An Efficient Computation of Mixed Volume

Ioannis Emiris
Computer Science Division
571 Evans Hall
University of California, Berkeley
Berkeley CA 94720
emiris@cs.berkeley.edu

Abstract

The Mixed Volume of n polytopes in n-dimensional space is a multilinear function with respect to Minkowski addition and scalar multiplication that generalizes the notion of volume. The current interest in efficient methods for computing it is mainly due to Bernstein's theorem [Be] which bounds the number of common roots of a system of polynomial equations by the Mixed Volume of the respective Newton polytopes. In this paper we propose an algorithm that uses polyhedral techniques to significantly improve upon the efficiency of the method based on evaluating the inclusion-exclusion formula and obtain a practical algorithm.

1 Introduction

The problem discussed in this paper is the following: Given n polytopes $\Delta_1, \ldots, \Delta_n$ in n-dimensional real space, compute their Mixed Volume. In a more general context, given are n point sets in $\mathbf{R}^n$ and the question is to find the Mixed Volume of the respective convex hulls $\Delta_1, \ldots, \Delta_n$.

Let $\alpha\Delta$, with $\alpha \in \mathbf{R}$, denote the set of all scalar multiples of points in polytope Δ by α and $\Delta_1 \oplus \Delta_2$ the set of all component-wise vector sums of points in Δ_1 and Δ_2 respectively; let $V_n(\Delta)$ denote the standard n-dimensional Lebesque measure. Then the Mixed Volume of $\Delta_1, \ldots, \Delta_n$ is defined as follows.

Definition 1.1 *The Mixed Volume is multilinear in each of $\Delta_1, \ldots, \Delta_n$ with respect to Minkowski sum, scaled so that the Mixed Volume of n identical polytopes Δ equals $V_n(\Delta)$.*

Alternatively, we have

Definition 1.2 *Consider the form expressing the n-dimensional volume $V_n(\alpha_1\Delta_1 \oplus \cdots \oplus \alpha_n\Delta_n)$, where α_i for $i \in \{1, \ldots, n\}$ is a symbolic real variable. Then the Mixed Volume is the coefficient of the $\alpha_1 \cdots \alpha_n$ term.*

It can be proven that another expression for the Mixed Volume is given by the classical Inclusion-Exclusion formula:

$$\frac{1}{n!}[V_n(\Delta_1 \oplus \cdots \oplus \Delta_n) - \sum_{t=1}^{n} V_n(\Delta_1 \oplus \cdots \oplus \Delta_{t-1} \oplus \Delta_{t+1} \oplus \cdots \oplus \Delta_n) + \cdots + (-1)^{n-1} \sum_{t=1}^{n} V_n(\Delta_t)],$$

which establishes the uniqueness of the Mixed Volume.

The importance of this problem stems from the fact that it provides an upper bound on the number of solutions for a polynomial system whose Newton polytopes are the $\Delta_1, \ldots, \Delta_n$ polytopes. For sparse systems this bound is usually significantly tighter than the Bezout bound. This topic is introduced and explored in a variety of papers [Be], [Ku], [Kh], [CaRo].

Applying directly the Inclusion-Exclusion formula implies an algorithm with worst-case asymptotic complexity at least $\Theta(m^{n^2})$, where m is the maximum number of vertices on any Δ_i and thus has to exceed n. Here we take a local approach to the problem and design a more efficient algorithm that computes the Mixed Volume in $\mathcal{O}(m^{2n}n^2)$. Our algorithm is conceptually straightforward, requires only basic linear algebra operations and makes no genericity assumption about the input. Another local method [Ma] does not differ in any essential way and offers no significant complexity improvement.

2 A Regular Decomposition of the Minkowski Sum

The basic idea is to decompose the Minkowski Sum into polytopes that contribute either their volume or zero to Mixed Volume and then design a local test for deciding the type of each component polytope. Naturally, each of these polytopes is the Minkowski Sum of certain faces of the original polytopes and the problem would then reduce to finding which combinations of faces contribute to Mixed Volume. There are more than one ways for partitioning so our first concern is to specify a unique partition with the above property.

Let Q be the convex hull of the Minkowski Sum $\Delta_1 \oplus \cdots \oplus \Delta_n$, which lies in $\mathbf{R}^n$. We define the *product polytope* P to be the convex hull of all points in n^2-dimensional space of the form $p_1 \times \cdots \times p_n = (p_1, \ldots, p_n)$, where $p_i \in \Delta_i,\ \forall i$; we shall also write $P = \Delta_1 \times \cdots \times \Delta_n$. We relate these two polytopes through a linear mapping π from $\mathbf{R}^{n^2}$ to $\mathbf{R}^n$, where $\pi : p = (p_1, \ldots, p_n) \mapsto p_1 + \cdots + p_n$; the corresponding matrix is $n^2 \times n$ and is composed of n blocks of $n \times n$ unit matrices. Then $Q = \pi(P)$; the dimension of the map's kernel is $n^2 - n$ and similarly for the preimage of any point q in Q.

$$(\Delta_1, \ldots, \Delta_n) \;\rightarrow\; P = \times_{i=1}^{n} \Delta_i \;\stackrel{\pi}{\rightarrow}\; Q = \oplus_{i=1}^{n} \Delta_i.$$

At this point we need some definitions and facts from [St] applied to the Minkowski Sum Q and the product polytope P; also helpful is [Br]. A *polyhedral complex* is any collection of polytopes such that the intersection of any two is a face of each and is in the collection. A *polyhedral subdivision* of the convex hull Q of a set of points $\mathcal{B}$ is a collection Π of proper subsets of $\mathcal{B}$ whose convex hulls form a polyhedral complex with their union being equal to Q. A polyhedral subdivision Π is *induced by* map $\pi : P \rightarrow Q$ if each cell $\sigma \in \Pi$ is of the form $\pi(F_\sigma)$ for some face F_σ of P. An induced subdivision is *tight* if σ and F_σ have equal dimension, for each $\sigma \in \Pi$. For a linear functional ψ on $\mathbf{R}^{n^2}$, a face of some polytope S is

$$S^\psi = \{x \in S : \psi^T x \geq \psi^T y,\ \forall y \in S\},$$

where $v^T w$ will denote the inner product of two column vectors v, w. We call a subdivision Π *coherent* if there exists a linear functional ψ on $\mathbf{R}^{n^2}$ such that

$$(\pi^{-1} q)^\psi = \pi^{-1} q \cap F_\sigma,\ \forall q \in \sigma \subset Q. \tag{1}$$

The choice of the linear functional uniquely determines a coherent subdivision; furthermore, for generic ψ, this subdivision is tight.

In the rest of this paper we assume that the linear functional ψ is generic. Let Π denote the coherent subdivision of Q obtained from this ψ and let σ denote an arbitrary maximal cell of Π. Clearly, the maximal cells are n-dimensional and, due to the genericity of ψ, they are the images under π of n-faces of P. Additionally, if σ is the Minkowski Sum of certain faces from the original polytopes, then F_σ is the product of these faces. This associates at least one n-face F_σ of P with every maximal cell σ. We wish to make this correspondence bijective, ie. choose a unique F_σ among those in the preimage of σ.

For every point $q \in Q$, the preimage $\pi^{-1}q$ is an $(n^2 - n)$-dimensional polytope. From the genericity of ψ it follows that $(\pi^{-1}q)^{\psi}$ is a vertex of this polytope which, by (1), lies on some n-face F_σ of P. This n-face is unique for a given point q and also unique and well-defined for a given cell σ.

It is straightforward to show that an induced tight coherent Π imposes some partition of Q into Mixed Volume and non-Mixed Volume cells. In other words, every maximal cell σ contributes either zero or its n-dimensional volume to Mixed Volume. The argument is based on the fact that every σ of the first type, unlike those of the second type, is the Minkowski Sum of edges.

When σ contributes to Mixed Volume then F_σ is the product of 1-faces, or edges, in the original polytopes. Assume that $F_\sigma = e_1 \times \cdots \times e_n$, where $e_i \in \Delta_i, \forall i$. Take any point $q \in \sigma$ and denote by p the maximal vertex $(\pi^{-1}q)^{\psi}$ in the intersection of its preimage and F_σ. There exist a unique n-tuple of points $p_1, \ldots, p_n$ such that $p_1 + \cdots + p_n = p$ and $p_i \in e_i$. To see this, if p is allowed to move within F_σ, each p_i may also move, on edge e_i, thus identifying a unique edge on each Δ_i.

We now have a regular way, given a generic linear functional, to decompose the Minkowski Sum into n-dimensional cells that contribute either zero or their volume to Mixed Volume and, for the latter kind, the set of edges producing it is unique. Moreover, every cell is in a bijective correspondence with some n-face of the product polytope, where every point on the latter is the product of a unique n-tuplet of points from the respective original polytopes.

3 Characterization of the Mixed Volume

In this section we study the properties of an arbitrary cell $\sigma \in Q$ that contributes to Mixed Volume and try to derive corresponding properties for the face F_σ, the unique n-face associated with σ. Our strategy is to formulate a local test on the input polytopes that checks whether an edge n-tuplet contributes to Mixed Volume or not.

The fact that p is maximal in $\pi^{-1}q$ implies that if it is allowed to move in any direction within this preimage, it would cause the value of ψ to decrease. Recall that the value of ψ at any other point of $\pi^{-1}q$ is strictly smaller than that at p; it suffices to check those directions of movement defined by the intersection of the preimage with the $(n+1)$-faces cobounding F_σ. Each of these $(n+1)$-faces is the product of $n-1$ edges and a 2-face from some original polytope; the (j,k)-th one is produced by 2-face $t_{j,k}$ that belongs to Δ_j, where e_j lies in the boundary of $t_{j,k}$ and k does not exceed the number of 2-faces cobounding e_j in Δ_j. By a dimension counting argument the dimension of $(\pi^{-1}q) \cap (e_1 \times \cdots \times t_{j,k} \times \cdots \times e_n)$ is 1 in general; to guarantee this, we think of q as being in general position within σ. Movement along this intersection corresponds to simultaneous movement of the p_i's on the respective edges e_i for all $i \neq j$, while p_j is allowed to move along a line in $t_{j,k}$ that is not parallel to e_j.

Let $G_{j,k}$ represent some $(n+1)$-face cobounding F_σ, for $1 \leq j \leq n$ and appropriate values of k. The conditions are stated in terms of some generic point $q \in \sigma$ which does not need be explicitly specified.

1	F_σ is the product of 1-faces.
2	$p = (\pi^{-1}q)^{\psi}$ belongs to F_σ
$\Leftrightarrow$	the value of ψ decreases as p moves away from F_σ on $\pi^{-1}q \cap G_{j,k}$, for all $G_{j,k}$
$\Leftrightarrow$	the value of ψ decreases as each p_i moves on $e_i, \forall i \neq j$ and p_j moves on $t_{j,k}$.
3	$\det E \neq 0$.

The last condition concerns the amount of the contribution of some set of edges $e_1, \ldots, e_n : e_i \in \Delta_i$, in other words, the n-dimensional volume of the Minkowski Sum $e_1 \oplus \cdots \oplus e_n$. It is easy to see that this is zero if and only if the n n-dimensional real vectors representing the edges are not independent when translated to a common origin. The volume equals the absolute value of the determinant of an $n \times n$ matrix E, that has as *columns* the respective entries of the n vectors.

4 The Algorithm

In this section we formalize a local test on every combination of edges from different polytopes that decides whether this combination forms a Mixed Volume cell. After giving basic algorithm I, we consider various tricks to speed it up which lead to improved algorithms II and III; the performance analysis is found in the next section.

We represent the amount of movement of point p_i by $\lambda_i \in \mathbf{R}$, assuming that the origin lies on e_i at p_i. In other words, the new position of p_i after moving is $\lambda_i e_i$. Point p_j on Δ_j must necessarily move away from e_j in the interior of cobounding 2-face $t_{j,k}$, for some such face. For a given $t_{j,k}$, this movement can be described as the linear combination of movement along e_j and movement along some vector $v_{j,k}$ measured by a positive amount $\mu \in \mathbf{R}_+^*$; $v_{j,k}$ can be any vector lying in 2-face $t_{j,k}$, rooted at a point of e_j and not parallel to it. Hence, the first condition that ensures that the movement occurs within the preimage $\pi^{-1}q$ can be written

$$\sum_{i=1}^{n} \lambda_i e_i + \mu v_{j,k} = 0. \tag{2}$$

The zero here, as will in certain occasions later, indicates the n-vector of zero entries.

To ensure that the value of the linear functional decreases as this movement occurs, the following scalar quantity Ψ must be negative. Regard ψ as a real n^2-vector and decompose it to n column n-vectors $\psi_1, \ldots, \psi_n$, then $\psi = (\psi_1^T, \ldots, \psi_n^T) \in \mathbf{R}^{n^2}$. Then we must have

$$\Psi = \sum_{i=1}^{n} \lambda_i e_i^T \psi_i + \mu v_{j,k}^T \psi_j < 0. \tag{3}$$

The last test is on det$E$, where E is the $n \times n$ real matrix formed by the vectors representing the respective edges ie. $E_{t,i} = e_{i,t}, 1 \leq t \leq n$. If the determinant vanishes, then no further test is required because the given combination, even if it forms a Mixed Volume cell, makes a zero contribution. Hence the third condition is written

$$\det E \neq 0. \tag{4}$$

We now manipulate these three conditions to a more amenable form. The first observation is that (2) is a system of n polynomials, from which we could derive expressions for all λ_i's, in terms of μ. Let λ represent the n-dimensional real vector of the λ_i's. Assuming that E is non-singular, (2) is equivalent to

$$E\lambda + \mu v_{j,k} = 0 \Leftrightarrow \lambda = -\mu E^{-1} v_{j,k}.$$

We denote the inverse of E by F and the i-th *row* of F by f_i. Then

$$\lambda_i = -\mu f_i^T v_{j,k} \Leftrightarrow \lambda_i = -\mu v_{j,k}^T f_i.$$

Now we substitute into (3) and eliminate the λ_i's.

$$\begin{aligned} \Psi &= \sum_{i=1}^{n} (-\mu v_{j,k}^T f_i) e_i^T \psi_i + \mu v_{j,k}^T \psi_j \\ &= \mu v_{j,k}^T (\psi_j - \sum_{i=1}^{n} f_i e_i^T \psi_i) \end{aligned}$$

Since μ is restricted to be positive, we can substitute the condition that $\Psi < 0$ by an equivalent one:

$$v_{j,k}^T(\psi_j - \sum_{i=1}^{n} f_i e_i^T \psi_i) < 0. \tag{5}$$

Notice that the expression in parentheses is an n-dimensional real vector that does not depend on the particular 2-face of Δ_j, but does depend on j, ie. on which polytope 2-face $t_{j,k}$ belongs. For convenience, let us call this expression ϕ_j:

$$\phi_j = \psi_j - \sum_{i=1}^{n} f_i e_i^T \psi_i.$$

We can now formulate the basic algorithm I for computing the Mixed Volume.

```
1  Select generic linear functional ψ ∈ R^{n^2}.
2  For every combination of edges e_1, ..., e_n such that e_i ∈ Δ_i do
3      Compute F = E^{-1} and obtain detE. If detE = 0 then goto step 2.
4      For j = 1 to n do
5          Compute φ_j = ψ_j − Σ_{i=1}^n f_i e_i^T ψ_i.
6          For each 2-face t_{j,k} cobounding e_j do
7              Select v_{j,k} ∈ t_{j,k} such that v_{j,k} ∦ e_j.
8              If v_{j,k}^T φ_j ≥ 0 then goto step 2.
9              If t_{j,k} is the last 2-face, increment the Mixed Volume by detE and goto step 2.
```

Step 1 in the algorithm can be implemented by a randomized procedure; we do not pursue this issue in detail because this step is simplified below. Step 5 can be made more efficient by checking whether $j = 1$ or not. Only in the former case the algorithm needs to calculate ϕ_j through the given formula; for all $j > 1$ we can instead use $\phi_j = \phi_{j-1} - \psi_{j-1} + \psi_j$. The implementation of step 7 depends on the representation of the original polytopes Δ_i; to define $v_{j,k}$ it suffices to fix its origin on e_j and know one point in $t_{j,k}$ that is not on the edge.

Now we explore some methods for simplifying as well as speeding up algorithm I, which have no effect on the worst-case asymptotic complexity, though. The first observation concerns the linear functional ψ, for which there have been no assumptions so far, and the idea is to exploit some restrictions on it. Suppose that the entries of ψ_1 are infinitely larger in magnitude than those of $\psi_2, \ldots, \psi_n$; formally, every polynomial in the entries of $\psi_2, \ldots, \psi_n$ with real coefficients has a smaller absolute value than that of any polynomial in some of the entries of ψ_1. This relation among the component vectors of ψ is not explicit, for only component ψ_1 is needed in the algorithm; step 5 then becomes

$$\phi_1 = \psi_1 - f_1 e_1^T \psi_1 \text{ and } \phi_j = -f_1 e_1^T \psi_1, \forall j \neq 1.$$

The justification behind using the new formula for the ϕ_j's is that they are only needed in step 8, where a sign is tested. Each of the terms of the corresponding expression involves exactly one entry of ψ and since those of ψ_1 are significantly larger, the respective terms determine the sign. Furthermore, not all of these terms can be zero if ψ_1 and every vector e_i or $v_{j,k}$ has at least one non-zero component. The first requirement can be imposed during selection of ψ_1, while e_i and $v_{j,k}$ are non-zero for any input set.

This idea can be taken one step further to impose some structure on ψ_1. What is needed is any n-vector whose entries are ordered so that each one is infinitesimal compared to any previous one. Let us choose, arbitrarily, the lexicographic vector with entries $\psi_{1,1} \gg \psi_{1,2} \gg \cdots \gg \psi_{1,n}$. Again, we can avoid explicitly coding this relation; in fact, we do not need any entry of ψ in numeric form if we combine steps 5 and 8 and are willing to perform, in one case, $\mathcal{O}(n)$ sign tests instead of one.

A second improvement comes from speeding up the loop at step 6. Intuitively, it suffices to check at most n independent vectors $v_{j,k}$; if one of their inner products is non-negative, the current edge combination is rejected. On the other hand, if the test is passed on a basis of vectors, then it must be passed on the positive combination of any subset of them, thus including any possible n-vector $v_{j,k}$. This trick requires having available, for every edge e_j, a basis of $v_{j,k}$ vectors with common origin on e_j.

These ideas lead to algorithm II, which does not rely on an explicit linear functional, but assumes instead the lexicographic one. Compared to the previous one and another one coming up, this algorithm is entirely deterministic and numeric.

1 For every combination of edges $e_1, \ldots, e_n$ such that $e_i \in \Delta_i$ do
2   Compute the first row f_1 of $F = E^{-1}$ and $\det E$. If $\det E = 0$ then goto step 1.
3   For $j = 1$ to n do
4   For each 2-face $t_{j,k}$ cobounding e_j such that $v_{j,k}$ is in the corresponding basis do
5   Retrieve $v_{j,k} \in t_{j,k}$, $v_{j,k} \not\parallel e_j$.
6   Compute $w = v_{j,k}^T f_1$.
7   If $j = 1$ then compute successive entries of $v_{j,k} - we_1$ until the first non-zero entry or until n are computed; if the last entry computed is non-negative, goto step 1.
            If $j > 1$ then find the first non-zero entry e_{1s} of e_1 once; if $-we_{1s} \geq 0$ goto step 1.
8   If $t_{j,k}$ is the last 2-face, increment the Mixed Volume by $\det E$ and goto step 1.

The central question is whether the exponential complexity factor that expresses the number of edge combinations can be beaten. We cannot answer this in the affirmative, but we describe one approach that leads to a decrease of the worst-case and, mainly, the average-case running-time. One way is to examine all of the edges in one polytope together. Pick Δ_1 to be this special polytope and represent any edge e_1 in it symbolically, by n variables $e_{11}, \ldots, e_{1n}$. We can execute algorithm II up to step 6; to avoid rational expressions in the inverse matrix, we can multiply by the determinant polynomial. Let $f_1' = \det E f_1$, which is real and directly calculated as fast as in the all-numeric case. We are now ready for algorithm III.

1 Let e_1 be a symbolic n-vector.
2 For every combination of edges $e_2, \ldots, e_n$ such that $e_i \in \Delta_i$ do
3   Compute $f_1' = \det E f_1$, where f_1 is the first row of $F = E^{-1}$ and $\det E$;
        if $\det E$ is identically zero then goto step 2.
4   For $j = 2$ to n do
5   For every 2-face $t_{j,k}$ cobounding e_j such that $v_{j,k}$ is in the corresponding basis do
6   Compute $v_{j,k}^T f_1'$; if it is zero or has different sign than $v_{21}^T f_1'$, then goto step 2.
7   For every e_1 in Δ_1
8   If $-(v_{21}^T f_1')(\det E)\, e_{1s} \geq 0$, where e_{1s} is the first non-zero entry of e_1, then goto step 7.
9   For $j = 1$ goto step 4 of algorithm II.

The determinant of E is a polynomial in the symbolic entries of e_1, therefore step 2 checks whether it is identically zero or not. This happens whenever any columns among those indexed 2 to n in E are dependent; compared to the test of the real $\det E$, this one misses only the cases that e_1 is dependent on the rest of the edge vectors.

The second source of average-case speedup is at step 6, where we complete all partial computation that does not involve e_1. Clearly if some $v_{j,k}^T f_1'$, for $j > 1$, is zero, the corresponding expression at step 7 of algorithm II is zero and the current edge combination can be rejected. Or, if two of these expressions have different signs, then one of them gives a positive sign at the above step and again the edge combination is rejected.

When both of these attempts for a shortcut fail, algorithm III reduces to algorithm II but is still preferable with respect to worst-case complexity as seen in the next section.

5 Complexity Analysis

Assuming that the input consists of n point sets, each in $\mathbf{R}^n$, an initial phase must compute the respective convex hulls $\Delta_1, \ldots, \Delta_n$. There exists an algorithm [Ch] that does this in $\mathcal{O}(m^{\lfloor n/2 \rfloor})$ time on a real RAM, where m is an upper bound on the number of input points. An alternative algorithm that is easier to implement, such as Beneath-Beyond as described in [Ed], has algebraic complexity $\mathcal{O}(m^{\lceil n/2 \rceil})$ but has the advantage to compute the hull's volume without any significant extra cost. In what follows, the floor function is understood as being applied to the exponent and is omitted in our bounds.

For algorithms II and III, preprocessing takes $\mathcal{O}(nm^{n/2})$ for constructing n convex hulls, plus the cost of the second stage during which a basis is selected among all vectors $v_{j,k}$, for every edge e_ji. The total number of edges is $\mathcal{O}(nm^2)$ because a polytope of $\mathcal{O}(m)$ vertices has $\mathcal{O}(m^2)$ edges and $\mathcal{O}(m^3)$ 2-faces. Hence, the second stage, which is a linear algebra procedure, takes polynomial time and is dominated by the first stage.

For comparison purposes, we briefly examine the complexity of the Inclusion-Exclusion approach. It computes the Minkowski sum, its convex hull and the hull's volume for each of the 2^n subsets of the set of the original polytopes. In the worst case, every Minkowski sum is such that all points lie on the convex hull surface; then, counting only the cost of the convex hull computations gives a total complexity proportional to

$$\sum_{s=1}^{n} \binom{n}{s} (m^t)^{n/2} = \sum_{s=1}^{n} \binom{n}{s} (m^{n/2})^t = \Theta((m^{n/2})^n) = \Theta(m^{n^2/2}).$$

We analyze algorithm II. The outer loop at step 1 is repeated l^n times, where l is the maximum number of edges in any of the original polytopes and step 2 takes $\mathcal{O}(n^{2.376})$ time [CoWi]. The loops at step 3 and 4 are each repeated n times and the sign test at 7 takes $\mathcal{O}(n)$ because it it linear-time for $j=1$ and constant-time for every $j>1$. The total asymptotic complexity is then

$$\mathcal{O}(l^n n^{2.376} + l^n n^3) = \mathcal{O}(l^n n^3).$$

Let m denote the maximum cardinality of the input point sets; in the worst case, the original polytopes have $\mathcal{O}(m)$ vertices, $\mathcal{O}(m^2)$ edges and $\mathcal{O}(m^3)$ 2-faces. Then the complexity becomes $\mathcal{O}(m^{2n}n^3)$, and dominates the complexity for preprocessing, where $m > n$ since the original polytopes lie in n-dimensional space.

Algorithm III has complexity $\mathcal{O}(l^{n-1}n^3)$ up to step 6. The back end that examines specific edges e_1 has a total cost $\mathcal{O}(l^n n^2)$ which is dominant. In terms of m, the worst-case complexity becomes $\mathcal{O}(m^{2n}n^2)$.

6 Future Directions

One source of possible improvement is to compress successive calculations of the inverse matrix F and the determinant of E, as only a single column changes from one instance of E to the next. Letting more than one edge be symbolic does not yield any immediate improvement. Of course, the main focus lies on whether we can avoid looking at every single combination of edges. A radically new approach might be necessary so that the Mixed Volume is not decomposed to cells produced by edge combinations, for these cells are, in a sense, too small and placed in a disorderly manner. Yet another prospect is randomization.

Acknowledgments

John Canny has substantially helped in this project as my advisor and instructor of course CS 275 during the spring of 1992, for which this project was undertaken. In addition, Bernd Sturmfels suggested looking at the product polytope.

References

[Be] D.N. Bernshtein, The Number of Roots of a System of Equations, translated from *Funktsional'nyi Analiz i Ego Prilozheniya*, Vol. 9, No. 3, pp. 1-4, Jul.-Sep. 1975.

[Br] A. Brøndsted, *An Introduction to Convex Polytopes*, Springer-Verlag, New York, 1983.

[CaRo] J. Canny and J.M. Rojas, An Optimal Condition for Determining the Exact Number of Roots of a Polynomial System, *Proc. of ACM Int. Symp. Symb. Algebr. Comp.*, pp. 96-102, Bonn, July 1991.

[Ch] B. Chazelle, An Optimal Convex Hull Algorithm and New Results on Cuttings, *Proc. 32nd Annual IEEE FOCS*, San Juan, pp. 29-38, Oct. 1991.

[CoWi] Coppersmith D. and S. Winograd, Matrix Multiplication via Arithmetic Progressions, *J. Symb. Comp.*, 9, pp. 251-280, 1990.

[Ed] H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer-Verlag, Heidelberg, 1987.

[Kh] A.G. Khovanskii, Newton Polyhedra and the Genus of Complete Intersections, translated from *Funktsional'nyi Analiz i Ego Prilozheniya*, Vol. 12, No. 1, pp. 51-61, Jan.-Mar. 1978.

[Ku] A.G. Kushnirenko, Newton Polytopes and the Bezout Theorem, translated from *Funktsional'nyi Analiz i Ego Prilozheniya*, Vol. 10, No. 3, Jul.-Sep. 1976.

[Ma] E. Mahalingam, Master's thesis, UC Berkeley, Fall 1992.

[St] B. Sturmfels, Fiber Polytopes: A Brief Overview, Manuscript, 1991.