On the Hardness of Minkowski Addition and Related Operations

Hans Raj Tiwary*
FR Informatik, Universität des Saarlandes
Saarbrücken, Germany
hansraj@cs.uni-sb.de

ABSTRACT

For polytopes $\mathcal{P}, \mathcal{Q} \subset \mathbb{R}^d$ we consider the intersection $\mathcal{P} \cap \mathcal{Q}$; the convex hull of the union $CH(\mathcal{P} \cup \mathcal{Q})$; and the Minkowski sum $\mathcal{P} + \mathcal{Q}$. We prove that given rational $\mathcal{H}$-polytopes $\mathcal{P}_1, \mathcal{P}_2, \mathcal{Q}$ it is impossible to verify in polynomial time whether $\mathcal{Q} = \mathcal{P}_1 + \mathcal{P}_2$, unless $P = NP$. In particular, this shows that there is no output sensitive polynomial algorithm to compute the facets of the Minkowski sum of two arbitrary $\mathcal{H}$-polytopes even if we consider only rational polytopes. Since the convex hull of the union and the intersection of two polytopes relate naturally to the Minkowski sum via the Cayley trick and polarity, similar hardness results follow for these operations as well.

Categories and Subject Descriptors: F.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems

General Terms: Algorithm, Theory

Keywords: Minkowski Addition, Extended Convex Hull, Polytope intersection, Polytopes, coNP-hardness, Turing reduction, Computational Geometry

1. INTRODUCTION

A *convex polyhedron* or simply *polyhedron* in d-dimensional euclidean space $\mathbb{R}^d$ is the intersection of a finite number of halfspaces. A polyhedron is called *pointed* if it does not contain any affine line in its interior and *bounded* if it does not contain any ray. A bounded polyhedron is also called a polytope. A very basic result in the theory of polyhedra states that a polyhedron can be described both as the intersection of a finite number of halfspaces as well as the Minkowski sum of $conv(V) + cone(Y)$, where V and Y are finite sets of points in $\mathbb{R}^d$. For a thorough treatment of polytopes Grünbaum [11] and Ziegler [16] are excellent sources.

*The author was supported by Graduiertenkolleg der FR Informatik, Universität des Saarlandes while working on this problem.

SCG'07, June 6–8, 2007, Gyeongju, South Korea.

In this paper, we will talk mostly about polytopes and refer to the unbounded case only towards the end. We refer to the two equivalent representations as $\mathcal{H}$-representation and $\mathcal{V}$-representation respectively. Accordingly a polytope described by its vertices is called a $\mathcal{V}$-polytope and a polytope described by its facets is called an $\mathcal{H}$-polytope. Many operations that are easy to perform on one description become difficult if the description is changed. To give a simple example, finding a point inside a polytope that maximizes the inner product with a given vector can be done trivially if the polytope is in $\mathcal{V}$-representation but for the $\mathcal{H}$-representation this amounts to Linear Programming for which only weak polynomiality is known [14].

In this paper, we study three fundamental operations on polytopes and provide hardness results for them. For polytopes $\mathcal{P}, \mathcal{Q}$ in $\mathbb{R}^d$, the Minkowski addition $\mathcal{P} + \mathcal{Q}$, the convex hull of the union $CH(\mathcal{P} \cup \mathcal{Q})$ and the intersection $\mathcal{P} \cap \mathcal{Q}$ are defined as:

$$\begin{aligned} \mathcal{P} + \mathcal{Q} &= \{x + y | x \in P, y \in Q\} \\ CH(\mathcal{P} \cup \mathcal{Q}) &= \{\lambda x + (1-\lambda) y | x \in P, y \in Q, 0 \le \lambda \le 1\} \\ \mathcal{P} \cap \mathcal{Q} &= \{x | x \in P \, and \, x \in Q\} \end{aligned}$$

Throughout the paper $\mathcal{P}, \mathcal{Q}$ refer to polytopes and d denotes the dimension of the affine space in which $\mathcal{P}, \mathcal{Q}$ live. Also, *size* of a polytope refers to the number of vertices (facets resp.) of $\mathcal{V}$-polytope ($\mathcal{H}$-polytope resp.) times the dimension d.

We are interested in the complexity of performing these operations and providing non-redundant description of the resulting polytope in appropriate representation. Since the worst case size of the output for all the three operations can be exponential in the size of input (see [12, 10]), it is natural to talk of *output sensitive* algorithms. The complexity of an output sensitive algorithm is measured in terms of the size of both input and the output. Thus, a polynomial output sensitive algorithm is one whose running time is polynomial in the size of input and output.

We consider only the case when the operand polytopes are in the same representation as the resulting polytope. For mixed representations one has to implicitly perform representation conversion. The problem of converting one representation to the other has been well studied. The problem of converting $\mathcal{V}$-representation to $\mathcal{H}$-representation is the well known Convex Hull problem and the reverse is known as the Vertex Enumeration problem. No polynomial output sensitive algorithm is known for this problem except for some special cases (see [1, 4, 5, 15]).

It is easy to see that computing the non-redundant $\mathcal{V}$-representation of $CH(\mathcal{P} \cup \mathcal{Q})$ and $\mathcal{P} + \mathcal{Q}$ is easy if $\mathcal{P}, \mathcal{Q}$ are $\mathcal{V}$-polytopes and an oracle for LP is given. Similarly, computing the $\mathcal{H}$-representation of $\mathcal{P} \cap \mathcal{Q}$ is easy if $\mathcal{P}, \mathcal{Q}$ are $\mathcal{H}$-polytopes and an oracle for LP is given. So, we are interested in other versions, namely the ones where the operands for $CH(\mathcal{P} \cup \mathcal{Q})$ and $\mathcal{P} + \mathcal{Q}$ are $\mathcal{H}$-polytopes and the operands for $\mathcal{P} \cap \mathcal{Q}$ are $\mathcal{V}$-polytopes. Before we proceed to the similarities between the three operations, a short description of polarity and the Cayley trick is in order.

1.1 Polarity

Let $\mathcal{P} = \{x | Ax \leq \mathbf{1}\}$ be a full dimensional polytope in $\mathbb{R}^d$ containing the origin in its interior. Here $A \in \mathbb{R}^{m \times d}$ and $\mathbf{1}$ is an $m \times 1$ column vector with all entries 1. The polar (also dual) of $\mathcal{P}$, denoted by $\mathcal{P}^*$, is obtained by treating the vectors $a_i \in A$ as points in dimension d and taking the convex hull of all these points. For a detailed treatment of this operation the reader is again referred to [11]. One interesting property of the polar operation is that points in the interior of $\mathcal{P}$ are mapped to hyperplanes that don't intersect $\mathcal{P}^*$. Similarily points on the boundary of $\mathcal{P}$ are mapped to hyperplanes that touch $\mathcal{P}^*$ and points outside $\mathcal{P}$ are mapped to hyperplanes that intersect the interior of $\mathcal{P}^*$.

The convex hull of the union; and intersection operations are related via polar duality. More precisely, if $\mathcal{P}, \mathcal{Q}$ are two full dimensional polytopes in $\mathbb{R}^d$ both containing origin in the relative interior, then $\mathcal{P} \cap \mathcal{Q}$ is the polar dual of $CH(\mathcal{P}^* \cup \mathcal{Q}^*)$.

1.2 The Cayley trick

The Cayley trick is a simple embedding of k d-dimensional polytopes $P_1, \cdots, P_k$ into $\mathbb{R}^{d+k-1}$. The embedding is obtained by appending $\mathbf{e}_{i-1}$ to every point in P_i, where $\mathbf{e}_0 = \mathbf{0}$ and $\mathbf{e}_i$ is the i-th unit vector of $\mathbb{R}^{k-1}$ and taking the convex hull of all the embedded copies. It is easy to see that the Minkowski sum of these polytopes (up to a scaling) can be obtained from the Cayley embedding by intersecting the polytope obtained after the embedding with a suitable d-flat. To illustrate this, consider the case when $k = 2$. The Cayley embedding is obtained by putting a copy of P_1 in the hyperplane $x_{d+1} = 0$ and a copy of P_2 in the hyperplane $x_{d+1} = 1$ and taking the convex hull of both embedded polytopes. The Minkowski sum (scaled by a factor half), then, is the intersect of the resulting polytope with the hyperplane $x_{d+1} = \frac{1}{2}$.

The rest of the paper is organized as follows. In the following section, we describe prior work related to performing these operations in appropriate representations and in Section 3 we establish a hardness result that explains why no polynomial algorithm exists for performing these operations on arbitrary polytopes given in arbitrary representation.

2. RELATED WORK

The problem of enumerating the facets of $CH(\mathcal{P} \cup \mathcal{Q})$, when both $\mathcal{P}$ and $\mathcal{Q}$ are given by their facets, has been studied in [2] and [8]. Balas [2] constructs polynomial algorithm for a special class of polytopes while Fukuda, Liebling and Lütolf [8] present polynomial algorithms with certain assumptions about the polytopes. A closely related problem is to determine whether the union of k polytopes (in either representation) is convex. As noted by Bemporad *et. al* in [3], the union is convex if and only if it coincides with the convex hull of the union. They also present a polynomial algorithm for the case when when $k = 2$.

Minkowski sums have been studied in much more detail compared to convex hull of union. They frequently come up in computational algebra [10], robotics and motion planning, geometric convexity, computer graphics and many other areas. Roughly speaking, Minkowski sums show up in two context - computing the facets (or vertices) of Minkowski sum of k-polytopes given by their facets (or vertices); and computing mixed volumes of k polytopes.

Gritzmann and Sturmfels [10] studied Minkowski sum in the context of computational algebra and gave (exponential) bounds on the number of faces of Minkowski sum. They also gave examples of cases where the bounds were tight. As noted before, exponential lower bounds motivate one to look for output sensitive algorithms for cases when the output is far from worst case. Fukuda and Weibel [9, 7] propose polynomial algorithms for enumerating faces of Minkowski sum of k polytopes where each polytopes is given by vertices. They explicitly avoid the case when the input polytopes are described by facets and the facets of Minkowski sum are to be enumerated. Fukuda and Weibel, in another work (see [12]), study Minkowski sums of special polytopes that are "well centered" and also provide better bounds on the number of faces.

In the next section, we prove that unless $P = NP$, there is no output sensitive polynomial algorithm that computes the facets of the Minkowski sum of two $\mathcal{H}$-polytopes. It follows, from the Cayley trick and polarity, that enumerating the facets of the convex hull of the union of two $\mathcal{H}$-polytopes; or enumerating the vertices of the intersection of two $\mathcal{V}$-polytopes is hard as well.

3. HARDNESS OF MINKOWSKI ADDITION

In this section we establish the hardness of enumerating the facets of the Minkowski sum of two $\mathcal{H}$-polytopes. We prove this by proving a possibly stronger statement. Consider the following decision version of the enumeration problem

PROBLEM MINKOWSKIVERIFY
INPUT: $\mathcal{H}$-Polytopes $\mathcal{P}_1, \mathcal{P}_2, \mathcal{Q}$.
OUTPUT: Yes, if $\mathcal{Q} = \mathcal{P}_1 + \mathcal{P}_2$. No, otherwise.

THEOREM 1. *There is no algorithm solving* MINKOWSKIVERIFY *in polynomial time, unless* $P = NP$.

We prove this by providing a polynomial-time Turing reduction from a coNP-Hard problem to this problem. A problem $\mathcal{A}$ is said to be polynomial-time Turing reducible to problem $\mathcal{B}$ if, one can construct an algorithm for problem $\mathcal{A}$ from an algorithm for $\mathcal{B}$ by using the latter as a subroutine and invoking it polynomially many times. The more common Karp reduction allows only one call to the oracle and that too at the end.

It was shown by Khachiyan *et. al*[13] that it is coNP-Hard to enumerate all vertices of a polyhedron given by its facets. The following theorem restates the result of [13].

THEOREM 2. *Given a polyhedron* $\mathcal{P}$ *in* $\mathcal{H}$*-representation and a set* $\mathcal{V}$ *of vertices of* $\mathcal{P}$*, it is* $coNP$*-complete to decide whether* $\mathcal{V}$ *is the complete vertex set of* $\mathcal{P}$.

Now, we prove that if we have an algorithm for deciding MINKOWSKIVERIFY for two arbitrary polytopes, then we can invoke the oracle polynomial number of times and decide for some set of vertices V and and an $\mathcal{H}$-polytope $\mathcal{P}$, whether $V = vert(\mathcal{P})$. Note that a very important fact for the Theorem 2 is that a polyhedron can have many more extreme rays than the number of vertices. For polyhedra that don't have too many extreme rays, the problem of enumerating vertices is effectively the same as the problem of enumerating all vertices of a polytope. The complexity status of the latter is unknown.

Our reduction works in the following way. Let $\mathcal{P} = \{x|Ax \leq b\}$ be a polyhedron in $\mathbb{R}^d$ and V be a subset of vertices $vert(P)$. We assume that no two vertices have the same x_d-coordinate. We also assume that $\mathcal{P} \cap \{x_d \geq -c\}$ is bounded for sufficiently large c. The first assumption ensures a unique ordering on the vertices of the polyhedron and the second assumption ensures that the polyhedron that are generated to be passed onto the MINKOWSKIVERIFY oracle are indeed polytopes *i.e.* bounded.

To see that the above assumptions can indeed be satisfied by a polynomial preprocessing, recall that a polyhedron has a unique representation as the minkowski sum of a polytope and a cone. Also, the cone of the polyhedron $\mathcal{P} = \{x|Ax \leq b\}$ is just $cone(\mathcal{P}) = \{x|Ax \leq 0\}$ with some inequalities possibly redundant. To satisfy the second assumption, we need to pick a vector c such that $cone(\mathcal{P}) \cap \{c \cdot x \leq 1\}$ is bounded; and align the negative x_d-axis along c. Picking a vector c from $cone(A)$[1] does the trick.

Now, if c is picked uniformly at random from $cone(A)$ then the probability of any two vertices of the polyhedron $\mathcal{P}$ having the same x_d-coordinate is 0 and so the first assumption is satisfied too. Random sampling from such convex objects can be done in time polynomial in the dimension using an algorithm of Dyer, Frieze and Kannan [6].

Now for our reduction, consider the vertices v_i of V in the order of their x_d-coordinate. That is, if $\mathbf{e}_d$ is the unit vector $(0, \cdots, 0, 1)$ in $\mathbb{R}^d$ and $\mathbf{e}_d$ is considered as *up* then the vertices are considered in the order of increasing height. Now, consider some v_i and v_{i+1} and define three polytopes in the following way:

$$\begin{aligned} \mathcal{P}_{-1} &= \mathcal{P} \cap \{x_d = v_i \cdot \mathbf{e}_d\} \\ \mathcal{P}_1 &= \mathcal{P} \cap \{x_d = v_{i+1} \cdot \mathbf{e}_d\} \\ \mathcal{P}_0 &= \mathcal{P} \cap \left\{x_d = \frac{v_i \cdot \mathbf{e}_d + v_{i+1} \cdot \mathbf{e}_d}{2}\right\} \end{aligned}$$

where the dot product $v_i \cdot \mathbf{e}_d$ is nothing but the x_d-coordinate of v_i.

The following lemma provides the necessary tool for using the oracle for MINKOWSKIVERIFY as a subroutine for verifying whether a set of vertices of an arbitrary polyhedron defines its vertex set completely.

LEMMA 1. *$2\mathcal{P}_0 \neq \mathcal{P}_{-1} + \mathcal{P}_1$ if and only if there exists some $v \in vert(\mathcal{P})$ that is not in V and $v_i \cdot \mathbf{e}_d < v \cdot \mathbf{e}_d < v_{i+1}.\mathbf{e}_d$.*

PROOF. We prove the non-trivial direction only. Suppose some vertex $v \in vert(\mathcal{P})$ is not in V and $v_i \cdot \mathbf{e}_d < v \cdot \mathbf{e}_d < v_{i+1}.\mathbf{e}_d$ for some i. Without loss of generality we can assume that v lies above the hyperplane containing $\mathcal{P}_0$. If so, there is an $u \in vert(\mathcal{P}_{-1})$ such that $\overrightarrow{uv}$ lies on some edge of $\mathcal{P}$. Clearly, $\overrightarrow{uv}$ intersects $\mathcal{P}_0$, say at w. We claim that $2w \notin \mathcal{P}_{-1} + \mathcal{P}_1$.

Assume for the sake of contradiction that $2w \in \mathcal{P}_{-1} + \mathcal{P}_1$. Then there are $x \in \mathcal{P}_{-1}$ and $y \in \mathcal{P}_1$ such that $2w = x + y$. Since, any point on an edge of a polytope can be *uniquely* represented as the convex combination of the vertices defining the edge, it follows that $x = u$ and y is a vertex of $\mathcal{P}_1$. This implies that v is a convex combination of x, y as well and hence, v can not be a vertex of $\mathcal{P}$, a contradiction. $\square$

Lemma 1 gives us a way to use the algorithm for Minkowski sum computation for deciding whether some vertex between v_i and v_{i+1} is missing. So by invoking the Minkowski sum algorithm $|V|$ times we can decide whether or not $V = vert(\mathcal{P})$. This gives us the following corollary to Theorem 1:

[1]This is the polar dual of $cone(\mathcal{P})$.

COROLLARY 1. *Given two $\mathcal{H}$-polytopes $\mathcal{P}, \mathcal{Q} \in \mathbb{R}^d$, there is no output sensitive polynomial algorithm that enumerates the facets of $\mathcal{P} + \mathcal{Q}$ unless $P = NP$.*

Similarly, if we define $\mathcal{P}_0$ as:

$$\mathcal{P}_0 = \mathcal{P} \cap \{x_d \geq v_i \cdot \mathbf{e}_d, x_d \leq v_{i+1} \cdot \mathbf{e}_d\}$$

then it is easy to see that $\mathcal{P}_0 \neq CH(\mathcal{P}_{-1} \cup \mathcal{P}_1)$ if and only if there exists some $v \in vert(\mathcal{P})$ that is not in V and $v_i \cdot \mathbf{e}_d < v \cdot \mathbf{e}_d < v_{i+1}.\mathbf{e}_d$. The proof is analogous to the proof of Lemma 1 and so we omit it. Thus, we have the following theorem:

THEOREM 3. *Given $\mathcal{H}$-polytopes $\mathcal{P}_1, \mathcal{P}_2, \mathcal{Q} \in \mathbb{R}^d$, there is no algorithm that verifies in polynomial time whether $CH(\mathcal{P}_1 \cup \mathcal{P}_2) = \mathcal{Q}$ unless $P = NP$.*

Note that, the operand polytopes here are not full dimensional. However, this can easily be fixed by taking a point p in the interior of $\mathcal{P}_0$ and constructing a pyramid over $\mathcal{P}_{-1}$ and $\mathcal{P}_1$ with p as the apex. Again, arguing as earlier, we have the following corollary:

COROLLARY 2. *Given two $\mathcal{H}$-polytopes $\mathcal{P}, \mathcal{Q} \in \mathbb{R}^d$, there is no polynomial output sensitive polynomial that enumerates the facets of $CH(\mathcal{P} \cup \mathcal{Q})$ unless $P = NP$.*

And lastly because of polar duality, we have the following corollary:

COROLLARY 3. *Given two $\mathcal{V}$-polytopes $\mathcal{P}, \mathcal{Q} \in \mathbb{R}^d$, there is no polynomial output sensitive polynomial that enumerates the vertices of $\mathcal{P} \cap \mathcal{Q}$ unless $P = NP$.*

4. CONCLUDING REMARKS

In this paper we proved that unless $P = NP$, there does not exist a polynomial output sensitive algorithm for

- enumerating all facets of the Minkowski sum of two $\mathcal{H}$-polytopes,
- enumerating all vertices of the intersection of two $\mathcal{V}$-polytopes or
- enumerating all facets of the convex hull of the union of two $\mathcal{H}$-polytopes.

In all the three cases, the input and the output polytopes have same representation. An interesting case is when the input and output representations are not same. Effectively these “mixed representation" versions need to efficiently solve the problem of converting $\mathcal{H}$-representation to $\mathcal{V}$-representation (Vertex Enumeration) or the other way round (Convex Hull). For this mixed version neither any polynomial time output sensitive algorithms are known nor any hardness results, analogous to the ones presented here, are known.

References

[1] D. Avis and K. Fukuda. A pivoting algorithm for convex hulls and vertex enumeration of arrangements and polyhedra. *Discrete & Computational Geometry*, 8:295–313, 1992.

[2] E. Balas. On the convex hull of the union of certain polyhedra. *Operations Research Letters*, 7:279–283, 1988.

[3] A. Bemporad, K. Fukuda, and F. D. Torrisi. Convexity recognition of the union of polyhedra. *Comput. Geom*, 18(3):141–154, 2001.

[4] D. Bremner, K. Fukuda, and A. Marzetta. Primal - dual methods for vertex and facet enumeration. *Discrete & Computational Geometry*, 20(3):333–357, 1998.

[5] M. R. Bussieck and M. E. Lübbecke. The vertex set of a 0/1-polytope is strongly P-enumerable. *Comput. Geom*, 11(2):103–109, 1998.

[6] M. E. Dyer, A. M. Frieze, and R. Kannan. A random polynomial time algorithm for approximating the volume of convex bodies. *J. ACM*, 38(1):1–17, 1991.

[7] K. Fukuda. From the zonotope construction to the minkowski addition of convex polytopes. *J. Symb. Comput.*, 38(4):1261–1272, 2004.

[8] K. Fukuda, T. M. Liebling, and C. Lutolf. Extended convex hull. *Computational Geometry*, 20(1-2):13–23, 2001.

[9] K. Fukuda and C. Weibel. Computing all faces of the minkowski sum of $\mathcal{V}$-polytopes. *In Proceedings of the 17th Canadian Conference on Computational Geometry*, 2005.

[10] Gritzmann and Sturmfels. Minkowski addition of polytopes: Computational complexity and applications to grobner bases. *SIJDM: SIAM Journal on Discrete Mathematics*, 6, 1993.

[11] B. Grünbaum. *Convex Polytopes*. Wiley Interscience Publ., London, 1967.

[12] K.Fukuda and C. Weibel. On f-vectors of minkowski additions of convex polytopes. *Technical report*, 2005. submitted to DCG.

[13] L. Khachiyan, E. Boros, K. Borys, K. M. Elbassioni, and V. Gurvich. Generating all vertices of a polyhedron is hard. In *SODA*, pages 758–765. ACM Press, 2006.

[14] L. G. Khachiyan. A polynomial algorithm in linear programming. *Soviet Mathematics Doklady*, 20:191–194, 1979.

[15] J. S. Provan. Efficient enumeration of the vertices of polyhedra associated with network LP's. *Math. Program*, 63:47–64, 1994.

[16] G. M. Ziegler. *Lectures on Polytopes*. Graduate Texts in Mathematics, No. 152. Springer-Verlag, Berlin, 1995.