A note on the minimal volume of almost cubic parallelepipeds*†

Daniele Micciancio‡

Abstract

We prove that the best way to reduce the volume of the n-dimensional unit cube by a linear transformation that maps each of the main vertices $\vec{e}_i$ to a point within distance $\varepsilon < \sqrt{(1/n) - (1/n^2)}$ from $\vec{e}_i$ is to shorten all edges by a factor $(1 - \varepsilon)$. In particular, the minimal volume of such an almost cubic parallelepiped is $(1 - \varepsilon)^n$. This problem naturally arises in the construction of lattice based one-way functions with worst-case/average-case connection.

1 Introduction

In this note we consider the following geometric problem. Let

$$\mathcal{C} = \left\{ \sum_{i=1}^{n} c_i \vec{e}_i : 0 \leq c_i \leq 1 \right\}$$

be the unit hypercube in n-dimensional Euclidean space, and deform it into an almost cubic parallelepiped

$$\mathcal{P} = \left\{ \sum_{i=1}^{n} c_i (\vec{e}_i + \vec{x}_i) : 0 \leq c_i \leq 1 \right\}$$

where a small perturbation $\vec{x}_i$ is added to each vertex $\vec{e}_i$. We want to determine the minimal volume of $\mathcal{P}$, when perturbations have length at most $\|\vec{x}_i\| \leq \varepsilon$. Using matrix notation, the problem can be reformulated as determining the value of the function

$$f_n(\varepsilon) = \min\{|\det(I + X)| : \max_i \|\vec{x}_i\| \leq \varepsilon\}$$

where $\vec{x}_i$ (for $i = 1, \ldots, n$) are the columns of matrix X. (Notice that f_n is well defined because $|\det(I+X)|$ is a continuous function of X, and the constraint $\max_i \|\vec{x}_i\| \leq \varepsilon$ defines a compact set.)

This is a natural geometric question, and it often arises in problems involving point lattices and volume estimations. For example, Cai and Nerurkar [3] use bounds on the volume of $\mathcal{P}$ to construct cryptographic one-way functions which are as hard to invert (on the average) as the worst case complexity of certain lattice approximation problems, improving on previous results of Ajtai [1]. (See also [4] for the best currently known construction.) Notice that since the columns of X have norm bounded by ε, all eigenvalues of X are at most $\varepsilon\sqrt{n}$ in absolute value. It follows that the eigenvalues of $I + X$ are at least $1 - \varepsilon\sqrt{n}$, and therefore $\det(I + X) \geq (1 - \varepsilon\sqrt{n})^n$ for all $\varepsilon \leq 1/\sqrt{n}$. This immediately gives lower bound

$$f_n(\varepsilon) \geq (1 - \varepsilon\sqrt{n})^n \tag{1}$$

for $\varepsilon \leq 1/\sqrt{n}$. (See Figure 1.) An asymptotically better solution (for small values of ε) is given in [3] which proves lower bound

$$f_n(\varepsilon) \geq (1 - \varepsilon n) \tag{2}$$

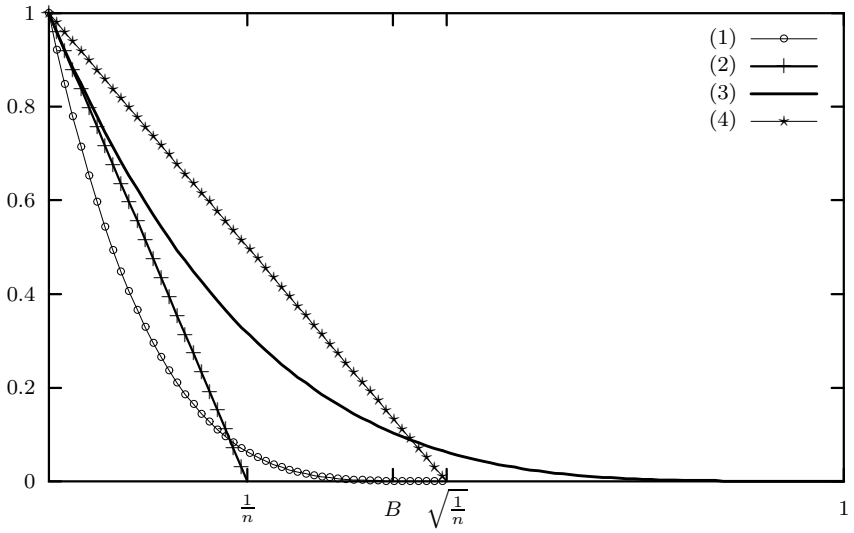

Figure 1: Bounds on $f_n(\varepsilon)$ with $n = 4$. (1) and (2) are lower bounds. (3) and (4) are upper bounds. Bound (3) is met with equality for all $\varepsilon < B = \sqrt{\frac{1}{n} - \frac{1}{n^2}}$.

for all $\varepsilon \leq 1/n$. (See Figure 1.) So, the eigenvalue bound (1) is not optimal for values of ε close to 0. Notice that the related question of computing the maximal volume has a very simple answer: here bounds on the eigenvalues of $I + X$ give inequality

$$|\det(I + X)| \leq (1 + \varepsilon)^n$$

which is optimal because matrix $X = \varepsilon I$ achieves the maximum $|\det(I + X)| = \det(I + X) = (1 + \varepsilon)^n$. This solution corresponds to the geometric intuition that the best way to increase the volume of a hypercube is to lengthen each edge by ε, without changing the direction of the edges. A similar solution to the minimization problem (setting $X = -\varepsilon I$) gives upper bound

$$f_n(\varepsilon) \leq (1 - \varepsilon)^n. \tag{3}$$

Notice that upper bound (3) is strictly bigger than lower bound (2) for all values of $\varepsilon > 0$. (See Figure 1.) So, it is natural to ask whether bound (2) can be improved, and, in particular, bound (3) holds with equality. In other words, we ask if the best way to reduce the volume of a hypercube is to shorten each edge by ε, or it is possible to do better by changing the angles between the edges. This question is explicitly asked by Cai in [2], where the author also points out possible difficulties in proving equality $f_n(\varepsilon) = (1 - \varepsilon)^n$ in high dimensions. In fact, for any fixed ε, inequality (3) is strict for all sufficiently large n. (See below for details.)

It immediately follows from (1) that lower bound (2) is not optimal for all $\varepsilon \in [1/n, 1/\sqrt{n})$. (In fact, by a continuity argument, lower bound (2) is not optimal when $\varepsilon \in (c, 1/\sqrt{n})$ for some c strictly smaller than $1/n$.) Interestingly, one can also prove upper bound

$$f_n(\varepsilon) \leq 1 - \varepsilon\sqrt{n}, \tag{4}$$

attained setting all entries of X to $-\varepsilon/\sqrt{n}$. Upper bound (4) shows that no nontrivial lower bound exists for $\varepsilon \geq 1/\sqrt{n}$, i.e., $f_n(\varepsilon) = 0$ for all $\varepsilon \geq 1/\sqrt{n}$.

*Research supported in part by NSF CAREER Award CCR-0093029

†A formatted version of this paper appears in Discrete and Computational Geometry 29(1):133-138, 2002.

‡Department of Computer Science and Engineering, University of California, San Diego. Mail Code: 0114. 9500 Gilman Drive, La Jolla, CA 92093-0114, USA. Email: daniele@cs.ucsd.edu

By continuity, (4) implies that bound (3) is strict for values of ε sufficiently close to $1/\sqrt{n}$, i.e., if ε is not too small, there are better ways to reduce the volume of a hypercube than simply shortening the edges. It remains to be determined where, in the triangular region between bounds (2), (4) and $f_n(\varepsilon) > 0$, the graph of $f_n(\varepsilon)$ exactly lies.

It should be noted that for small values of ε bounds (2) and (3) are the same up to lower order terms, so improving lower bound (2) is unlikely to have substantial implication in most computational settings, including the proof of security of lattice based one-way functions from [3]. Still, it would be nice to determine if the intuitive solution of bound (3) is optimal at least in a sufficiently small neighborhood of the origin.

In this paper we resolve this question in the affirmative, and prove that (3) in fact holds with equality for most of the interval $[0, 1/\sqrt{n})$, as stated in the following theorem.

Theorem 1 *The minimal volume of an almost cubic parallelepiped satisfies $f_n(\varepsilon) = (1-\varepsilon)^n$ for all*

$$\varepsilon < \sqrt{\frac{1}{n} - \frac{1}{n^2}}.$$

Moreover, $0 < f_n(\varepsilon) \leq \min\{(1-\varepsilon)^n, 1-\varepsilon\sqrt{n}\}$ for all $\varepsilon < \sqrt{1/n}$ and $f_n(\varepsilon) = 0$ for all $\varepsilon \geq \sqrt{1/n}$.

2 Proof of the theorem

We already discussed bounds $f_n(\varepsilon) \leq (1-\varepsilon)^n$ and $f_n(\varepsilon) \leq 1-\varepsilon\sqrt{n}$. In the rest of this section we prove the lower bound $f_n(\varepsilon) \geq (1-\varepsilon)^n$ for all $\varepsilon < \sqrt{(1/n) - (1/n^2)}$. Notice that (by the argument used in the proof of bound (1)) if $\varepsilon < 1/\sqrt{n}$, then all eigenvalues of $I+X$ are strictly positive. Therefore, $\det(I+X) > 0$ and $|\det(I+X)| = \det(I+X)$. This proves that for all interesting values of ε, function $f_n(\varepsilon)$ can be equivalently redefined as the minimum of the determinant $\det(I+X)$, without taking the absolute value.

Let X be a matrix with $\|\vec{x}_i\| \leq \varepsilon$ for all i such that $\det(I+X)$ is minimized. (Remember that for $\varepsilon < \sqrt{1/n}$, $\det(I+X) > 0$, and therefore $|\det(I+X)| = \det(I+X)$.) We want to prove that X is diagonal. It will follow that

$$f_n(\varepsilon) = \det(I+X) = \prod_i (1+x_{i,i}) \geq (1-\varepsilon)^n.$$

Assume without loss of generality that the diagonal elements of X form a monotonically increasing sequence $x_{1,1} \leq x_{2,2} \leq \ldots \leq x_{n,n}$ (this can always be achieved changing the order of basis vectors $\vec{e}_1, \ldots, \vec{e}_n$) and let T be the diagonal matrix with entries $x_{1,1}, \ldots, x_{n,n}$.

For any $i \leq n$, fix the value of $\vec{x}_j$ for all $j \neq i$, and consider $\det(I+X)$ as a function of $\|\vec{x}_i\| \leq \varepsilon$. For the volume to be minimum, $\vec{e}_i + \vec{x}_i$ must be as close as possible to the hyper-plane spanned by vectors $\vec{e}_j + \vec{x}_j$ (for $j \neq i$). Therefore $\vec{x}_i$ is a vector of length $\|\vec{x}_i\| = \varepsilon$ orthogonal to all $\vec{e}_j + \vec{x}_j$, for $j \neq i$. (We recall that since $\det(I+X)$ is strictly positive, vector $\vec{e}_j + \vec{x}_j$ cannot be on or beyond the hyper-plane, and $\vec{e}_j + \vec{x}_j$ is closest to the plane when the length of $\vec{x}_j$ is maximal.) Using matrix notation, this can be expressed as

$$(I+X)^T X = \varepsilon^2 I + T.$$

Rearranging the terms

$$X = \varepsilon^2 I + T - X^T X.$$

This proves that $X = X^T$ is a symmetric matrix. We want to show that X is actually diagonal. Since X is symmetric, there exist linearly independent eigenvectors $U = [\vec{u}_1, \ldots, \vec{u}_n]$ such that $XU = UD$ for some diagonal matrix D. Since $T = X^2 + X - \varepsilon^2 I$ is polynomial in X, the columns of U are also eigenvectors for T. But T is diagonal, so, by appropriately permuting the columns of U, we can assume that $TU = UT$. Details follow. Let $\lambda_1, \ldots, \lambda_l$ be the eigenvalues of T, and let n_i be the multiplicity of λ_i for all $i = 1, \ldots, l$. Remember that matrix T is the diagonal of matrix X, and the columns of the latter have been sorted according to the value of diagonal elements $x_{i,i}$. It follows that T can be written as a block diagonal matrix with blocks $T_i = \lambda_i I$ (for $i = 1, \ldots, l$), where the ith block has dimension $n_i \times n_i$. Since vectors $\vec{u}_j$ are

eigenvectors for T, for each $j = 1, \ldots, n$ there is a $i_j \in \{1, \ldots, l\}$ such that $T\vec{u}_j = \lambda_{i_j}\vec{u}_j$. Moreover, since vectors $\vec{u}_j$ form a basis, for every i there are exactly n_i vectors $\vec{u}_j$ such that $T\vec{u}_j = \lambda_{i_j}\vec{u}_j$. Therefore, if the columns $\vec{u}_j$ are sorted according to their eigenvalues λ_{i_j}, we get $TU = UT$. (Notice that permuting the columns of U also permutes the diagonal elements of D, but this change is clearly immaterial.)

Since eigenvectors with different eigenvalues are orthogonal, matrix U is also block-diagonal. To see this, divide matrix U into blocks $U_{i,j}$,of size $n_i \times n_j$. From matrix equation $TU = UT$, and the block structure of T, we get $\lambda_i U_{i,j} = \lambda_j U_{i,j}$ for all $1 \le i, j \le l$. Since $\lambda_i \neq \lambda_j$ for all $i \neq j$, this proves that $U_{i,j} = 0$ for all $i \neq j$, i.e., U is block diagonal with blocks $U_{i,i}$.

Now consider matrix D. Since D is diagonal, we can write it in block diagonal form, where each block D_i (for $i = 1, \ldots, l$) is an $n_i \times n_i$ diagonal matrix. (Notice however that the blocks of D are not necessarily scalar, because eigenvalues of D_i may correspond to the two distinct solutions of equation $x^2 + x - \varepsilon^2 = \lambda_i$.)

Finally, consider the original matrix X. We know that $XU = UD$, or, equivalently, $X = UDU^{-1}$. From the block diagonal structure of U and D we get that X is also block diagonal, with blocks $X_i = U_{i,i}D_iU_{i,i}^{-1}$ of size $n_i \times n_i$ for $i = 1, \ldots, l$. Moreover, since T is the diagonal of X, each block X_i is constant on the diagonal, with all diagonal elements equal to λ_i. We need to prove that each block X_i is in fact a diagonal matrix, i.e., $X_i = \lambda_i I$.

We first consider the case when X consists of a single block, i.e., $l = 1$ and the diagonal of X is a scalar matrix $T = tI$. Clearly, it must be $|t| \le \varepsilon$. We already proved that $X^T = X$ is symmetric and $X = (\varepsilon^2 + t)I - X^2$. We want to prove that $X = tI$. From $|t| \le \varepsilon$, we get $\varepsilon^2 + t + 1/4 \ge (\varepsilon - 1/2)^2 > 0$. If we define $Z = (X + (1/2)I)/\sqrt{\varepsilon^2 + t + 1/4}$, we see that X can be expressed as $X = -(1/2)I + cZ$ where Z is a symmetric matrix satisfying $Z^2 = I$ and constant on the diagonal. Let sI be the diagonal of Z. Notice that the trace of Z is ns. The trace is equal to the sum of the eigenvalues, and since $Z^2 = I$, all eigenvalues of Z are 1 or -1. So, the trace is $ns = 2a - n$ where a is the multiplicity of the $+1$ eigenvalue. Let us compute the lengths of the columns of X. The squared lengths of these vectors are the diagonal elements of $X^TX = X^2 = (1/4)I + c^2Z^2 - cZ$, i.e., $1/4 + c^2 - cs$ which is minimized when $c = s/2$ with minimum

$$\varepsilon^2 \ge \frac{1 - s^2}{4}.$$

Substituting $s = 2a/n - 1$ we get

$$\varepsilon^2 \ge \frac{1}{4}\left(1 - \frac{4a^2}{n^2} - 1 + \frac{4a}{n}\right) = \frac{a}{n} - \frac{a^2}{n^2} = \frac{a}{n}\left(1 - \frac{a}{n}\right).$$

Unless $a = 0$ or $a = n$, this expression is at least $(1/n)(1 - (1/n))$, contradicting the assumption $\varepsilon^2 < (1/n) - (1/n^2)$. So, all eigenvalues of Z are equal and $Z = \pm I$. This proves that $X = -(1/2)I + cZ$ is a diagonal matrix, and therefore $X = T$.

Now assume T is not scalar, and X consists of more than one block. Let $X_1, \ldots, X_l$ be the blocks of X. Since the determinant of $I + X$ equals the product of the determinants of the blocks $I + X_i$, in order to minimize $\det(I + X)$ we need to minimize the determinant $\det(I + X_i)$ of each block independently. We prove that all blocks X_i are diagonal. Let X_i be any of the blocks and let us minimize $\det(I + X_i)$ subject to the constraint that all columns of X_i have length bounded by $\varepsilon < (1/n) - (1/n^2)$. If the size n_i of the block is 1, then X_i is clearly diagonal. Otherwise, $n_i \ge 2$ and $(1/n) - (1/n)^2 < (1/n_i) - (1/n_i^2)$. It follows from the proof of the case when X consists of a single block that $\det(I + X_i)$ is minimized when $X_i = -\varepsilon I$. This completes the proof that X is a diagonal matrix, and therefore $X = -\varepsilon I$.

3 Conclusion

For every dimension n, if ε is sufficiently close to $1/\sqrt{n}$ there are better ways to reduce the volume of an n-dimensional hypercube other than shortening all edges by ε. In particular, if $\varepsilon \ge 1/\sqrt{n}$, then the volume can be reduced to zero by appropriately changing the angles between the edges. We proved that as soon as ε gets smaller than $\sqrt{(1/n) - (1/n^2)}$, then the minimal volume is $f_n(\varepsilon) = (1 - \varepsilon)^n$, i.e., the optimum is

achieved by shortening each edge by ε. This leaves open the question of what is the value of $f_n(\varepsilon)$ when $\varepsilon \in [\sqrt{(1/n) - (1/n^2)}, \sqrt{1/n}]$. We leave the complete characterization of the value of $f_n(\varepsilon)$ as an open problem.

References

[1] M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing*, pages 99–108, Philadelphia, Pennsylvania, 22–24 May 1996. ACM.

[2] J.-Y. Cai. Some recent progress on the complexity of lattice problems. In *Proceedings of the Fourteenth Annual IEEE Conference on Computational Complexity*, pages 158–178, Atlanta, Georgia, 4–6 Mar. 1999. IEEE.

[3] J.-Y. Cai and A. P. Nerurkar. An improved worst-case to average-case connection for lattice problems (extended abstract). In *38th Annual Symposium on Foundations of Computer Science*, pages 468–477, Miami Beach, Florida, 20–22 Oct. 1997. IEEE.

[4] D. Micciancio. Improved cryptographic hash functions with worst-case/average-case connection (extended abstract). In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, Montreal, Canada, 19–21 May 2002. ACM. pp. 609–618.