

Edge-Graph Diameter Bounds for Convex Polytopes with Few Facets

David Bremner and Lars Schewe

CONTENTS

- 1. Introduction
- 2. Path Complexes
- 3. Generation of Oriented Matroids Using SAT Solvers
- 4. Shortcuts
- 5. The Case $\Delta(4, 11)$
- 6. Conclusion
- Note Added in Proof
- Acknowledgments
- References

We show that the edge graph of a 6-dimensional polytope with 12 facets has diameter at most 6, thus verifying the d -step conjecture of Klee and Walkup in the case $d = 6$. This implies that for all pairs (d, n) with $n - d \leq 6$, the diameter of the edge graph of a d -polytope with n facets is bounded by 6, which proves the Hirsch conjecture for all $n - d \leq 6$. We prove this result by establishing this bound for a more general structure, so-called matroid polytopes, by reduction to a small number of satisfiability problems.

1. INTRODUCTION

It is a longstanding open problem to determine the maximal diameter $\Delta(d, n)$ of a d -dimensional polytope with n facets. Not much is known even in small dimensions. The *Hirsch conjecture* states that $\Delta(d, n) \leq n - d$. The special case $\Delta(d, 2d) \leq d$ is known as the *d -step conjecture*. Proving the d -step conjecture for a fixed d implies (by an argument in [Klee and Walkup 67]) that the Hirsch conjecture holds for all pairs (n', d') with $n' - d' = d$. To date, the d -step conjecture has been proved for all $d \leq 5$ by Klee and Walkup.

We show that the d -step conjecture is true in dimension 6 as well. We derive this result by considering a more general class of objects, namely *matroid polytopes*, i.e., oriented matroids, which, if realizable, correspond to convex polytopes. We show that no 6-dimensional matroid polytope with 12 vertices and a shortcut-free facet path of length 7 exists. Then $\Delta(6, 12) = 6$ follows by considering polarity and the already known bounds.

To show that $\Delta(6, 12) \leq 6$, we first give combinatorial conditions for matroid polytopes that violate this bound. This is achieved through the study of path complexes (cf. Section 2). We then show that these conditions cannot be satisfied by an oriented matroid. To prove this, we use a satisfiability solver to produce the desired contradiction (see Section 3). We will use the same method to show that $\Delta(4, 11) = 6$, which settles another special case of

d	$n - d$			
	4	5	6	7
4	4	5	5	{6, 7}
5	<i>4</i>	5	6	[7, 9]
6	<i>4</i>	5	{6, 7}	[7, 9]
7	<i>4</i>	5	{6, 7}	[7, 10]

TABLE 1. Bounds on $\Delta(d, n)$ known before the work described in this article. The values printed in italics follow from observation (iv) of Lemma 1.1.

the Hirsch conjecture. The latter result allows us also to improve the upper bound on $\Delta(5, 12)$ from 9 to 8.

For small parameters there are known general bounds that allow us to compute or at least bound the diameter of polytopes. We summarize them in the following lemma. For an overview of the known bounds, we refer to the books [Grünbaum 03, Ziegler 95] and the survey [Klee and Kleinschmidt 87].

Lemma 1.1. [Klee 64, Klee and Walkup 67, Holt 04] *The following relations hold for the maximal diameter $\Delta(d, n)$ of a d -polytope with n facets:*

- (i) $\Delta(3, n) = \lfloor \frac{2}{3}n \rfloor - 1$.
- (ii) $\Delta(d, 2d + k) \leq \Delta(d - 1, 2d + k - 1) + \lfloor \frac{k}{2} \rfloor + 1$ for all d and $k = 0, 1, 2, 3$.
- (iii) $\Delta(d, n) \leq \Delta(n - d, 2(n - d))$ for all (d, n) .
- (iv) $\Delta(d, n) = \Delta(n - d, 2(n - d))$ for all (d, n) with $n \leq 2d$.
- (v) $\Delta(d, n) \geq n - d$ for all $n > d \geq 7$

Apart from these general results, some special cases have been solved in [Goodey 72]:

Lemma 1.2. [Goodey 72] *The following relations hold for the maximal diameter $\Delta(d, n)$ of a d -polytope with n facets:*

- (i) $\Delta(4, 10) = 5$ and $\Delta(5, 11) = 6$.
- (ii) $\Delta(6, 13) \leq 9$ and $\Delta(7, 14) \leq 10$.

Table 1 summarizes the bounds on $\Delta(d, n)$ that follow from Lemmas 1.1 and 1.2. The values printed in italics follow from observation (iv) of Lemma 1.1. The results in this article make it possible to strengthen the bounds; we give an overview in Table 2.

It is not difficult to see (e.g., by a perturbation argument) that $\Delta(d, n)$ is always attained by a simple poly-

d	$n - d$			
	4	5	6	7
4	4	5	5	6
5	4	5	6	{7, 8}
6	4	5	6	[7, 9]
7	4	5	6	[7, 10]

TABLE 2. Summary of bounds for $\Delta(d, n)$. The bold entries are from the computations described in this article.

tope. Thus, it is sufficient for our purposes to restrict our attention to this class of polytopes. It will also be useful to investigate the problem in a polar setting. Thus, we will be looking at d -dimensional simplicial polytopes with n vertices. In this setting $\Delta(d, n)$ is just the maximal length of a shortest facet path in the polytope.

The rest of this article is organized as follows: We will first explain the notion of path complexes (combinatorial generalizations of facet paths). It will turn out that in the cases we are interested in, there are many fewer relevant types of path complexes than of simplicial polytopes. We then outline the method that allows us to transform each of the resulting (non)realizability problems into an (un)satisfiability problem. The generation of a particular set of constraints—the forbidden shortcuts constraints—is derived in a dedicated section. We conclude with a discussion of possible extensions of our approach.

2. PATH COMPLEXES

In this section we outline a combinatorial model for the facet paths of simplicial polytopes: path complexes. Since we are proving our main result by making a case distinction whereby each candidate path complex needs to be investigated, it is important to restrict the number of cases we need to deal with. We adopt the terminology of [Bremner et al. 05]. We also mention several propositions from the same source; for the most part we repeat proofs here to make this article self-contained.

Recall that in a pure simplicial complex of dimension $d - 1$, the 0, $d - 1$, and $d - 2$ simplices are called respectively *vertices*, *facets*, and *ridges*. The generalization of a facet path in this setting is a *path complex*: a pure simplicial complex whose dual graph (with *facets* as nodes and *ridges* shared by two facets as edges) is a path.

We will use the path complex given in Figure 1 as a running example.

It is known [Klee and Walkup 67] that when $n \geq 2d$, the maximum diameter of an (n, d) -polytope is always re-




FIGURE 1. Example of a path complex; canonical representation $\langle 1, 2, 1, 3, 1 \rangle$; pivot sequence $(1, 4), (2, 5), (4, 6), (3, 7), (6, 8)$; restricted growth string $1, 1, 2, 1$.

alized by some *end-disjoint* path (i.e., one whose end vertices do not share a facet). We thus assume that the start and end facets of all path complexes considered here are vertex disjoint; we sometimes call such complexes *end-disjoint path complexes* to emphasize this feature. It will be convenient to discuss *directed* path complexes with a fixed ordering on the facets. Each undirected path complex corresponds to at most two directed path complexes.

We fix notation for path complexes as follows. Let \mathcal{F} be a $(d - 1)$ -dimensional directed path complex with facets F_0, \dots, F_k . Without loss of generality, we fix $F_0 = \{1, \dots, d\}$.

From the definition of a path complex, we know that F_j , $j > 0$, can be obtained from F_{j-1} by a *pivot* (l_j, e_j) , where $F_j = F_{j-1} \setminus \{l_j\} \cup \{e_j\}$. A path complex is thus encoded by a *pivot sequence* $(l_1, e_1), (l_2, e_2), \dots, (l_k, e_k)$. The pivot sequence of the path complex in Figure 1, for instance, is given by $(1, 4), (2, 5), (4, 6), (3, 7), (6, 8)$.

A *revisit* occurs when a vertex leaves F_p and reenters in F_q . A *nonrevisiting* path complex is thus one in which no entering vertex e_j is equal to leaving vertex l_k ; equivalently, $|\cup_{j=0}^k F_j| = d + k$.

The path complex in Figure 1 is nonrevisiting. An example of a path complex with one revisit can be found in Figure 2.

In the case of a nonrevisiting path, we may give an alternative representation for the pivot sequence that

is more convenient for computational purposes. Write F_0, \dots, F_k as the rows of a $d \times (k + 1)$ table so that l_j and e_j are in the same column. We call the list of column indices corresponding to pivots an *index sequence* for the path complex. A basic property of index sequences is that successive indices are distinct, since a repeated index would cause three successive facets to intersect in a ridge, which is impossible in a path complex (cf. [Bremner et al. 05, Section 2.1]).

In our example (Figure 1), we have written the rows of the table directly in each facet of the complex; the column that is affected by the pivot is underlined. In this case, we can read off the index sequence as $\langle 1, 2, 1, 3, 1 \rangle$.

Many path complexes with distinct pivot sequences are actually symmetric copies of each other (i.e., isomorphic as simplicial complexes). We thus need a method to remove symmetric copies of a path complex. One way to do so would be to associate a colored graph with each path complex and remove copies when these graphs are isomorphic. Another way is to use explicit rules to identify redundant complexes. We have implemented the first way (using the graph isomorphism tool nauty [McKay 05]). We will, however, describe only the second approach in this article: in the cases described here, the number of path complexes is small enough to check the reduction directly using the relevant rules (see the discussion in the following paragraph and Lemmas 2.2 and 5.1). We used the graph-isomorphism variant, however, to check the correctness of the resulting list of path complexes.

There are at least two different kinds of symmetry of a path complex to be considered. In the first case, we may relabel the vertices of the initial simplex in $d!$ ways. This symmetry can be removed by insisting on a particular labeling. We call an index sequence whose column indices occur in order, i.e., the vertex of F_0 in column c leaves before the vertex in column $c + 1$, a *canonical index sequence*.

The second kind of symmetry to be considered is the choice of initial facet. In general, the two choices may lead to different canonical index sequences. For nonrevisiting paths, we simply keep the lexicographically smaller canonical index sequence.

To give an example, we have provided the canonical index sequences on three symbols of length 5 in Table 3.

If two nonrevisiting directed path complexes (with the same dimension and number of elements) have the same index sequence, then a compatible (i.e., inclusion-preserving) labeling of the two complexes can be constructed inductively as follows. Start with a common la-

Canonical	Reverse	Complex
$\langle 1, 2, 1, 2, 3 \rangle$	$\langle 1, 2, 3, 2, 3 \rangle$	$\{\{1, 2, 3\}, \{2, 3, 4\}, \{3, 4, 5\}, \{3, 5, 6\}, \{3, 6, 7\}, \{6, 7, 8\}\}$
$\langle 1, 2, 1, 3, 1 \rangle$	$\langle 1, 2, 1, 3, 1 \rangle$	$\{\{1, 2, 3\}, \{2, 3, 4\}, \{3, 4, 5\}, \{3, 5, 6\}, \{5, 6, 7\}, \{5, 7, 8\}\}$
$\langle 1, 2, 1, 3, 2 \rangle$	$\langle 1, 2, 3, 1, 3 \rangle$	$\{\{1, 2, 3\}, \{2, 3, 4\}, \{3, 4, 5\}, \{3, 5, 6\}, \{5, 6, 7\}, \{6, 7, 8\}\}$
$\langle 1, 2, 3, 1, 2 \rangle$	$\langle 1, 2, 3, 1, 2 \rangle$	$\{\{1, 2, 3\}, \{2, 3, 4\}, \{3, 4, 5\}, \{4, 5, 6\}, \{5, 6, 7\}, \{6, 7, 8\}\}$
$\langle 1, 2, 3, 2, 1 \rangle$	$\langle 1, 2, 3, 2, 1 \rangle$	$\{\{1, 2, 3\}, \{2, 3, 4\}, \{3, 4, 5\}, \{4, 5, 6\}, \{4, 6, 7\}, \{6, 7, 8\}\}$

TABLE 3. Canonical index sequences on three symbols of length 5 and the corresponding simplicial complexes. The second column lists noncanonical index sequences for the same complex.

bel set for the initial facets. For each vertex that enters according to the index sequence, choose a previously unused label.

On the other hand, if we have two isomorphic directed path complexes, then a compatible labeling of vertices labels the pivots in the same way. Consider the table representation of the two (compatibly labeled) directed path complexes (with columns canonically ordered by first pivot). At least the first d rows of the two tables are identical. Let j be the first step at which the index sequences differ. Since the first j rows of the two tables are identical, the two leaving labels are distinct, a contradiction.

We will next sketch the enumeration of all paths with at most one revisit as derived in [Bremner et al. 05]. The authors’ first result concerns the number of directed non-revisiting d -paths of length l . They show that this number can be expressed as a Stirling number of the second kind; we denote these numbers by $\left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\}$. The basic recurrence of the Stirling numbers can then be used to give a recursive algorithm for generating these paths. The proof uses one intermediate structure: *restricted growth strings*, i.e., k -ary strings whose symbols occur in order, and all k symbols occur. Or put more formally: a restricted growth string is a sequence e_1, \dots, e_n of symbols from $\{1, \dots, k\}$ that starts with $e_1 = 1$ and has the property that for each element $e_j, j > 1$, the property $e_j = l$ holds only if there exists an element e_i with $i < j$ such that $e_i = l - 1$.

Lemma 2.1. [Bremner et al. 05] *The number of directed nonrevisiting d -paths of length l is $\left\{ \begin{smallmatrix} l-1 \\ d-1 \end{smallmatrix} \right\}$.*

Sketch of the proof. We first argue that there is a bijection between these directed nonrevisiting d -paths of length l and restricted growth strings on $d - 1$ symbols of length $l - 1$. The bijection between directed facet paths and index sequences was discussed above. Here we consider the correspondence between index sequences and restricted growth strings. Given a canonical index sequence p_1, \dots, p_l , we can output a restricted growth

string s_1, \dots, s_{l-1} by setting $s_1 = 1$ and s_{j-1} to the rank of p_j in $\{1, \dots, d\} \setminus p_{j-1}$ for $j > 2$. This transformation is evidently a bijection.

The Stirling numbers of the second kind $\left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\}$ count the number of partitions of an n -element set into k nonempty parts. The bijection between these partitions and restricted growth strings of length n on k elements can be given as follows. Let $\pi = (\pi_i)_{i=1}^k$ be a partition of the n -element set with k parts. We may assume that the k parts are ordered according to their minimal elements, i.e., $\min \pi_1 = 1$ and $\min \pi_i \leq \min \pi_{i+1}$ for all $i \in \{1, \dots, k - 1\}$.

Then we can construct a restricted growth string $(e_i)_{i=1}^n$ of length n on k elements by setting $e_i = j$ if $i \in \pi_j$. We omit the verification that this mapping is a bijection. \square

In the case $l = d + 1$, the number of canonical index sequences is $\left\{ \begin{smallmatrix} d \\ d-1 \end{smallmatrix} \right\} = \binom{d}{2}$, since the partition of a d -set into $d - 1$ parts is determined by what pair of elements are grouped together. More directly, the canonical index sequences are determined by the $\binom{d+1}{2} - d = \binom{d}{2}$ choices of nonadjacent pairs of equal indices.

Single-revisit paths are generated from nonrevisiting paths on one more vertex by identifying two vertices. We represent this by partitioning the index sequence into three possibly empty parts: the *prefix*, the *loop*, and the *suffix*. The loop represents the actual revisit, where the first pivot is the vertex in question leaving the facet and the last pivot in the loop is the vertex returning to the facet.

The following lemma gives two necessary conditions for identifications of vertices of path complexes; thus we do not omit any valid path complexes by pruning according to them.

Lemma 2.2. [Bremner et al. 05] *Let $P|L|S$ be an index sequence of a nonrevisiting path. Let us identify the first and last elements of L . Then the following conditions are necessary for the resulting complex to be an end-disjoint path complex:*




FIGURE 2. Example of a single-revisit path complex.

- (i) The loop L must contain at least three distinct symbols.
- (ii) Either the first symbol of L must appear in P or the last symbol of L must appear in S .

Proof. The first condition prevents the creation of a new ridge, which would violate the condition that the dual graph is a path. The second condition makes sure that the identified vertex is not on both the first and last facets, which would violate end-disjointness. Note that both disjuncts can hold, which just means that the identified vertex is on neither the first nor the last facet. \square

For revisiting paths, rather than keeping the lexicographically smaller canonical index sequence, it seems to be better to use a different symmetry-breaking strategy that uses the second condition of Lemma 2.2.

Lemma 2.3. [Bremner et al. 05] *Every combinatorial type of end-disjoint single-revisit path has an encoding as an index sequence without a revisit on the first facet.*

Proof. Consider a path complex with an index sequence π with an identification in the first facet. Since the path complex is end-disjoint, π must not have an identification in the last facet. Thus, the reverse index sequence has no revisit on the first facet. Since both of these sequences describe the same combinatorial type, we may choose the latter. \square

As with Lemma 2.2, it is clear that the condition of (the proof of) Lemma 2.3 is necessary, and we lose no combinatorial types of path complexes in filtering by it. Note that in general, our two symmetry-breaking strategies for choice of an initial facet are incompatible, and we must choose one. In the case of $l = d + 1$, insisting that the first symbol of L be contained in P , along with the restriction $|L| \geq 3$, yields that the repeated symbol must

1	2	1	3	4	5	6
1	2	3	1	4	5	6
1	2	3	4	1	5	6
1	2	3	2	4	5	6
1	2	3	4	2	5	6
1	2	3	4	2	5	6

TABLE 4. Permissible canonical index sequences for a single revisit, the case $l = 7, d = 6$.

1	2	(1	3	4)	5	6
1	2	(1	3	4	5)	6
1	2	(1	3	4	5	6)
1	2	3	(1	4	5)	6
1	2	3	(1	4	5	6)
1	2	3	4	(1	5	6)
1	2	3	(2	4	5)	6
1	2	3	(2	4	5	6)
1	2	3	4	(2	5	6)
1	2	3	4	(3	5	6)

TABLE 5. Canonical index sequences (with loops indicated) for $l = 7, d = 6$.

be within the first $d + 1$ positions. Applying this restriction for $d = 6$, we derive the canonical index sequences in Table 4.

Since the start of L is fixed at the first repeated index, the remaining choice is the end of L . Making loops of length at least 3 yields the sequences of Table 5.

In Table 6 we have listed the pivot sequences for the canonical index sequences of Table 5, i.e., possible path complexes of length 7 for polytopes of dimension 6 on 12 vertices according to the results so far. We summarize our discussion (and the previous bounds on $\Delta(6, 12)$) in the following proposition. We will give an algorithmic way to prove that the necessary condition outlined in the proposition is true. This is the topic of the next section.

Proposition 2.4. *If none of the path complexes in Table 6 can be realized as a matroid polytope, then $\Delta(6, 12) = 6$ holds.*

3. GENERATION OF ORIENTED MATROIDS USING SAT SOLVERS

To show that a given path complex cannot be completed to a simplicial polytope, we prove the stronger statement that it cannot be completed to a matroid polytope, i.e., there exists no oriented matroid with the given path complex in its boundary.

(1, 7)	(2, 8)	(7, 9)	(3, 10)	(4, 7)	(5, 11)	(6, 12)
(1, 7)	(2, 8)	(7, 9)	(3, 10)	(4, 11)	(5, 7)	(6, 12)
(1, 7)	(2, 8)	(7, 9)	(3, 10)	(4, 11)	(5, 12)	(6, 7)
(1, 7)	(2, 8)	(3, 9)	(7, 10)	(4, 11)	(5, 7)	(6, 12)
(1, 7)	(2, 8)	(3, 9)	(7, 10)	(4, 11)	(5, 12)	(6, 7)
(1, 7)	(2, 8)	(3, 9)	(4, 10)	(7, 11)	(5, 12)	(6, 7)
(1, 7)	(2, 8)	(3, 9)	(8, 10)	(4, 11)	(5, 8)	(6, 12)
(1, 7)	(2, 8)	(3, 9)	(8, 10)	(4, 11)	(5, 12)	(6, 8)
(1, 7)	(2, 8)	(3, 9)	(4, 10)	(8, 11)	(5, 12)	(6, 8)
(1, 7)	(2, 8)	(3, 9)	(4, 10)	(9, 11)	(5, 12)	(6, 9)

TABLE 6. Possible pivot sequences in the (6, 12) case.

This section is devoted to defining oriented matroids and explaining how the nonexistence of certain oriented matroids (and consequently the nonexistence of certain point sets) can be proved using SAT solvers.

The use of SAT solvers to generate oriented matroids was first described in [Schewe 07] and [Schewe 10]. The method used there is the basis for the approach outlined in this section. In our computations we used the SAT solver Minisat [Eén and Sörensson 03].

Oriented matroids have been used before to treat diameter questions of polytopes; one reference that is particularly interesting is the thesis [Schuchert 95], where among other things, the author confirms that $\Delta(4, 11) = 6$ in the special case of neighborly matroid polytopes.

Oriented matroids are a combinatorial abstraction of point configurations in \mathbb{R}^d . We will use the chirotope axioms of oriented matroids in the sequel; for further axiom systems and proofs of equivalence we refer to [Björner et al. 99, Chapter 3].

Since we are dealing only with simplicial polytopes, we may always assume that our oriented matroids are uniform, i.e., $\chi(b) \neq 0$ for all $(d + 1)$ -sets b . This further simplifies the axioms, so that we need to check the following axioms.

Definition 3.1. Let $E = \{1, \dots, n\}$, $r \in \mathbb{N}$, and $\chi : E^r \rightarrow \{-1, +1\}$. We call $\mathcal{M} = (E, \chi)$ a *uniform oriented matroid of rank r* if the following conditions are satisfied:

- (B1) The mapping χ is alternating.
- (B2) For all $\sigma \in \binom{E}{r-2}$ and all subsets $\{x_1, \dots, x_4\} \subseteq E \setminus \sigma$,

$$\{\chi(\sigma, x_1, x_2)\chi(\sigma, x_3, x_4), -\chi(\sigma, x_1, x_3)\chi(\sigma, x_2, x_4), \chi(\sigma, x_1, x_4)\chi(\sigma, x_2, x_3)\} = \{-1, +1\}.$$

These relations can be seen as abstractions of the Grassmann–Plücker relations on determinants [Björner et al. 99].

We also need to express the fact that the path complex we are given is in the boundary of the oriented matroid. The fact that an ordered d -set F is a facet of the matroid polytope can be expressed by enforcing that $\chi(F, e)$ have the same sign for all $e \in E \setminus F$. We call an oriented matroid every element of which is contained in a facet a *matroid polytope*. Since the chirotope axioms are invariant under negation, we may always assume that the sign of one base is positive. Using the fact that the facet incidence graph of the path complex is connected, we can infer the signs of the other bases that contain a facet of the path complex.

In the case $n = 2d$ (we consider more-general situations below), end-disjointness implies that every point in the oriented matroid must be on some facet of the input path complex.

This property implies that we do not need to add additional constraints enforcing convexity; the points of the oriented matroid are already in convex position. To make sure that the starting path complex is actually a geodesic path in the boundary complex of the oriented matroid, we need to forbid “shortcuts,” i.e., shorter paths that connect the end facets of our starting path complex. To enforce this, we need to ensure that for each such possible shortcut, at least one facet is missing. The generation of these shortcuts is the subject of the next section.

As mentioned above, we use the approach of [Schewe 07, Schewe 10] to transform these conditions into an instance of SAT. We introduce variables $[b]$ for each r -tuple b . The interpretation of these variables is that in a satisfying assignment, $[b]$ should be true if $\chi(b) = +$. It follows from condition (B1) that we will need only $\binom{n}{d}$ of these variables in the final SAT instance.

It turns out that axiom (B2) yields $16\binom{n}{r-2}\binom{n-r+2}{4}$ CNF constraints, which are shown in Table 7. For details we refer to [Schewe 07, Schewe 10].

It remains to explain how to encode the facet path and the forbidden shortcuts. Given a d -tuple F and an ordering x_1, \dots, x_{n-d} of the set $X = \{1, \dots, n\} \setminus F$, to enforce that F is a facet we need to add the following clauses:

$$\left(\bigwedge_{i=1}^{n-d-1} [F, x_i] \vee \neg[F, x_{i+1}] \right) \wedge \left(\bigwedge_{i=1}^{n-d-1} \neg[F, x_i] \vee [F, x_{i+1}] \right).$$

$$\begin{aligned}
 GP(\alpha, \beta, \gamma, \delta, \epsilon, \zeta) &= (\neg[\alpha] \vee \neg[\beta] \vee \neg[\gamma] \vee [\delta] \vee \neg[\epsilon] \vee \neg[\zeta]) \\
 &\wedge (\neg[\alpha] \vee \neg[\beta] \vee \neg[\gamma] \vee [\delta] \vee [\epsilon] \vee [\zeta]) \\
 &\wedge (\neg[\alpha] \vee \neg[\beta] \vee [\gamma] \vee \neg[\delta] \vee \neg[\epsilon] \vee \neg[\zeta]) \\
 &\wedge (\neg[\alpha] \vee \neg[\beta] \vee [\gamma] \vee \neg[\delta] \vee [\epsilon] \vee [\zeta]) \\
 &\wedge (\neg[\alpha] \vee [\beta] \vee \neg[\gamma] \vee \neg[\delta] \vee \neg[\epsilon] \vee [\zeta]) \\
 &\wedge (\neg[\alpha] \vee [\beta] \vee \neg[\gamma] \vee \neg[\delta] \vee [\epsilon] \vee \neg[\zeta]) \\
 &\wedge (\neg[\alpha] \vee [\beta] \vee [\gamma] \vee [\delta] \vee \neg[\epsilon] \vee [\zeta]) \\
 &\wedge (\neg[\alpha] \vee [\beta] \vee [\gamma] \vee [\delta] \vee [\epsilon] \vee \neg[\zeta]) \\
 &\wedge ([\alpha] \vee \neg[\beta] \vee \neg[\gamma] \vee \neg[\delta] \vee \neg[\epsilon] \vee [\zeta]) \\
 &\wedge ([\alpha] \vee \neg[\beta] \vee \neg[\gamma] \vee \neg[\delta] \vee [\epsilon] \vee \neg[\zeta]) \\
 &\wedge ([\alpha] \vee \neg[\beta] \vee [\gamma] \vee [\delta] \vee \neg[\epsilon] \vee [\zeta]) \\
 &\wedge ([\alpha] \vee \neg[\beta] \vee [\gamma] \vee [\delta] \vee [\epsilon] \vee \neg[\zeta]) \\
 &\wedge ([\alpha] \vee [\beta] \vee \neg[\gamma] \vee [\delta] \vee \neg[\epsilon] \vee \neg[\zeta]) \\
 &\wedge ([\alpha] \vee [\beta] \vee \neg[\gamma] \vee [\delta] \vee [\epsilon] \vee [\zeta]) \\
 &\wedge ([\alpha] \vee [\beta] \vee [\gamma] \vee \neg[\delta] \vee \neg[\epsilon] \vee \neg[\zeta]) \\
 &\wedge ([\alpha] \vee [\beta] \vee [\gamma] \vee \neg[\delta] \vee [\epsilon] \vee [\zeta])
 \end{aligned}$$

TABLE 7. CNF constraints corresponding to a given σ , x_1, \dots, x_4 in axiom (B2). Each of α, \dots, ζ is a $(d+1)$ -set of indices. $[\alpha] \dots [\zeta]$ are the corresponding SAT variables.

Here we use the convention that for a d -tuple $F = (f_1, \dots, f_d)$, the variable $[F, x]$ denotes the variable $[(f_1, \dots, f_d, x)]$.

In order to enforce that some $F \in \mathcal{F}$ is not a facet, we first construct constraints implied by all $F \in \mathcal{F}$ being on the boundary, then negate them. Let $\tau(Y) = (-1)^k$, where k transpositions are required to sort tuple Y .

Lemma 3.2. *Let $\mathcal{F} = \{F_1, F_2, \dots, F_m\}$ be a path complex. For $i > 1$, let $e_i = F_i \setminus F_{i-1}$, $l_i = F_{i-1} \setminus F_i$. Let $\sigma_1 = 1$, and for $i > 1$, let $\sigma_i = \tau(F_{i-1}, e_i)\tau(F_i, l_i)\sigma_{i-1}$. If F_i and F_{i-1} are both facets, then*

$$\sigma_i \chi(F_i, x) = \sigma_{i-1} \chi(F_{i-1}, y), \quad x \notin F_i, y \notin F_{i-1}.$$

Proof. Suppose that F_{i-1} and F_i are both facets, and let $x \notin F_i$, $y \notin F_{i-1}$, $T_i = F_i \cup F_{i-1}$. Then

$$\begin{aligned}
 \chi(F_{i-1}, e_i)\chi(F_{i-1}, y) &= \chi(F_i, l_i)\chi(F_i, x) = 1, \\
 \chi(F_{i-1}, e_i)\chi(F_i, l_i) &= \tau(F_{i-1}, e_i)\tau(F_i, l_i)\chi(T_i)^2 \\
 &= \tau(F_{i-1}, e_i)\tau(F_i, l_i).
 \end{aligned}$$

It follows that

$$\tau(F_{i-1}, e_i)\tau(F_i, l_i)\chi(F_i, x) = \chi(F_{i-1}, y). \quad (1)$$

The lemma now follows from the definition of σ_i and (1):

$$\begin{aligned}
 \sigma_i \chi(F_i, x) &= \sigma_{i-1} \tau(F_{i-1}, e_i)\tau(F_i, l_i)\chi(F_i, x) \\
 &= \sigma_{i-1} \chi(F_{i-1}, y).
 \end{aligned}$$

□

From Lemma 3.2, it follows that to force path complex $\mathcal{F} = \{F_1, F_2, F_3, \dots, F_m\}$ not to be entirely on the boundary, it suffices that

$$\{\sigma_j \chi(F_j, x) \mid 1 \leq j \leq m, x \notin F_j\} = \{+1, -1\}.$$

Let $z_i(x)$ be the CNF literal corresponding to $\sigma_i \chi(F_i, x)$. Then the two corresponding CNF constraints are

$$\left(\bigvee_{\substack{i \in \{2, \dots, |\mathcal{F}|-1\} \\ x \in \{1, \dots, n\} \setminus F_i}} z_i(x) \right) \wedge \left(\bigvee_{\substack{i \in \{2, \dots, |\mathcal{F}|-1\} \\ x \in \{1, \dots, n\} \setminus F_i}} \neg z_i(x) \right).$$

In our setting, we may assume that the first and last elements of \mathcal{F} are facets.

We get a large number of these clauses, and each of them—on its own—is quite weak (since it contains many literals). However, taken together, they lead to the desired contradiction. We note that since we are looking for a contradiction, it is not necessary to generate all of these clauses.

4. SHORTCUTS

In the context of testing a k -step path complex $\Delta = F_0, \dots, F_k$ for geodesic (non)realizability, we call any path complex (on the same set of vertices) from F_0 to F_k with fewer than k pivots a *shortcut*. Our general scheme in establishing that a given path complex is not geodesically realizable is to find a set of shortcuts S and show that no matroid polytope (and hence no convex polytope) can contain Δ but no element of S .

A path $\pi = v_0, v_1, \dots, v_k$ in a graph G is called *inclusion-minimal* if no proper subset of $\{v_0, \dots, v_k\}$ is a path from v_0 to v_k . Every inclusion-minimal path is evidently simple, and every geodesic (shortest path) is inclusion-minimal.

Lemma 4.1. *The inclusion-minimal (s, t) -paths of length k in a graph G are exactly π, t , where π is an inclusion-minimal (s, t') -path of length $k-1$, t' is a neighbor of t , and no other $v \in \pi$ is adjacent to t .*

Proof. Let $\pi = v_0, \dots, v_k$ be an inclusion-minimal path of length k . The path $\pi' = v_0, \dots, v_{k-1}$ must be inclusion-minimal, for otherwise, π would also not be inclusion-minimal. Similarly, if v_k is adjacent to some v_j , $j < k-1$, then π is not inclusion-minimal.

Suppose, on the other hand, we have an inclusion-minimal path $\pi' = v_0, \dots, v_{k-1}$ such that v_k is

adjacent to v_{k-1} , but not to any other vertex in π' . Then $\pi = v_0, \dots, v_k$ is a path from v_0 to v_k , and if this path is not inclusion-minimal, then a shortcut must exist in π' , which is a contradiction. \square

From this lemma we can derive an algorithm that generates all the potential shortcuts for the given path complex, taking as the graph G the so-called *pivot graph*, whose nodes are d -sets and whose edges are pivots.

As we noted at the end of the previous section, we need only find sufficient shortcuts such that every candidate long facet path is shown to have at least one of them. We therefore implemented a variant of the oriented matroid search algorithm that finds shortcuts in a current candidate realization in the style of cutting-plane algorithms. This is the algorithm we used to prove Proposition 2.4. It yields nonrealizability for all cases in Table 6. Thus, we get the following proposition, which, together with Proposition 2.4, proves that $\Delta(6, 12) = 6$.

Proposition 4.2. *None of the path complexes of Table 6 can be realized as part of the boundary complex of a matroid polytope.*

5. THE CASE $\Delta(4, 11)$

The method outlined in the sections above can also be used to show that $\Delta(4, 11) = 6$. The difference is the generation of the path complexes. Since we can restrict ourselves to end-disjoint paths, the number of revisits is bounded by 3. The paths with up to one revisit can be generated as outlined in Section 2.

Here we can use the symmetry-breaking methods outlined previously. The paths with two and three revisits are generated similarly, but this time we get two (respectively three) loops. Here we did not use any symmetry reduction.

We need, however, an additional result from [Bremner et al. 05].

Lemma 5.1. *If $\Delta(d - 1, n - 1) < l - 1$, then an index sequence of length l on d symbols in which (a) the symbol 1 appears uniquely at the beginning or in which (b) d appears uniquely at the end cannot correspond to a shortest path.*

Sketch of proof. Either of the given conditions means that the corresponding nonrevisiting path complex contains a common vertex on the last $l - 1$ facets in case (a) and the first $l - 1$ facets in case (b). \square

# revisits	# path complexes
0	35
1	185
2	354
3	96

TABLE 8. Number of path complexes after symmetry reduction for the case $(4, 11)$.

Note that the pruning of Lemma 5.1 takes place before any identifications, but the condition of having a common vertex in $l - 1$ facets is not changed by identifying vertices. Further, note that Lemma 5.1 eliminates candidate complexes only when we have a sufficiently strong bound on $\Delta(d - 1, n - 1)$.

We can combine the results of this section to generate a set of possible index sequences that satisfy all of the given criteria. The number of path complexes generated can be found in Table 8.

We can then use the methods outlined in Sections 3 and 4 to check that none of these path complexes can be realized as part of a matroid polytope. One complication is that revisiting paths need not use all of the vertices when $n > 2d$. Although it is relatively straightforward to add constraints to enforce that all points of the oriented matroid be contained in some facet (see [Bokowski et al. 09] for details), here we rely on the observation that any realization of a k -path that fails to have all of the points on the boundary is also a realization for fewer points. Since it is known [Goodey 72] that for $k \leq 10$, $\Delta(4, k) \leq 5$, we may ignore such a possibility.

6. CONCLUSION

The actual SAT computations took less than one hour for each of the ten cases in the $(6, 12)$ case on a regular desktop computer. However, the problems get more difficult for higher parameters. The key parameter seems to be $n - d$; so far, we have not been able to finish the computations for the $(5, 12)$ case in reasonable time, even though they are in a lower dimension. In these cases, the SAT solver can be made to produce an actual proof of infeasibility. However, the produced proofs are too unwieldy to be checked manually.

To be able to finish our computation, it is crucial that we do not rely on an explicit enumeration of all matroid polytopes that attain a given bound. All neighborly matroid polytopes of dimension 4 with 11 vertices are enumerated in [Schuchert 95]. Schuchert found 6 492 neighborly matroid polytopes with these parameters having diameter 6.

For larger computations, the SAT problems become considerably harder. However, it might be more interesting to study the path complexes in more detail. It would be potentially useful to give criteria that further reduce the number of equivalence classes of path complexes to be considered. However, the number of classes will probably still grow too fast to make dealing with cases with much larger parameters feasible.

NOTE ADDED IN PROOF

Recently, Francisco Santos announced a 43-dimensional counterexample to the d -step conjecture [Santos 10]. There remains a great deal of interest in the status of the Hirsch conjecture in lower dimensions, and in whether some polynomial bound on the diameter of convex polytopes holds in general.

ACKNOWLEDGMENTS

This research was carried out while the authors were guests of the Hausdorff Research Institute for Mathematics, Bonn. The research of the first author was partly supported by NSERC Canada. Computational resources were supplied by the ACEnet consortium.

REFERENCES

- [Björner et al. 99] Anders Björner, Michel Las Vergnas, Bernd Sturmfels, Neil White, and Günter M. Ziegler. *Oriented Matroids*, 2nd ed., Encyclopedia of Mathematics and Its Applications 46. Cambridge, UK: Cambridge University Press, 1999.
- [Bokowski et al. 09] Jürgen Bokowski, David Bremner, and Gábor Gévay. “Symmetric Matroid Polytopes and Their Generation.” *European J. Combin.* 30:8 (2009), 1758–1777.
- [Bremner et al. 05] David Bremner, Fred B. Holt, and Victor Klee. “Path Complexes and Polytope Diameters.” Unpublished manuscript, 2005.
- [Eén and Sörensson 03] Niklas Eén and Niklas Sörensson. “An Extensible SAT-Solver.” In *SAT (2003)*, edited by Enrico Giunchiglia and Armando Tacchella, Lecture Notes in Computer Science 2919, pp. 502–518. New York: Springer, 2004.
- [Goodey 72] P. R. Goodey. “Some Upper Bounds for the Diameters of Convex Polytopes.” *Israel J. Math.* 11 (1972), 380–385.
- [Grünbaum 03] Branko Grünbaum. *Convex Polytopes*, 2nd ed., Graduate Texts in Mathematics 221. New York: Springer-Verlag, 2003.
- [Holt 04] Fred B. Holt. “Blending Simple Polytopes at Faces.” *Discrete Math.* 285:1–3 (2004), 141–150.
- [Klee 64] Victor Klee. “Diameters of Polyhedral Graphs.” *Canad. J. Math.* 16 (1964), 602–614.
- [Klee and Kleinschmidt 87] Victor Klee and Peter Kleinschmidt. “The d -Step Conjecture and Its Relatives.” *Math. Oper. Res.* 12:4 (1987), 718–755.
- [Klee and Walkup 67] Victor Klee and David W. Walkup. “The d -Step Conjecture for Polyhedra of Dimension $d < 6$.” *Acta Math.* 117 (1967), 53–78.
- [McKay 05] Brendan D McKay. “Nauty.” Available online (<http://cs.anu.edu.au/people/bdm/nauty/>), 2005.
- [Santos 10] Francisco Santos. “A Counterexample to the Hirsch Conjecture.” arXiv:1006.2814v1, 2010.
- [Schewe 07] Lars Schewe. “Satisfiability Problems in Discrete Geometry.” PhD thesis, TU Darmstadt, 2007.
- [Schewe 10] Lars Schewe. “Non-realizable Minimal Vertex Triangulations of Surfaces: Showing Non-realizability Using Oriented Matroids and Satisfiability Solvers.” *Disc. Comput. Geom.* 43:2 (2010), 289–302.
- [Schuchert 95] Peter Schuchert. “Matroid-Polytope und Einbettungen kombinatorischer Mannigfaltigkeiten.” PhD thesis, TU Darmstadt, 1995.
- [Ziegler 95] Günter M. Ziegler. *Lectures on Polytopes*, Graduate Texts in Mathematics 152. New York: Springer-Verlag, 1995.

David Bremner, Faculty of Computer Science, University of New Brunswick, Box 4400, Fredericton NB, E3B 5A3, Canada (bremner@unb.ca)

Lars Schewe, Department of Mathematics, University Erlangen-Nuremberg, Erlangen, Germany (lars.schewe@math.uni-erlangen.de)

Received September 14, 2009; accepted June 6, 2010.