-dimensional polytope with n facets cannot have diameter greater than $n - d$. The number n of facets is the minimum number of closed half-spaces needed to form the polytope and the conjecture asserts that one can go from any vertex to any other vertex using at most $n - d$ edges.

Despite being one of the most fundamental, basic and old problems in polytope theory, what we know is quite scarce. Most notably, no polynomial upper bound is known for the diameters that are conjectured to be linear. In contrast, very few polytopes are known where the bound $n - d$ is attained. This paper collects known results and remarks both on the positive and on the negative side of the conjecture. Some proofs are included, but only those that we hope are accessible to a general mathematical audience without introducing too many technicalities.

E.D. Kim
Department of Mathematics, University of California, Davis, Davis, CA 95616, USA
e-mail: ekim@math.ucdavis.edu
url: http://www.math.ucdavis.edu/~ekim/

F. Santos (✉)
Departamento de Matemáticas, Estadística y Computación, Universidad de Cantabria, 39005 Santander, Spain
e-mail: francisco.santos@unican.es
url: http://personales.unican.es/santosf/

1 Introduction

Convex polytopes generalize convex polygons (of dimension two). More precisely, a *convex polyhedron* is any intersection of finitely many affine closed half-spaces in $\mathbb{R}^d$. A *polytope* is a bounded polyhedron. The long-standing Hirsch conjecture is the following very basic statement about the structure of polytopes. Besides its implications in linear programming, which motivated the conjecture in the first place, it is one of the most fundamental open questions in polytope theory.

Conjecture 1.1 (Hirsch conjecture) *Let $n > d \geq 2$. Let P be a d-dimensional polytope with n facets. Then* $\mathrm{diam}(G(P)) \leq n - d$.

Facets are the faces of dimension $d - 1$ of P, so that the number n of them is the minimum number of half-spaces needed to produce P as their intersection (assuming P is full-dimensional). The number $\mathrm{diam}(G(P)) \in \mathbb{N}$ is the *diameter of the graph of P*. Put differently, the conjecture states that we can go from any vertex of P to any other vertex traversing at most $n - d$ edges.

Consider the following examples; all of them satisfy the inequality strictly, except for the cube where it is tight:

P	n	d	$\mathrm{diam}(G(P))$
Polygon	n	2	$\lfloor n/2 \rfloor$
Cube	6	3	3
Icosahedron	20	3	3
Soccer-ball	32	3	9

Polytopes and polyhedra are the central objects in the area of geometric combinatorics, but they also appear in diverse mathematical fields: From the applications' point of view, a polyhedron is the *feasibility region* of a *linear program* [14]. This is the context in which the Hirsch conjecture was originally posed (see below). In *toric geometry*, to every (rational) polytope one associates a certain projective variety (see, e.g., [50]). The underlying interaction between combinatorics and algebraic geometry has proved extremely fruitful for both areas, leading for example to a complete characterization of the possible numbers of faces (vertices, edges, facets, ...) that a *simplicial polytope* can have. The same question for arbitrary polytopes is open in dimension four and higher [60]. Polytopes with special symmetries, such as regular ones and variations of them arise naturally from Coxeter groups and other algebraic structures [6, 25]. Last but not least, counting integer points in polytopes with integer vertex coordinates has applications ranging from number theory and representation theory to cryptography, integer programming, and statistics [5, 16].

In this paper we review the current status of the Hirsch conjecture and related questions. Some proofs are included, and many more appear in an appendix which is available electronically [38]. Results whose proof can be found in [38] are marked with an asterisk. An earlier survey of this topic, addressed to a more specialized audience, was written by Klee and Kleinschmidt in 1987 [40].

1.1 A Bit of Polytope Theory

We now review several concepts that will appear throughout this paper. For further discussion, we refer the interested reader to [18, 59].

A *polyhedron* is the intersection of a finite number of closed half-spaces and a polytope is a bounded polyhedron. A *polytope* is, equivalently, the convex hull of a finite collection of points. Although the geometric objects are the same, from a computational point of view it makes a difference whether a certain polytope is represented as a convex hull or via linear inequalities: the size of one description cannot be bounded polynomially in the size of the other, if the dimension d is not fixed. The *dimension* of a polytope is the dimension of its affine hull aff(P). A d-dimensional polytope is called a *d-polytope*.

If H is a closed half-space containing P, then the intersection of P with the boundary of H is called a *face* of P. Every non-empty face is the intersection of P with a supporting hyperplane. Faces are themselves polyhedra of lower dimension. A face of dimension i is called an *i-face*. The 0-faces are the *vertices* of P, the 1-faces are *edges*, the $(d-2)$-faces are *ridges*, and the $(d-1)$-faces are called *facets*. In its irredundant description, a polytope is the convex hull of its vertices, and the intersection of its facet-defining half-spaces.

For a polytope P, we denote by $G(P)$ its *graph* or 1-*skeleton*, consisting of the vertices and edges of P: the vertices of the graph $G(P)$ are indexed by the vertices of the polytope P, and two vertices in the graph $G(P)$ are connected by an edge exactly when their corresponding vertices in P are contained in a 1-face. The *distance* between two vertices in a graph is the minimum number of edges needed to go from one to the other, and the *diameter* of a graph is the maximum distance between its vertices. (For an unbounded polyhedron, the graph contains only the *bounded* edges. The unbounded 1-faces are called *rays*.)

Example 1.2 Examples of polytopes one can build in every dimension are the following:

1. The d-simplex. The convex hull of $d+1$ points in $\mathbb{R}^d$ that do not lie on a common hyperplane is a d-dimensional simplex. It has $d+1$ vertices and $d+1$ facets. Its graph is complete, so its diameter is 1.
2. The d-cube. The vertices of the d-cube, the product of d segments, are the 2^d points with ± 1 coordinates. Its facets are given by the $2d$ inequalities $-1 \le x_i \le 1$. Its graph has diameter d: the number of steps needed to go from a vertex to another equals the number of coordinates in which the two vertices differ.
3. Cross polytope. This is the convex hull of the d standard basis vectors and their negatives, which generalizes the 3-dimensional octahedron. It has 2^d facets, one in each orthant of $\mathbb{R}^d$. Its graph is almost complete: the only edges missing from it are those between opposite vertices.

See Fig. 1.

Their numbers m of vertices, n of facets, dimension d and diameter are:

P	m	n	d	$\text{diam}(G(P))$
d-Simplex	$d+1$	$d+1$	d	1
d-Cube	2^d	$2d$	d	d
d-Crosspolytope	$2d$	2^d	d	2

Of special importance are the simple and simplicial polytopes. A d-polytope is called *simple* if every vertex is the intersection of exactly d facets. Equivalently, a d-polytope is simple if every vertex in the graph $G(P)$ has degree exactly d. We note that the d-simplices and d-cubes are simple, but cross-polytopes are not simple starting in dimension three. Any polytope or polyhedron P, given by its facet-description, can be perturbed to a simple one P' by a generic and small change in the coefficients of its defining inequalities. This will make non-simple vertices "explode" and become clusters of new vertices, all of which will be simple. This process can not decrease the diameter of the graph, since we can recover the graph of P from that of P' by collapsing certain edges. Hence, to study the Hirsch conjecture, one only needs to consider the simple polytopes:

Lemma 1.3 *The diameter of any polytope P is bounded above by the diameter of some simple polytope P' with the same dimension and number of facets.*

Graphs of simple polytopes are better behaved than graphs of arbitrary polytopes. Their main property in the context of the Hirsch conjecture is that if u and v are vertices joined by an edge in a simple polytope then there is a single facet containing u and not v, and a single facet containing v and not u. That is, at each step along the graph of P we enter a single facet and leave another one.

Every polytope P (containing the origin in its interior, which can always be assumed by a suitable translation) has a polar polytope P^* whose vertices (respectively facets) correspond to the facets (respectively vertices) of P. More generally, every $(d-i)$-face of P^* corresponds to a face of P of dimension $i-1$, and the incidence relations are reversed.

The polars of simple polytopes are called *simplicial*, and their defining property is that every facet is a $(d-1)$-simplex. As an example, the d-dimensional cross polytope is the polar of the d-cube. Since cubes are simple polytopes, cross polytopes are simplicial. The polar of a simplex is a simplex, and simplices are the only polytopes

Fig. 1 Basic examples of polytopes
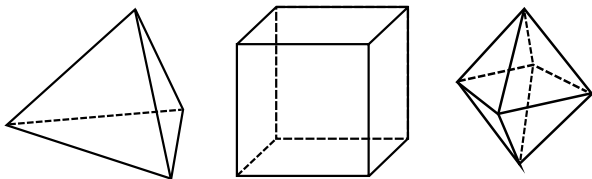
of dimension greater than two which are at the same time simple and simplicial. Since all faces of a simplex are themselves simplices, all faces of a simplicial polytope are simplices. From this viewpoint, one can forget the geometry of P^* and look only at the combinatorics of the simplicial complex formed by its faces, the boundary of P^*. Topologically, this simplicial complex is a sphere of dimension $d - 1$.

For simplicial polytopes we can state the Hirsch conjecture as asking how many ridges do we need to cross in order to walk between two arbitrary facets, if we are only allowed to move from one facet to another via a ridge. This suggests defining the *dual graph* $G^{\Delta}(P)$ of a polytope: The undirected graph having as nodes the facets of P and in which two nodes are connected by an edge if and only if their corresponding facets intersect in a ridge of P. In summary, $G^{\Delta}(P) = G(P^*)$.

1.2 Relation to Linear Programming

The original motivation for the Hirsch conjecture comes from its relation to the simplex algorithm for linear programming. In linear programming, one is given a system of linear equalities and inequalities, and the goal is to maximize (or minimize) a certain linear functional. Every such problem can be put in the following standard form, where A is an $m \times n$ real matrix A, and $\mathbf{b} \in \mathbb{R}^m$ and $c \in \mathbb{R}^n$ are two real vectors:

$$\text{Maximize } c \cdot \mathbf{x}, \text{ subject to } A\mathbf{x} = \mathbf{b} \text{ and } \mathbf{x} \geq 0.$$

Suppose the matrix A has full row rank $m \leq n$. Then, the equality $A\mathbf{x} = \mathbf{b}$ defines a d-dimensional affine subspace ($d = n - m$), whose intersection with the linear inequalities $\mathbf{x} \geq 0$ gives the *feasibility polyhedron P*:

$$P := \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b} \text{ and } \mathbf{x} \geq 0\}.$$

One typically desires not only the maximum value of $c \cdot \mathbf{x}$ but also (the actual coordinates of) a vector $\mathbf{x} \in P$ where the maximum is attained. It is easy to prove that such an $\mathbf{x}$, if it exists, can be found among the vertices of P. If P is unbounded and $c \cdot \mathbf{x}$ does not have an upper bound on it one considers the problem "solved" by describing a ray of P where the value $c \cdot \mathbf{x}$ goes to infinity.

In 1979, Hačijan [30] proved that linear programming problems can be solved in polynomial time via the so-called *ellipsoid method*. In 1984, Karmarkar [37] devised a different approach, the *interior point method*. Although the latter is more applicable (easier to implement, better complexity) than the former, still to this day the most commonly used method for linear programming is the *simplex method* devised by G. Dantzig in 1947. For a complete account of the complexity of linear programming, see the survey [48] by Megiddo.

In geometric terms, the simplex method first finds an arbitrary vertex in the feasibility polyhedron P. Then, it moves from vertex to adjacent vertex in such a way that the value $c \cdot \mathbf{x}$ of the linear functional increases at every step. These steps are called *pivots* and the rule used to choose one specific adjacent vertex is called the *pivot rule*. When no pivot step can increase the functional, convexity implies that we have arrived to the global maximum.

Clearly, a lower bound for the performance of the simplex method under *any* pivot rule is the diameter of the polyhedron P. The converse is not true, since knowing that

P has a small graph diameter does not in principle tell us how to go from one vertex to another in a small number of steps. In particular, many of the results on diameters of polyhedra do not help for the simplex method.

In fact, the complexity of the simplex method depends on the local rule (the *pivot rule*) chosen to move from vertex to vertex. The *a priori* best pivot rule, the one originally proposed by Dantzig, is "move along the edge with maximum gradient", but Klee and Minty [41] showed in 1972 that this can lead to paths of exponential length, even in polytopes combinatorially equivalent to cubes. The same worst-case exponential behavior has been proved for essentially every deterministic rule devised so far, although there are subexponential, but yet not polynomial, randomized pivot algorithms (see Theorem 2.8). However, the simplex algorithm is highly efficient *in practice* on most linear optimization problems.

There is another reason why investigating the complexity of the simplex method is important, even if we already know polynomial time algorithms. The algorithms of Khachiyan and Karmarkar are polynomial in the *bit length* of the input; but it is of practical importance to know whether a polynomial algorithm for linear programming in the *real number machine* model of Blum, Cucker, Shub, and Smale [7] exists. That is, is there an algorithm that uses a number of arithmetic operations that is polynomial on the number of coefficients of the linear program, rather than on their total bit-length; or, better yet, a *strongly polynomial algorithm*, i.e., one that is polynomial both in the arithmetic sense and the bit sense? These two related problems were included by Smale in his list of "mathematical problems for the next century" [54]. A polynomial pivot rule for the simplex method would solve them in the affirmative.

In this context, the following *polynomial version* of the conjecture is relevant, if the linear one turns out to be false. See, for example, [35]:

Conjecture 1.4 (Polynomial Hirsch conjecture) *Is there a polynomial function $f(n, d)$ such that for any polytope (or polyhedron) P of dimension d with n facets,* $\text{diam}(G(P)) \leq f(n, d)$?

1.3 Overview of This Paper

Our initial purpose with this paper was two-fold: on the one hand, we thought it is about time to have in a single source an overview of the state of the art concerning the Hirsch conjecture and related issues, putting up to date the 20 year old survey by Klee and Kleinschmidt [40]. On the other hand, since there seems to be agreement on the fact that the Hirsch conjecture is probably false (with opinions about the polynomial version of it being divided) we wanted to give a fresh look at the past attempts to *disprove* the conjecture.

These two goals turned out to be in conflict, or at least too ambitious, so the first version of the paper was too long and too technical for the intended readership. After wise comments from our editor Jörg Rambau and an anonymous referee, we decided to take most of the proofs out of the main paper, and compiled them in the companion paper [38]. Results whose proof can be found in [38] are marked with an asterisk.

Our two-fold intentions are still reflected in the two quite distinct parts that the paper has, Sects. 2 and 3. The first one is devoted to positive results, comprising general upper bounds for polytope diameters, special cases where the Hirsch conjecture

is known, etc., and the second one contains mainly constructions and results aimed at disproving the Hirsch conjecture.

The two sections differ in several respects: Sect. 2 is written in an informative style. No proofs are included (although some appear in [38]) since this section covers quite different topics and the techniques and ideas used are too technical and, more importantly, too diverse. In Sect. 3, on the contrary, we provide proofs for essentially all the results (some here and some in [38]); on the one hand the tools needed are more homogeneous and elementary; on the other hand, in this section we feel that having a new look at old results is useful. We have tried to identify the basic ingredients in each construction, obtaining in some cases much simpler (to our taste) proofs and expositions than the original ones. In particular, the main novelty in this survey, if any, is probably in our descriptions of the non-Hirsch polyhedron and Hirsch-sharp polytope found by Klee and Walkup in 1967 (Sect. 3.3) and of the non-Hirsch sphere found by Mani and Walkup in 1980 (Sect. 3.6).

Another difference between the two sections is that Sect. 2 contains several very recent developments, while all of Sect. 3, with the single exception of Theorem 3.11, refers to results that are at least 25 years old.

Let us now give a brief roadmap for the paper.

Section 2.1 lists the pairs of parameters (n, d) such that the Hirsch Conjecture is known to hold for all d-polytopes with n facets. That is, denoting $H(n, d)$ the maximum diameter of d-polytopes with n facets, we list all pairs for which the Hirsch inequality $H(n, d) \leq n - d$ is known to hold. This comprises the cases $d \leq 3$ (Klee [39]), $n - d \leq 6$ (Bremner and Schewe [11]), and $(n, d) \in \{(11, 4), (12, 4)\}$ (Bremner et al. [10]).

The next section lists general upper bounds on $H(n, d)$: a linear one in fixed dimension (Barnette and Larman) [4, 44] and a quasi-polynomial one of $n^{\log_2(d)+1}$ (Kalai-Kleitman [36]). These bounds hold not only for diameters of polytopes but also for much more abstract and general objects. A very recent development by Eisenbrand, Hähnle, Razborov and Rothvoß [24] is the identification of one such class for which the proofs of these two bounds work but which admit objects with quadratic diameter. This may be considered evidence against the Hirsch conjecture.

In Sect. 2.3 we concentrate on algorithmic aspects. For example, we state two algorithmic analogues of the two bounds mentioned above: the proof by Meggido [47] that linear programming can be done in linear time if the dimension is fixed, and randomized pivot rules for the simplex method in arbitrary dimension that finish in $O(\exp(K\sqrt{d \log n}))$ (Kalai [34], Matoušek, Sharir and Welzl [46]).

We then turn our attention to special polytopes for which good bounds are known. Polytopes with 0-1 coordinates and linear programming duals of transportation polytopes are known to satisfy the Hirsch conjecture (Naddef [49], Balinski [3]). Diameters of network-flow polytopes, which include transportation polytopes, have quadratic bounds [13, 28, 51].

Section 2 finishes with an account of recent work of Deza, Terlaky and Zinchenko [20–22] on a continuous analogue of the Hirsch conjecture that arises in the context of interior point methods for linear programming.

Almost all of the results in Sect. 3 revolve around two basic ingredients. The first one is the *wedge* operation, which we describe in Sect. 3.1. Wedging is a very simple

operation that increases both the dimension and number of facets of a polytope by one maintaining (or increasing) its diameter. Using this, it is easy to prove the following fundamental result:

Theorem 1.5 (Klee and Walkup [42]) $H(d + k, d) \leq H(2k, k)$, *with equality if* (*but not only if*) $k < d$.

In particular, to prove (or disprove) the Hirsch conjecture one can concentrate on the case where the number of facets equals twice the dimension. This case is sometimes referred to as the *d-step Conjecture*, since the Hirsch conjecture is saying that we can go from any vertex to any other vertex in d-steps. Via wedging, the Hirsch conjecture is also equivalent to the following *non-revisiting Conjecture*: If u and v are two arbitrary vertices of a simple polytope P, then there is a path from u to v which at every step enters a facet of P that was not visited before. We prove the equivalence of the three conjectures (Hirsch, d-step and non-revisiting) in Sect. 3.2.

The second ingredient is the construction by Klee and Walkup [42] of a 4-dimensional polytope with 9 facets that meets the Hirsch bound with equality (that is, whose diameter equals $9 - 4 = 5$). Polytopes with this property are called *Hirsch-sharp*. They are easy to construct with a number of facets not exceeding twice their dimension (e.g., cubes). Klee and Walkup's Hirsch-sharp polytope is the smallest "non-trivial" Hirsch-sharp polytope, with more facets than twice its dimension. In fact, it is also the starting block to the construction of every other Hirsch-sharp polytope with $n > 2d$ known to date. In Sect. 3.3 we give our own description and coordinatization (much smaller than the original one) of the Klee-Walkup polytope.

In Sect. 3.4 we recount the state of the art on the existence of Hirsch-sharp polytopes, following work of Fritzsche, Holt and Klee [26, 31, 32]. Their results, combined with what is known for small dimension or number of facets, are summarized in Table 1, which gives a "plot" of the function $H(n, d) - (n - d)$. The horizontal coordinate is $n - 2d$, so that the column labeled "0" corresponds to the polytopes relevant to the d-step conjecture. The cases where we know $H(n, d)$ exactly are marked "=" or "<" depending on whether Hirsch-sharp polytopes exist or not. The cases where Hirsch-sharp polytopes are known to exist but for which the Hirsch conjecture is not proved are marked "≥". Cases where we neither know the Hirsch conjecture nor the existence of Hirsch-sharp polytopes are marked "?" and appear only in dimensions 4, 5 and 6. The diagonal dots in the left column reflect the equality case of Theorem 1.5.

The Klee-Walkup polytope is also instrumental in the construction by Klee, Walkup and Todd [42, 56] of counter-examples to two generalizations of the Hirsch conjecture that are quite natural in the context of linear programming: the case of perhaps-*unbounded* polyhedra (which was the original conjecture by Hirsch) and a *monotone* version in which we look at the maximum number of monotone steps with respect to a given linear function that are needed to go from any vertex of a polytope P to an optimal vertex. We show these constructions in Sect. 3.5, and show in Sect. 3.6 a counter-example, by Mani and Walkup [45] to a third, topological, version of the conjecture.

Table 1 $H(n, d)$ versus $n - d$, the state of the art

d \ $n - 2d$	$\cdots$	0	1	2	3	4	5	6	7	$\cdots$
2		=	<	<	<	<	<	<	<	$\cdots$
3	$\cdot\cdot^{\cdot}$	=	<	<	<	<	<	<	<	$\cdots$
4	$\cdot\cdot^{\cdot}$	=	=	<	<	<	?	?	?	$\cdots$
5	$\cdot\cdot^{\cdot}$	=	=	?	?	?	?	?	?	$\cdots$
6	$\cdot\cdot^{\cdot}$	=	?	?	?	?	?	?	?	$\cdots$
7	$\cdot\cdot^{\cdot}$	$\geq$	$\geq$	$\geq$	$\geq$	$\geq$	$\geq$	$\geq$	$\geq$	$\cdots$
8	$\cdot\cdot^{\cdot}$	$\geq$	$\geq$	$\geq$	$\geq$	$\geq$	$\geq$	$\geq$	$\geq$	$\cdots$
$\vdots$	$\cdot\cdot^{\cdot}$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$

2 Bounds and Algorithms

In this section, we present special cases for which the Hirsch conjecture holds, upper bounds for diameters of polytopes and subexponential complexity results for the simplex method. We also summarize recent work on analogues of the conjecture for hyperplane arrangements and for paths of interior point methods.

2.1 Small Dimension or Few Facets

The following statements exhaust all pairs (n, d) for which the maximum diameter $H(n, d)$ of d-polytopes with n facets is known. We omit the cases $n < 2d$, because $H(d + k, d) = H(2k, k)$ for all $k < d$ (see Theorem 1.5), and the trivial case $d \leq 2$. Remember that an asterisk in front of a statement denotes the proof can be found in [38].

***Theorem 2.1** (Klee [39]) $H(n, 3) = \lfloor \frac{2n}{3} \rfloor - 1$.

Theorem 2.2

- $H(8, 4) = 4$ (Klee [39]).
- $H(9, 4) = H(10, 5) = 5$ (Klee-Walkup [42]).
- $H(10, 4) = 5$, $H(11, 5) = 6$ (Goodey [29]).
- $H(11, 4) = H(12, 6) = 6$ (Bremner-Schewe [11]).
- $H(12, 4) = 7$ (Bremner et al. [10]).

Since $\max_d H(d + k, d) = H(2k, k)$ (see Theorem 1.5 again), the results for $H(8, 4)$, $H(10, 5)$ and $H(12, 6)$ imply:

Corollary 2.3 *The Hirsch conjecture holds for polytopes with at most six facets more than their dimension*.

It is easy to generalize one direction of Theorem 2.1, giving the following lower bound for $H(n, d)$. Observe that the formula gives the exact value of $H(n, d)$ for $d \in \{1, 2, 3\}$.

****Proposition 2.4**

$$H(n, d) \geq \left\lfloor \frac{d-1}{d} n \right\rfloor - (d-2).$$

2.2 General Upper Bounds on Diameters

Diameters of polytopes admit a *linear* upper bound when the dimension d is fixed. This was first noticed by Barnette [4] and then improved by Larman [44]:

****Theorem 2.5** (Larman [44]) *For every $n > d \geq 3$, $H(n, d) \leq n2^{d-3}$.*

But when the number of facets is not much bigger than d, a much better upper bound was given by Kalai and Kleitman [36], with a surprisingly simple and elegant proof (the paper is just two pages!).

****Theorem 2.6** (Kalai-Kleitman [36]) *For every $n > d$, $H(n, d) \leq n^{\log_2(d)+1}$.*

The proofs of Theorems 2.6 and 2.5 use very limited properties of graphs of polytopes. For example, Klee and Kleinschmidt (see Sect. 7.7 in [40]) show that Theorem 2.5 holds for the ridge-graphs of all pure simplicial complexes, and even more general objects. In the same vein, Eisenbrand, Hähnle, Razborov and Rothvoß [24] have recently shown the following generalization of Theorems 2.6 and 2.5:

Theorem 2.7 (Eisenbrand et al. [24]) *Let G be a graph whose vertices are certain subsets of size d of an n-element set. Assume that between every pair of vertices u and v in G there is a path using only vertices that contain $u \cap v$.*
Then, $\mathrm{diam}(G) \leq \min\{n^{1+\log d}, n2^{d-1}\}$.

The novelty in [24] is that the authors show that there are graphs with the hypotheses of Theorem 2.7 and with $\mathrm{diam}(G) \geq cn^2/\log n$, for arbitrarily large n and a certain constant c. It is not clear whether this is support against the Hirsch conjecture or it simply indicates that the arguments in the proofs of Theorems 2.6 and 2.5 do not take advantage of properties that graphs of polytopes have and which prevent their diameters from growing. For example, observe that *any* connected graph is valid for the case $d = 1$ of Theorem 2.7.

2.3 Subexponential Simplex Algorithms

Since the Hirsch conjecture is strongly motivated by the simplex algorithm of linear programming, it is natural to ask about the number of iterations needed under particular pivot rules. Most of the proofs for the upper bounds in the previous sections do not give a clue on how to find a short path towards the vertex maximizing a given functional, or even an explicit path between any pair of given vertices.

Kalai [34] and, independently, Matoušek, Sharir and Welzl [46] proved the existence of randomized pivot rules for the simplex method with subexponential running time for arbitrary linear programs.

Theorem 2.8 (Kalai [34], Matoušek, Sharir and Welzl [46]) *There exist randomized simplex algorithms where the expected number of arithmetic operations needed in the worst case is at most* $\exp(K\sqrt{d\log n})$, *where* K *is a fixed constant.*

If we consider Theorem 2.8 as an algorithmic analogue of Theorem 2.6, then the following result of Megiddo is the analogue of Theorem 2.5. It says that linear programming can be performed in linear time in fixed dimension:

Theorem 2.9 (Meggido [47]) *There are pivot rules for the simplex algorithm that run in* $O(2^{2^d}n)$ *time.*

It is also known that random polytopes have polynomial diameter, which explains why the simplex method seems to work well in practice. The first results in this direction were proved by Borgwardt [8] and, independently, Smale [53], who analyzed the "average case" complexity of the simplex method. Average case means that we are looking at a linear program

$$\text{Maximize } c \cdot \mathbf{x}, \text{ subject to } A\mathbf{x} = b \text{ and } \mathbf{x} \geq 0,$$

but the entries of A, b and c are considered random variables with respect to certain spherically symmetric probability distributions. In Borgwardt's model the simplex method runs in expected polynomial time in the size of the input. Smale shows that in his model, if one of the parameters d and $n - d$ is fixed and the other is allowed to grow, the expected running time is only polylogarithmic. The latter was improved to constant by Megiddo, see [48] for details.

Even more surprising is the fact that *every* linear program can be slightly perturbed to one that can be solved in polynomial time. Let us formalize this. Let P be the feasibility polyhedron

$$P = \{\mathbf{x} \in \mathbb{R}^d \mid \langle a_i, \mathbf{x} \rangle \leq b, (i = 1, \ldots, n)\}$$

of a certain linear program. If we replace the vectors $a_i \in \mathbb{R}^d$ and $b \in \mathbb{R}^n$ with independent Gaussian random vectors with means $\mu_i = a_i$ and $\mu = b$ (respectively), and standard deviations $\sigma \max_i \|(\mu_i, \mu)\|$ we say that we have *perturbed* P *randomly within a parameter* σ. In [55], Spielman and Teng proved that the expected diameter of a linear program that is perturbed within a parameter σ is polynomial in d, n, and σ^{-1}. In [57], Vershynin improved the bound to be polylogarithmic in n and polynomial only in d and σ^{-1}.

Theorem 2.10 (Vershynin [57]) *If a linear program is perturbed randomly within a parameter* σ, *then its expected diameter of its feasibility polyhedron is in* $O(\log^7 n(d^9 + d^3\sigma^{-4}))$.

As mentioned above, this result is not only structural. The simplex method can *find* a path of that expected length in the perturbed polyhedron.

2.4 Some Polytopes from Combinatorial Optimization

There are some classes of polytopes of special interest and for the diameters of which we know polynomial upper bounds.

2.4.1 Small Integer Coordinates

Of special importance in combinatorial optimization are the 0-1 polytopes, in which every vertex has coordinates 0 or 1. They satisfy the Hirsch conjecture.

****Theorem 2.11** (Naddef [49]) *If P is a 0-1 polytope then*

$$\text{diam}(P) \leq \#facets(P) - \dim(P).$$

As a generalization, Kleinschmidt and Onn [43] prove the following bound on the diameter of lattice polytopes in $[0, k]^d$. A polytope is called a *lattice polytope* if every coordinate of every vertex is integral.

Theorem 2.12 (Kleinschmidt and Onn [43]) *The diameter of a lattice polytope contained in $[0, k]^d$ cannot exceed kd.*

However, existence of a polynomial pivot rule for the simplex method in 0-1 polytopes is open. The proof of Theorem 2.11 constructs a short path from u to v only assuming that we know the coordinates of both.

2.4.2 Network-Flow Polytopes

A network flow polytope is defined by an arbitrary directed graph $G = (V, E)$ with weights given to its vertices. Negative weights represent demands and positive weights represent supplies. A flow is an assignment of non-negative numbers to the edges so as to cancel all the demands and supplies. See [13], [28], and [51] for details.

For any network G with e edges and v vertices, every sufficiently generic set of vertex weights produces a simple $(e - v + 1)$-dimensional polytope with at most $2e$ facets. Its diameter has the following almost quadratic upper bound. The proof yields a polynomial time pivot rule for the simplex method on these polytopes.

Theorem 2.13 [13, 28, 51] *The diameter of the network flow polytope on a directed graph $G = (V, E)$ is in $O(ev \log v)$. This, in turn, is in $O(n^2 \log n)$, where n is the number of facets of the polytope.*

The matrices defining network flow polytopes are examples of *totally unimodular* matrices, meaning that all its subdeterminants are 0, 1, or -1. Polytopes defined by these matrices still have polynomially bounded diameters, although the degree in the bound is much worse than the one for network flow polytopes:

Theorem 2.14 (Dyer and Frieze [23]) *For any totally unimodular $n \times d$ matrix A the diameter of the polyhedron $\{\mathbf{x} \in \mathbb{R}^d : A\mathbf{x} \leq c\}$ is in $O(d^{16} n^3 (\log(dn))^3)$.*

2.4.3 Transportation and Dual Transportation Polytopes

Given vectors $a \in \mathbb{R}^p$ and $b \in \mathbb{R}^q$, the $p \times q$ transportation polytope defined by a and b is the set of all $p \times q$ non-negative matrices with row sums given by a and column sums given by b:

$$T_{p,q}(a,b) = \left\{ (x_{ij}) \in \mathbb{R}^{p \times q} \;\middle|\; \sum_j x_{ij} = a_i, \sum_i x_{ij} = b_j, x_{ij} \geq 0 \right\}.$$

As an example, the Birkhoff polytope, whose vertices are the permutation matrices, is the transportation polytope obtained with $p = q$ and $a = b = (1, \dots, 1)$.

It is easy to show that (generically) $T_{p,q}(a,b)$ is a $(p-1)(q-1)$-dimensional polytope with at most pq facets. Thus, the Hirsch conjecture translates to its diameter being at most $p + q - 1$.

Transportation polytopes are a special case of network flow polytopes; they arise when the network is a complete bipartite graph on p and q nodes with all edges directed in the same direction. In particular, Theorem 2.13 gives an almost quadratic bound for their diameters. But Brightwell et al. [12] have recently proved a *linear* bound, with a multiplicative factor of eight. This has now been improved to:

Theorem 2.15 (Hurkens [33]) *The diameter of any $p \times q$ transportation polytope is at most $3(p + q - 1)$.*

In the context of linear programming, for every d-polyhedron with n facets there is a dual $(n-d)$-polyhedron with the same number of facets. Every linear program can be solved in its "primal" or "dual" polyhedron. The optimum achieved is the same in both, but the complexity of the algorithm may not.

The linear programming duals of $p \times q$ transportation polytopes are $(p + q - 1)$-polyhedra with pq facets. Balinski [3] proved the Hirsch conjecture for them.

Theorem 2.16 (Balinski [3]) *Let C be a $p \times q$ matrix. The diameter of the dual transportation polytope $D_{p,q}(C)$ is at most $(p-1)(q-1)$. This bound is the best possible and it yields a polynomial time dual simplex algorithm.*

2.4.4 3-Way Transportation Polytopes

Another seemingly harmless generalization of transportation polytopes comes from considering 3-way tables $(x_{ijk})_{ijk} \in \mathbb{R}^{p \times q \times r}$ instead of matrices. In fact, there are two different such generalizations. An *axial 3-way transportation polytope* consists of all non-negative tables with fixed sums in 2-d slices. A *planar 3-way transportation polytope* consists of all non-negative tables with fixed sums in 1-d slices. (The names seem wrong, but they reflect the fact that an axial transportation polytope is defined by three *vectors* of sizes p, q and r, and a planar one by three *matrices* of sizes $p \times q$, $p \times r$, and $q \times r$.)

Despite their definition being so close to that of transportation polytopes, 3-way transportation polytopes are *universal* in the following sense:

Theorem 2.17 (De Loera and Onn [17]) *Let P be a rational convex polytope.*

- *There is a 3-way planar transportation polytope Q isomorphic to P.*
- *There is a 3-way axial transportation polytope Q which has a face F isomorphic to P.*

In both cases there is a polynomial time algorithm to construct Q (and F).

Isomorphic here means affinely (and rationally) equivalent. In particular, that the polytope Q or its face F have the same edge-graph as P. Thus, it was interesting to try to apply to the 3-way case the methods that gave polynomial upper bounds for the graphs of transportation polytopes. This was attempted in [15], where a quadratic upper bound was obtained but only for axial transportation polytopes. A generalization of this result to *faces* of them or to *planar* transportation polytopes would prove the polynomial Hirsch conjecture.

Theorem 2.18 (De Loera, Kim, Onn, Santos [15]) *The diameter of every 3-way axial $p \times q \times r$ transportation polytope is at most $2(p + q + r - 3)^2$.*

2.5 A Continuous Hirsch Conjecture

Here we summarize some recent work of Deza, Terlaky and Zinchenko [20–22] in which they propose continuous analogues of the Hirsch and d-step conjectures related to the central path method—a variant of interior point methods—of linear programming. (For a complete description of the method we refer the reader to [9, 52].) The analogy comes from analyzing the total curvature $\lambda_c(P)$ of the central path with respect to a certain cost function c for the polyhedron P. By analogy with $H(n, d)$, let $\Lambda(n, d)$ denote the largest total curvature of the central path over all polytopes P of dimension d defined by n inequalities and over all linear objective functions c.

It had been conjectured that $\lambda_c(P)$ is bounded by a constant for each dimension d, and that it grows at most linearly with varying d. Deza et al. have disproved both statements: in [20], they construct polytopes for which $\lambda_c(P)$ grows exponentially with d. More strongly, in [22] they construct a family of polytopes that show that λ_c cannot be bounded only in terms of d:

Theorem 2.19 [22] *For every fixed dimension $d \geq 2$, $\liminf_{n \to \infty} \frac{\Lambda(n,d)}{n} \geq \pi$.*

Deza et al. consider this result a continuous analogue of the existence of Hirsch-sharp polytopes. Motivated by this they pose the following conjecture:

Conjecture 2.20 (Continuous Hirsch conjecture) $\Lambda(n, d) \in O(n)$. *That is, there is a constant K such that $\Lambda(n, d) \leq Kn$ for all n and d.*

Theorem 2.19 says that if the continuous Hirsch conjecture is true, then it is tight, modulo a constant factor. Deza et al. also conjecture a continuous variant of the d-step conjecture, and show it to be equivalent to the continuous Hirsch conjecture, thus providing an analogue of Theorem 3.2:

Conjecture 2.21 (Continuous d-step conjecture) *The function $\Lambda(2d, d)$ grows linearly in its input. That is to say, $\Lambda(2d, d)$ is in $O(d)$.*

Theorem 2.22 [21] *The continuous Hirsch conjecture is equivalent to the continuous d-step conjecture. That is, if $\Lambda(2d, d) \in O(d)$ for all d, then $\Lambda(n, d) \in O(n)$ for all d and n.*

The best upper bound known for $\Lambda(n, d)$ is a bound in $O(n^d)$, derived from Theorem 2.23 below. This theorem refers to the central path curvature for hyperplane arrangements, as studied by Dedieu, Malajovich and Shub [19].

An arrangement $\mathcal{A}$ of n hyperplanes in dimension d is called *simple* if every n hyperplanes intersect at a unique point. It is easy to show that any simple arrangement of n hyperplanes in $\mathbb{R}^d$ has exactly $s = \binom{n-1}{d}$ bounded full-dimensional cells. For a simple arrangement $\mathcal{A}$ with bounded cells $P_1, \ldots, P_s$ and a given objective function c, Dedieu et al. consider the quantity: $\lambda_c(\mathcal{A}) = \frac{1}{s} \sum_{i=1}^{s} \lambda_c(P_i)$. That is, the average total curvature of central paths of all bounded cells in the arrangement. They prove:

Theorem 2.23 [19] $\lambda_c(\mathcal{A}) \leq 2\pi d$, *for every simple arrangement.*

Put differently, even if individual cells can give total curvature linear in n by Theorem 2.19, the average over all cells of a given arrangement is bounded by a function of d alone.

Turning the analogy back to polytope graphs, Deza et al. [22] consider the average diameter of the graphs of all bounded cells in a simple arrangement $\mathcal{A}$. Denote it diam($\mathcal{A}$) and let $\mathcal{H}(n, d)$ be the maximum of diam($\mathcal{A}$) over all simple arrangements defined by n hyperplanes in dimension d. They relate $\mathcal{H}(n, d)$ to the Hirsch conjecture, as follows:

Proposition 2.24 [22] *The Hirsch conjecture implies $\mathcal{H}(n, d) \leq d + \frac{2d}{n-1}$.*

3 Constructions

We now move to interesting constructions of polytopes motivated by or related to the Hirsch conjecture. All the proofs that are not included in this section, plus additional comments, can be found in [38].

3.1 The Wedge Operation

Wedging is a very basic, yet extremely fruitful, operation that one can do to a polytope. Its simplicial counter-part is the *one-point suspension* (see [18, 38]).

Roughly speaking, the wedge of P at a facet F of it is the polytope, of one dimension more, obtained gluing two copies of P along F. See Fig. 2 for an example. More formally, let $f(x) \leq b$ be an inequality defining the facet F. The wedge of P over F is the polytope

$$\mathrm{W}_F(P) := P \times [0, \infty) \cap \{(x, t) \in \mathbb{R}^d \times \mathbb{R} : f(x) + t \leq b\}.$$

Fig. 2 A 5-gon and a wedge on its top facet
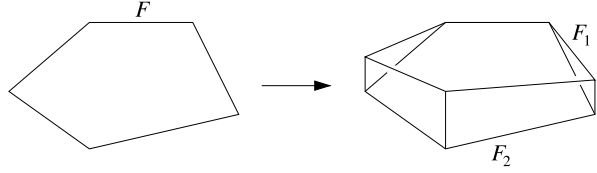
Put differently, $W_F(P)$ is formed by intersecting the half-cylinder $C := P \times [0, \infty)$ with a closed half-space J in $\mathbb{R}^{d+1}$ such that:

- the intersection $J \cap C$ is bounded and has nonempty interior, and
- $\partial J \cap C = F$.

Lemma 3.1 *Let P be a d-polytope with n facets. Let $W_F(P)$ be its wedge on a certain facet F. Then, $W_F(P)$ has dimension $d + 1$, $n + 1$ facets, and*

$$\mathrm{diam}(W_F(P)) \geq \mathrm{diam}(P).$$

Proof The wedge increases both the dimension and the number of facets by one. Indeed, $W_F(P)$ has a vertical facet projecting to each facet of P other than F, plus the two facets that cut the cylinder $P \times \mathbb{R}$, and whose intersection projects to F. The diameter of $W_F(P)$ is at least that of P, since every edge of $W_F(P)$ projects either to an edge of P or to a vertex of P. □

In particular, if P is Hirsch-sharp then $W_F(P)$ is either Hirsch-sharp or a counterexample to the Hirsch conjecture. The properties that P would need for the latter to be the case will be made explicit in Remark 3.5.

As a corollary of Lemma 3.1 we get that in order to prove (or disprove) the Hirsch conjecture it is sufficient to restrict attention to the case when the number of facets equals twice the dimension:

Theorem 3.2 (Klee-Walkup [42]) $H(k + d, d) \leq H(2k, k)$, *with equality if (but not necessarily only if) $k < d$.*

Proof By Lemma 3.1,

$$H(n, d) \leq H(n + 1, d + 1), \quad \forall n, d \tag{1}$$

so we only need to show that

$$H(n, d) \leq H(n - 1, d - 1) \quad \forall n < 2d. \tag{2}$$

Let P be a polytope with $n < 2d$ and let u and v be vertices of it. Since each vertex is incident to at least d facets, u and v lie in a common facet. This facet F has dimension $d - 1$, and each facet of it is the intersection of F with another facet of P. Hence, F has at most $n - 1$ facets itself. Since every path on F is also a path on P, we get (2). □

Fig. 3 The pentagon P and the wedge $P' = W_F(P)$ over its facet F: the upper pentagonal facet of P' is F_1 and the lower pentagonal facet is F_2
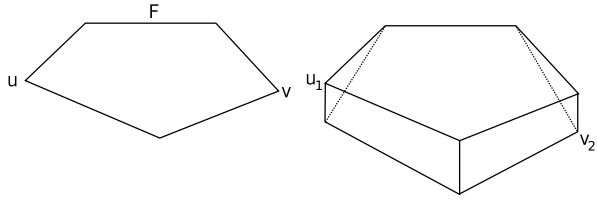
3.2 The d-Step and Non-revisiting Conjectures

The intuition behind the Hirsch conjecture is that to go from vertex u to vertex v of a polytope P, one does not expect to have to enter and leave the same facet several times. This suggests the following conjecture:

Conjecture 3.3 (The non-revisiting conjecture) *Let P be a simple polytope. Let u and v be two arbitrary vertices of P. Then, there is a path from u to v which at every step enters a facet of P that was not visited before.*

Paths with the conjectured property, that they do not revisit any facet, are called *non-revisiting paths*. (In the literature, they are also called W_v paths and Conjecture 3.3 is also known as the W_v conjecture.) Non-revisiting paths are never longer than $n - d$: at each step, we must enter a different facet, and the d facets that the initial vertex lies in cannot be among them. Thus, the non-revisiting conjecture implies the Hirsch conjecture. It turns out both are equivalent. A first step in the proof is the following analogue of Theorem 3.2 for the non-revisiting conjecture:

Theorem 3.4 *If all k-polytopes with $2k$ facets has the non-revisiting property, then the same holds for all d-polytopes with $d + k$ facets, for all d.*

Proof Let P be a polytope with $n \neq 2d$ and suppose it does not have the non-revisiting property. That is, there are vertices u and v such that every path from u to v revisits some facet that it previously abandons. We will construct another polytope P' without the non-revisiting property and with:

- One less facet and dimension than P if $n < 2d$, and
- One more facet and dimension than P if $n > 2d$.

In the first case, u and v lie in a common facet F and we simply let $P' = F$. In the second case, let F be a facet not containing u nor v and let $P' = W_F(P)$ be the wedge over F. Let F_1 and F_2 be the two facets of P' whose intersection projects to F. Let u_1 and v_2 be the vertices of P' that project to u and v and lie, respectively, on F_1 and F_2 (see Fig. 3). Now, consider a path from u_1 to v_2 on P' and project it to a path from u to v on P:

- If the path on P revisits a facet (call it G) other than F, then the path on P' revisits the facet that projects to G.
- If the path on P revisits F, then the path on P' revisits either F_1 or F_2.

$\qquad\square$

Remark 3.5 In the proof of Lemma 3.1 we noted that, applied to a Hirsch-sharp polytope P the wedge operator produced either another Hirsch-sharp polytope or a counterexample to the Hirsch conjecture. The last proof shows that the latter can happen only if P does not have the non-revisiting property.

Theorems 3.2 and 3.4 say that both in the Hirsch and the non-revisiting conjectures the crucial case is that of $n = 2d$. It is not surprising then that they are equivalent, since in this case they both almost restrict to the following:

Conjecture 3.6 (The d-step conjecture) *Let P be a simple d-polytope with $2d$ facets and let u and v be two complementary vertices (i.e., vertices not lying in a common facet). Then, there is a path of length d from u to v.*

There is still something to be proved, though. If u and v are not complementary vertices in a d-polytope with $2d$ facets then the d-step conjecture does not directly imply the other two. But in this case u and v lie in a common facet, so the proof of equivalence is not hard to finish via induction:

*Theorem 3.7** (Klee-Walkup [42]) *The Hirsch, non-revisiting, and d-step Conjectures* 1.1, 3.3, *and* 3.6 *are equivalent.*

3.3 The Klee-Walkup Polytope Q_4

In their seminal 1967 paper [42] on the Hirsch conjecture and related issues, Klee and Walkup describe a 4-polytope Q_4 with nine facets and diameter five. Innocent as this might look, this first "non-trivial" Hirsch-sharp polytope is at the basis of the construction of every remaining Hirsch-sharp polytope known to date (see Sect. 3.4.2). It is also instrumental in disproving the unbounded and monotone variants of the Hirsch conjecture, which we will discuss in Sect. 3.5. Moreover, its existence is something of an accident: Altshuler, Bokowski and Steinberg [2] list all combinatorial types of simplicial spheres with nine vertices (there are 1296, 1142 of them polytopal); among them, the polar of Q_4 is the only one that is Hirsch-sharp.

Here we describe Q_4 in the polar view. That is, we will describe a simplicial 4-polytope Q_4^* with nine vertices and show that its ridge-diameter is five. The vertices of Q_4^* are:

$$a := (-3, 3, 1, 2), \qquad e := (3, 3, -1, 2),$$
$$b := (3, -3, 1, 2), \qquad f := (-3, -3, -1, 2),$$
$$c := (2, -1, 1, 3), \qquad g := (-1, -2, -1, 3),$$
$$d := (-2, 1, 1, 3), \qquad h := (1, 2, -1, 3),$$
$$w := (0, 0, 0, -2).$$

[The simple polytope Q_4 is obtained converting each vertex v of Q_4^* into an inequality $v \cdot \mathbf{x} \leq 1$. For example, the inequality corresponding to vertex a above is $-3x_1 + 3x_2 + x_3 + 2x_4 \leq 1$.]
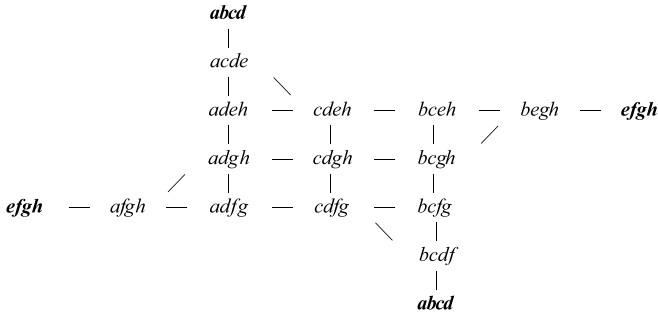
The key property of this polytope is that:

abcd
|
acde
|
adeh — cdeh — bceh — begh — **efgh**
|
adgh — cdgh — bcgh
|
efgh — afgh — adfg — cdfg — bcfg
|
bcdf
|
abcd

Fig. 4 The dual graph of the subcomplex K

Theorem 3.8 (Klee-Walkup [42]) *Any path in Q_4^* from the tetrahedron abcd to the tetrahedron efgh needs at least five steps.*

To prove this, you may simply input these coordinates into any software able to compute the (dual) graph of a polytope. Our suggestion for this is `polymake` [27]. But we believe that fully understanding this polytope can be the key to the construction of counter-examples to the Hirsch conjecture, so it is worth presenting a hybrid computer-human proof. It is worth mentioning that the coordinates we use for Q_4^*, much smaller than the original ones in [42], were obtained as a by-product of the description of Q_4^* contained in this proof.

Proof Paths through some intermediate tetrahedron containing the vertex w necessarily have at least five steps: apart of the step that introduces w, four more are needed to introduce, one by one, the four vertices e, f, g and h.

This means we can concentrate on the subcomplex K of ∂Q_4^* consisting of tetrahedra that do not use w. This subcomplex is called the *anti-star* of w in ∂Q_4^*. We claim (without proof, here is where you need your computer) that this subcomplex consists of the 15 tetrahedra in Fig. 4.

Figure 4 shows adjacencies among tetrahedra; that is, it shows the dual graph of K. The two tetrahedra *abcd* and *efgh* that we want to join are in boldface and appear repeated in the figure, to better reflect symmetry. The proof finishes by noticing that there is no intermediate tetrahedron that can be reached in two steps from both *abcd* and *efgh*. Hence, five steps are needed to go from one to the other. □

3.4 Many Hirsch-Sharp Polytopes?

Recall that we call a d-polytope (or polyhedron) with n facets *Hirsch-sharp* if its diameter is exactly $n - d$, as happens with the Klee-Walkup polytope of the previous section. Here we describe several ways to construct them.

3.4.1 Trivial Hirsch-Sharp Polytopes

Constructing Hirsch-sharp d-polytopes with any number of facets not exceeding $2d$ is easy. For this reason we call such Hirsch-sharp polytopes *trivial*:

1. *Product.* If P and Q are Hirsch-sharp, then so is their Cartesian product $P \times Q$. Indeed, the dimension, number of facets, and diameters of $P \times Q$ are the sum of

those of P and Q. For the diameter, if we want to go from vertex (u_1, v_1) to vertex (u_2, v_2) we can do so by going from (u_1, v_1) to (u_2, v_1) along $P \times \{v_1\}$ and then to (u_2, v_2) along $\{u_2\} \times Q$; there is no better way.

In particular, any product of *simplices* of any dimension is Hirsch-sharp. The dimension of $\Delta_{i_1} \times \cdots \times \Delta_{i_k}$, where Δ_i denotes the i-simplex, is $\sum_{j=1}^{k} i_j$, its number of facets is $\sum_{j=1}^{k}(i_j + 1)$, and its diameter is k.

Corollary 3.9 *For every $d < n \leq 2d$ there are simple d-polytopes with n facets and diameter $n - d$.*

Proof Let $k = n - d \leq d$ and let $i_1, \ldots, i_k$ be any partition of d into k positive integers (that is, $i_1 + \cdots + i_k = d$). Let $P = \Delta_{i_1} \times \cdots \times \Delta_{i_k}$.   □

2. *Intersection of two affine orthants.* Let $k = n - d \leq d$ and let u be the origin in $\mathbb{R}^d$. Let $v = (1, \ldots, 1, 0, \ldots, 0)$ be the point whose first k coordinates are 1 and whose remaining $d - k$ coordinates are 0.

Consider the polytope P defined by the following $d + k$ inequalities:

$$x_i \geq 0, \quad \forall i; \qquad \psi_j(\mathbf{x}) \geq 0, \quad j = 1, \ldots, k,$$

where the ψ_i are affine linear functionals that vanish at v and are positive at u. No matter what choice we make for the ψ_j's, as long as they are sufficiently generic to make P simple, P will have diameter (at least) k; to go from v to u we need to enter the k facets $x_j = 0$, $j = 1, \ldots, k$, and each step gets you into at most one of them. In principle, P may be an unbounded polyhedron; but if one of the ψ_j's is, say, $k - \sum x_i$, then it will be bounded.

Hirsch-sharp unbounded polyhedra with any number of facets are also easy to obtain: setcounterpropositionn9

****Proposition 3.10** *For every $n \geq d$ there are simple d-polyhedra with n facets and diameter $n - d$.*

### 3.4.2 Non-trivial Hirsch-Sharp Polytopes

****Theorem 3.11** (Fritzsche-Holt-Klee [26, 31, 32]) *Hirsch-sharp d-polytopes with n facets exist in at least the following cases*: (1) $n \leq 3d - 3$; *and* (2) $d \geq 7$.

Both parts are proved using the Klee-Walkup polytope Q_4 as a starting block, from which more complicated polytopes are obtained. In a sense, Q_4 *is the only non-trivial Hirsch-sharp polytope we know*.

The case $n \leq 3d - 3$ was first proved in 1998 [32], and follows from the iterated application of the next lemma to the Klee-Walkup polytope Q_4. Part 2 was proved in [26] for $d \geq 8$ and was improved to $d \geq 7$ in [31]. We sketch its proof in [38].

Lemma 3.12 (Holt-Klee [32]) *If there are Hirsch-sharp d-polytopes with $n > 2d$ facets, then there are also Hirsch-sharp $(d + 1)$-polytopes with $n + 1$, $n + 2$, and $n + 3$ facets.*

Fig. 5 After wedging in a Hirsch-sharp polytope, we can truncate twice
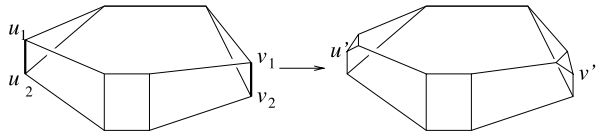
Proof Let u and v be vertices at distance $n - d$ in a simple d-polytope with n-facets. Let F be a facet not containing any of them, which exists since $n > 2d$. When we wedge on F we get two edges u_1u_2 and v_1v_2 with the properties that the distance from any u_i to any v_i is again (at least) d. We can then truncate one or both of u_1 and v_1 to obtain one or two more facets in a polytope that is still Hirsch-sharp. See Fig. 5. □

Hirsch-sharp polytopes of dimensions two and three exist only when $n \leq 2d$ (see Sect. 2.1). Existence of Hirsch-sharp polytopes with many facets in dimensions four to six remains open. We do know that they do not exist in dimension four with 10, 11 or 12 facets (see Theorem 2.2), which may well indicate that Q_4 is the only Hirsch-sharp 4-polytope.

3.5 The Unbounded and Monotone Hirsch Conjectures

In the Hirsch conjecture as we have stated it, we only consider bounded polytopes. However, in the context of linear programming the feasible region may well not be bounded, so the conjecture is equally relevant for *unbounded* polyhedra. In fact, that is how W. Hirsch originally posed the question.

Moreover, for the simplex method in linear programming one follows *monotone paths*: starting at an initial vertex u one does pivot steps (that is, one crosses edges) always increasing the value of the linear functional ϕ to be maximized, until one arrives at a vertex v where no pivot step gives a greater value to ϕ. Convexity then implies that v is the global maximum for ϕ in the feasible region. This raises the question whether a *monotone* variant of the Hirsch conjecture holds: given two vertices u and v of a polyhedron P and a linear functional that attains its maximum on P at v, is there a ϕ-monotone path of edges from u to v whose length is at most $n - d$?

Both the unbounded and monotone variants of the Hirsch conjecture fail, and both proofs use the Klee-Walkup Hirsch-sharp polytope Q_4 described in Sect. 3.3. In fact, knowing the mere existence of such a polytope is enough. The proofs do not use any property of Q_4 other than the fact that it is Hirsch-sharp, simple, and has $n > 2d$. Simplicity is not a real restriction since it can always be obtained without decreasing the diameter (Lemma 1.3). The inequality $n > 2d$, however, is essential.

Theorem 3.14 (Klee-Walkup [42]) *There is a simple unbounded polyhedron $\tilde{Q}_4$ with eight facets and dimension four and whose graph has diameter 5.*

Proof Let Q_4 be the simple Klee-Walkup polytope with nine facets, and let u and v be vertices of Q_4 at distance five from one another. By simplicity, u and v lie in (at most) eight facets in total and there is (at least) one facet F not containing u nor v. Let $\tilde{Q}_4$ be the unbounded polyhedron obtained by a projective transformation

Fig. 6 Disproving the
unbounded Hirsch conjecture
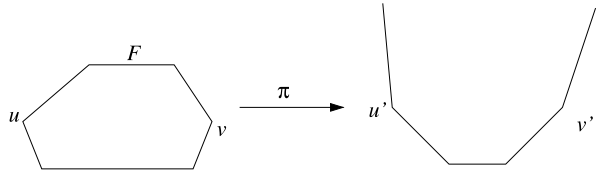
that sends this ninth facet to infinity. The graph of $\tilde{Q}_4$ contains both u and v, and
is a subgraph of that of $\tilde{Q}_4$, hence its diameter is still at least five. See Fig. 6 for a
schematic rendition of this idea.   □

Remark 3.15 The "converse" of the above proof also works: from any non-Hirsch
unbounded 4-polyhedron $\tilde{Q}$ with eight facets, one can build a bounded 4-polytope
with nine facets and diameter five, as follows:

Let u and v be vertices of $\tilde{Q}$ at distance five from one another. Construct the
polytope Q by cutting $\tilde{Q}$ with a hyperplane that leaves all the vertices of $\tilde{Q}$ on the
same side. This adds a new facet and changes the graph, by adding new vertices and
edges on that facet. But u and v will still be at distance five: to go from u to v either
we do not use the new facet F that we created (that is, we stay in the graph of $\tilde{Q}_4$)
or we use a pivot to enter the facet F and at least another four to enter the four facets
containing v.

We now turn to the monotone variant of the Hirsch conjecture:

*****Theorem 3.16** (Todd [56]) *There is a simple* 4-*polytope* P *with eight facets*, *two
vertices* u *and* v *of it*, *and a linear functional* ϕ *such that*:

1. v *is the only maximal vertex for* ϕ.
2. *Any edge-path from* u *to* v *and monotone with respect to* ϕ *has length at least five*.

In both the constructions of Theorems 3.14 and 3.16 one can glue several copies of
the initial block Q_4 to one another, increasing the number of facets by four and the
diameter by five, per Q_4 glued:

Theorem 3.17 (Klee-Walkup, Todd)

1. *There are unbounded* 4-*polyhedra with* $4 + 4k$ *facets and diameter* $5k$, *for every
$k \geq 1$.*
2. *There are bounded* 4-*polyhedra with* $4 + 4k$ *facets and vertices* u *and* v *of them
with the property that any monotone path from* u *to* v *with respect to a certain
linear functional* ϕ *maximized at* v *has length at least* $5k$.

This leaves the following open questions:

- Improve these constructions so as to get the ratio "diameter versus facets" bigger
than 5/4. A ratio bigger than two for the unbounded case would probably yield
counter-examples to the bounded Hirsch conjecture.

- Ziegler [59, p. 87] poses the following *strict* monotone Hirsch conjecture: "for every linear functional ϕ on a d-polytope with n facets there is a ϕ-monotone path of length at most $n - d$ from the minimal to the maximal vertex". Put differently, in the monotone Hirsch conjecture we add the requirement that not only v but also u has a supporting hyperplane where ϕ is constant.

3.6 The Topological Hirsch Conjecture is False

Another natural variant of the Hirsch conjecture is topological. Since (the boundary of) every simplicial d-polytope is a topological triangulation of the $(d - 1)$-dimensional sphere, we can ask whether the simplicial version of the Hirsch conjecture, the one where we walk from simplex to simplex rather than from vertex to vertex, holds for arbitrary triangulations of spheres. The first counterexample to this statement was found by Walkup in 1979 (see [58]), and a simpler one was soon constructed by Walkup and Mani [45].

Both constructions are based on the equivalence of the Hirsch conjecture to the non-revisiting conjecture (Theorem 3.7). The proof of the equivalence is purely combinatorial, so it holds true for topological spheres. Walkup's initial example is a 4-sphere without the non-revisiting property, and Mani and Walkup's is a 3-sphere:

*__Theorem 3.18__ (Mani-Walkup [45]) *There is a triangulated 3-sphere with 16 vertices and without the non-revisiting property. Wedging on it eight times yields an 11-sphere with 24 vertices and with ridge-graph diameter at least 13.*

This triangulated 3-sphere would give a counterexample to the Hirsch conjecture if it were *polytopal*. That is, if it were combinatorially isomorphic to the boundary complex of a four-dimensional simplicial polytope. Altshuler [1] has shown that (for the explicit completion of the subcomplex K given in [45]) this is not the case. As far as we know, it remains an open question to show that *no completion of $K \cup \{abcd, mnop\}$ to the 3-sphere is polytopal*, but we believe that to be the case. Even more strongly, we believe that $K \cup \{abcd, mnop\}$ cannot be embedded in $\mathbb{R}^3$ with linear tetrahedra, a necessary condition for polytopality by the well-known Schlegel construction [59].

As in the monotone and bounded cases, several copies of the construction can be glued to one another. Doing so provides triangulations of the 11-sphere with $12 + 12k$ vertices and diameter at least $13k$, for any k.

Note Added in Proof

On May 10, 2010, while this paper was already in press, the second author has announced a counter-example to the Hirsch Conjecture: there is a 43-dimensional polytope with 86 facets and whose graph has diameter at least 44. See http://gilkalai. wordpress.com/2010/05/10/ for the announcement. By the time this paper is published a preprint will most probably be available at http://www.arxiv.org.

References

1. Altshuler, A.: The Mani-Walkup spherical counterexamples to the W_v-path conjecture are not polytopal. Math. Oper. Res. **10**(1), 158–159 (1985)
2. Altshuler, A., Bokowski, J., Steinberg, L.: The classification of simplicial 3-spheres with nine vertices into polytopes and non-polytopes. Discrete Math. **31**, 115–124 (1980)
3. Balinski, M.L.: The Hirsch conjecture for dual transportation polyhedra. Math. Oper. Res. **9**(4), 629–633 (1984)
4. Barnette, D.: W_v paths on 3-polytopes. J. Comb. Theory **7**, 62–70 (1969)
5. Beck, M., Robins, S.: Computing the Continuous Discretely: Integer-point Enumeration in Polyhedra. Undergraduate Texts in Mathematics. Springer, Berlin (2007)
6. Björner, A., Brenti, F.: Combinatorics of Coxeter Groups. Graduate Texts in Mathematics, vol. 231. Springer, Berlin (2005)
7. Blum, L., Cucker, F., Shub, M., Smale, S.: Complexity and Real Computation. Springer, Berlin (1997)
8. Borgwardt, K.H.: The average number of steps required by the simplex method is polynomial. Z. Oper. Res. **26**, 157–177 (1982)
9. Boyd, S., Vandenberghe, L.: Convex Optimization. Cambridge University Press, Cambridge (2004)
10. Bremner, D., Deza, A., Hua, W., Schewe, L.: More bounds on the diameters of convex polytopes. Preprint, 8 p. (2009). arXiv:0911.4982v1
11. Bremner, D., Schewe, L.: Edge-graph diameter bounds for convex polytopes with few facets. Preprint, 9 p. (2008). arXiv:0809.0915v3
12. Brightwell, G., van den Heuvel, J., Stougie, L.: A linear bound on the diameter of the transportation polytope. Combinatorica **26**(2), 133–139 (2006)
13. Cunningham, W.H.: Theoretical properties of the network simplex method. Math. Oper. Res. **4**, 196–208 (1979)
14. Dantzig, G.B.: Linear Programming and Extensions. Princeton University Press, Princeton (1963)
15. De Loera, J.A., Kim, E.D., Onn, S., Santos, F.: Graphs of transportation polytopes. J. Comb. Theory Ser. A **116**(8), 1306–1325 (2009)
16. De Loera, J.A.: The many aspects of counting lattice points in polytopes. Math. Semesterber. **52**(2), 175–195 (2005)
17. De Loera, J.A., Onn, S.: All rational polytopes are transportation polytopes and all polytopal integer sets are contingency tables. In: Proc. 10th Ann. Math. Prog. Soc. Symp. Integ. Prog. Combin. Optim. (Columbia University, New York, NY, June 2004). Lec. Not. Comp. Sci., vol. 3064, pp. 338–351. Springer, New York (2004)
18. De Loera, J.A., Rambau, J., Santos, F.: Triangulations: Structures for Algorithms and Applications. Algorithms Comput. Math. vol. 25 (to appear), Springer-Verlag
19. Dedieu, J.-P., Malajovich, G., Shub, M.: On the curvature of the central path of linear programming theory. Found. Comput. Math. **5**, 145–171 (2005)
20. Deza, A., Terlaky, T., Zinchenko, Y.: Central path curvature and iteration-complexity for redundant Klee-Minty cubes. Adv. Mech. Math. **17**, 223–256 (2009)
21. Deza, A., Terlaky, T., Zinchenko, Y.: A continuous d-step conjecture for polytopes. Discrete Comput. Geom. **41**, 318–327 (2009)
22. Deza, A., Terlaky, T., Zinchenko, Y.: Polytopes and arrangements: diameter and curvature. Oper. Res. Lett. **36**(2), 215–222 (2008)

23. Dyer, M., Frieze, A.: Random walks, totally unimodular matrices, and a randomised dual simplex algorithm. Math. Program. **64**, 1–16 (1994)
24. Eisenbrand, F., Hähnle, N., Razborov, A., Rothvoß, T.: Diameter of polyhedra: limits of abstraction. Preprint, available at http://people.cs.uchicago.edu/~razborov/research.html. A preliminary version appeared in: Proceedings of the 25th Annual ACM Symposium on Computational Geometry (SoCG'09), ACM, New York (2009)
25. Fomin, S., Zelevinsky, A.: Y-systems and generalized associahedra. Ann. Math. **158**(2), 977–1018 (2003)
26. Fritzsche, K., Holt, F.B.: More polytopes meeting the conjectured Hirsch bound. Discrete Math. **205**, 77–84 (1999)
27. Gawrilow, E., Joswig, M.: Polymake: a framework for analyzing convex polytopes. In: Polytopes—combinatorics and computation (Oberwolfach, 1997). DMV Sem., vol. 29, pp. 43–73. Birkhäuser, Basel (2000)
28. Goldfarb, D., Hao, J.: Polynomial simplex algorithms for the minimum cost network flow problem. Algorithmica **8**, 145–160 (1992)
29. Goodey, P.R.: Some upper bounds for the diameters of convex polytopes. Isr. J. Math. **11**, 380–385 (1972)
30. Hačijan, L.G.: A polynomial algorithm in linear programming. Dokl. Akad. Nauk SSSR **244**(5), 1093–1096 (1979) (in Russian)
31. Holt, F.B.: Blending simple polytopes at faces. Discrete Math. **285**, 141–150 (2004)
32. Holt, F., Klee, V.: Many polytopes meeting the conjectured Hirsch bound. Discrete Comput. Geom. **20**, 1–17 (1998)
33. Hurkens, C.: Personal communication (2007)
34. Kalai, G.: A subexponential randomized simplex algorithm. In: Proceedings of the 24th Annual ACM Symposium on the Theory of Computing, pp. 475–482. ACM, New York (1992)
35. Kalai, G.: Online blog http://gilkalai.wordpress.com. See, for example, http://gilkalai.wordpress.com/2008/12/01/a-diameter-problem-7/, December 2008
36. Kalai, G., Kleitman, D.J.: A quasi-polynomial bound for the diameter of graphs of polyhedra. Bull. Am. Math. Soc. **26**, 315–316 (1992)
37. Karmarkar, N.: A new polynomial time algorithm for linear programming. Combinatorica **4**(4), 373–395 (1984)
38. Kim, E.D., Santos, F.: Companion to "An update on the Hirsch conjecture". Preprint, 22 p. (2009). arXiv:0912.4235v1
39. Klee, V.: Paths on polyhedra II. Pac. J. Math. **17**(2), 249–262 (1966)
40. Klee, V., Kleinschmidt, P.: The d-step conjecture and its relatives. Math. Oper. Res. **12**(4), 718–755 (1987)
41. Klee, V., Minty, G.J.: How good is the simplex algorithm? In: Inequalities, III: Proc. Third Sympos., Univ. California, Los Angeles, Calif., 1969; dedicated to the memory of Theodore S. Motzkin, pp. 159–175. Academic Press, New York (1972)
42. Klee, V., Walkup, D.W.: The d-step conjecture for polyhedra of dimension $d < 6$. Acta Math. **133**, 53–78 (1967)
43. Kleinschmidt, P., Onn, S.: On the diameter of convex polytopes. Discrete Math. **102**(1), 75–77 (1992)
44. Larman, D.G.: Paths of polytopes. Proc. Lond. Math. Soc. **20**(3), 161–178 (1970)
45. Mani, P., Walkup, D.W.: A 3-sphere counterexample to the W_v-path conjecture. Math. Oper. Res. **5**(4), 595–598 (1980)
46. Matoušek, J., Sharir, M., Welzl, E.: A subexponential bound for linear programming. In: Proceedings of the 8th Annual Symposium on Computational Geometry, pp. 1–8 (1992)
47. Megiddo, N.: Linear programming in linear time when the dimension is fixed. J. Assoc. Comput. Mach. **31**(1), 114–127 (1984)
48. Megiddo, N.: On the complexity of linear programming. In: Bewley, T. (ed.) Advances in Economic Theory: Fifth World Congress, pp. 225–268. Cambridge University Press, Cambridge (1987)
49. Naddef, D.: The Hirsch conjecture is true for (0, 1)-polytopes. Math. Program. **45**, 109–110 (1989)
50. Oda, T.: Convex Bodies and Algebraic Geometry. Springer, Berlin (1988)
51. Orlin, J.B.: A polynomial time primal network simplex algorithm for minimum cost flows. Math. Program. **78**, 109–129 (1997)
52. Renegar, J.: A Mathematical View of Interior-Point Methods in Convex Optimization. SIAM, Philadelphia (2001)
53. Smale, S.: On the average number of steps of the simplex method of linear programming. Math. Program. **27**, 241–262 (1983)

54. Smale, S.: Mathematical problems for the next century. In: Mathematics: Frontiers and Perspectives, pp. 271–294. American Mathematics Society, Providence (2000)
55. Spielman, D.A., Teng, S.: Smoothed analysis of algorithms: why the simplex algorithm usually takes polynomial time. J. ACM **51**(3), 385–463 (2004)
56. Todd, M.J.: The monotonic bounded Hirsch conjecture is false for dimension at least 4. Math. Oper. Res. **5**(4), 599–601 (1980)
57. Vershynin, R.: Beyond Hirsch conjecture: walks on random polytopes and smoothed complexity of the simplex method. In: IEEE Symposium on Foundations of Computer Science. NN, vol. 47, pp. 133–142. IEEE, New York (2006)
58. Walkup, D.W.: The Hirsch conjecture fails for triangulated 27-spheres. Math. Oper. Res. **3**, 224–230 (1978)
59. Ziegler, G.M.: Lectures on Polytopes. Graduate Texts in Mathematics, vol. 152. Springer, Berlin (1995)
60. Ziegler, G.M.: Face numbers of 4-polytopes and 3-spheres. In: Proceedings of the International Congress of Mathematicians, vol. III, Beijing, 2002, pp. 625–634. Higher Ed. Press, Beijing (2002)

Edward D. Kim is finishing his Ph.D. studies at the University of California, Davis under the supervision of Jesús A. De Loera. His research is in combinatorial optimization, convex polytopes, Lagrangian matroids, and multi-objective optimization. He is about to start his post-doctoral career with positions at Technische Univ. Delft in the Netherlands and at POSTECH (Pohang Univ. of Science and Technology) in Korea.

Francisco Santos is full professor of mathematics at the Universidad de Cantabria in Spain. He received the Young Researcher award from the Universidad Complutense de Madrid in 2003 and was an invited speaker in the Combinatorics Section of the International Congress of Mathematicians in 2006. He is a specialist in Geometric Combinatorics, especially in the theory of polytopes and triangulations. He is an editor of the journal "Discrete and Computational Geometry".