

Finding the vertices of the convex hull, even unordered, takes $\Omega(n \log n)$ time—a proof by reduction from ε -closeness

Herman Haverkort*

28 November 2018

Abstract

We consider the problem of computing, given a set S of n points in the plane, which points of S lie on the boundary of the convex hull of S . This paper provides a simple proof that the time complexity of this problem in the algebraic decision tree model is $\Omega(n \log n)$.

The problem In the algebraic decision tree model of computation, it takes $\Omega(n \log n)$ steps in the worst case to compute the convex hull of a set of n points in the plane. In text books [2, 3] this is proven by reduction from sorting as follows. Consider a set $A = a_1, \dots, a_n$ of n distinct real numbers that need to be sorted. Now let $L = L_1, \dots, L_n$ be the set of points given by $L_i = (a_i, a_i^2)$. The convex hull of L , as a closed counterclockwise polygonal chain, contains the points of L in order from left to right. Therefore, by computing the convex hull and then reading the x -coordinates of its vertices in order, we can obtain A in sorted order. Hence, lower bounds for sorting also apply to the computation of convex hulls.

However, if we do not require the convex hull to be produced as a counterclockwise sequence of vertices (or edges), but settle for obtaining its vertices in arbitrary order, may we then be able to do so in $o(n \log n)$ time, that is, less than $\Omega(n \log n)$ time? The answer is no, as we will see now, using a reduction from the ε -closeness problem.

The proof In its simplest form, the algebraic decision tree model captures computer programs that take as input a sequence of real numbers, and consist of a number of instructions I_1, I_2, I_3, \dots , each of which is of one of three types: 1. evaluate a polynomial function of the input variables, and use the sign of the result to choose which instruction to execute next; 2. output “yes”; 3. output “no”. The polynomial functions must have degree at most d , for some constant d : that is, each function is a sum of terms, where each term is a product of a constant and at most d factors, where each factor is an input variable.

The ε -closeness problem is the following: given an unordered multiset A of real numbers a_1, \dots, a_n and a positive real number ε , decide whether A contains two numbers that differ less than ε . In other words: is there a pair of indices $i \neq j$ such that $0 \leq a_j - a_i < \varepsilon$, or equivalently, $a_i \leq a_j < a_i + \varepsilon$? Any algebraic decision tree algorithm that solves the ε -closeness problem takes $\Omega(n \log n)$ time in the worst case [3].

Let us define the *any-point-inside* problem as follows: given an unordered multiset S of n points in the plane (given as $2n$ real coordinates), decide whether any point of S lies in the interior of the convex hull of S . Let $api(n)$ be the number of steps it takes in the worst case to solve this problem in the algebraic decision tree model.

*Institut für Informatik, Universität Bonn, Germany, cs.herman@haverkort.net




Figure 1: Left: if all numbers in A are at least ϵ apart, then all points of L and T are on the boundary of the convex hull. Right: if $a_i < a_j < a_i + \epsilon$, then T_i lies in the interior of the triangle $L_i L_j T_j$.

Lemma 1. *Given a multiset of real numbers $A = a_1, \dots, a_n$ and a real number ϵ , we can decide in $\text{api}(2n)$ time whether there are two indices $i \neq j$ such that $0 < a_j - a_i < \epsilon$, or equivalently, $a_i < a_j < a_i + \epsilon$.*

Proof. Let $S(A, \epsilon)$ be $L \cup T$, where $L = L_1, \dots, L_n$ is the set of points given by $L_i = (a_i, a_i^2)$, and $T = T_1, \dots, T_n$ is the set of points given by $T_i = (a_i + \epsilon/2, (a_i + \epsilon/2)^2 + \epsilon^2/4)$. We will see shortly that $S(A, \epsilon)$ is constructed such that all of its points will appear on the boundary of its convex hull, unless some point T_i is "hidden" by a nearby point L_j where $0 < a_j - a_i < \epsilon$. Thus, to determine whether there are two indices $i \neq j$ such that $0 < a_j - a_i < \epsilon$, we simply solve the any-point-inside problem on $S(A, \epsilon)$ in $\text{api}(2n)$ time.

To prove the correctness of this reduction, we start with the following observations. For $i \in \{1, \dots, n\}$, let R_i be the point $(a_i + \epsilon, (a_i + \epsilon)^2)$. Note that L_i and R_i lie on the parabola $y = x^2$, which we call the *outer* parabola. The point T_i lies exactly half-way between L_i and R_i on the parabola $y = x^2 + \epsilon^2/4$, which we call the *inner* parabola. Moreover, the slope of the inner parabola at T_i is $2a_i + \epsilon$, which is exactly the slope of the line ℓ_i through L_i, T_i and R_i , so ℓ_i is a tangent to the inner parabola.

Now I claim the following: (i) if there are no i and j such that $a_i < a_j < a_i + \epsilon$, then all $2n$ points of $S(A, \epsilon)$ are on the boundary of the convex hull of $S(A, \epsilon)$; (ii) if there are i and j such that $a_i < a_j < a_i + \epsilon$, then $S(A, \epsilon)$ contains at least one point, namely T_i , that is not on the boundary of the convex hull of $S(A, \epsilon)$. Both claims are illustrated in Figure 1.

Proof of (i): assume there are no i and j such that $a_i < a_j < a_i + \epsilon$. We will show that both L_i and T_i lie on the boundary of the convex hull for each i . Recall that the line ℓ_i through L_i, T_i and R_i is a lower tangent to the inner parabola and intersects the outer parabola in L_i and R_i . So, if $S(A, \epsilon)$ would contain any points that lie strictly below ℓ_i , these points would have to be points L_j on the outer parabola strictly between L_i and R_i . However, looking at the x -coordinates of these points, we would then find $a_i < a_j < a_i + \epsilon$, contradicting our assumptions. It follows that both L_i and T_i lie on the boundary of the convex hull, as there is a closed half-plane with L_i and T_i on the boundary and all other points of $S(A, \epsilon)$ inside.

Proof of (ii): we will show that T_i lies in the interior of the triangle $L_iL_jT_j$. First, because $a_i < a_j < a_i + \varepsilon$, the point L_j lies on the outer parabola strictly between L_i and R_i , so L_j lies strictly under the line ℓ_i through L_i , T_i , and R_i , and vice versa, T_i lies strictly above the line through L_i and L_j . Second, recall that the line through L_j and T_j is a lower tangent to the inner parabola, touching it in T_j , so T_i lies strictly above it. Third, the line through L_i and T_i is a lower tangent to the inner parabola, so T_j lies strictly above it, and vice versa, T_i lies strictly below the line through L_i and T_j . It follows that T_i lies in the interior of the triangle $L_iL_jT_j$, and thus, not on the boundary of the convex hull of $S(A, \varepsilon)$.

Thus, $S(A, \varepsilon)$ contains a point that lies in the interior of the convex hull of $S(A, \varepsilon)$ if and only if there are i and j such that $a_i < a_j < a_i + \varepsilon$. This proves the correctness of the reduction and thus proves the lemma. \square

Lemma 2. *Given a multiset of real numbers $A = a_1, \dots, a_n$ and a real number ε , we can decide ε -closeness in $\text{api}(2n)$ time, that is, we can decide in $\text{api}(2n)$ time whether there are two indices $i \neq j$ such that $0 \leq a_j - a_i < \varepsilon$, or equivalently, $a_i \leq a_j < a_i + \varepsilon$.*

Proof. We first decide, in $\text{api}(2n)$ time by Lemma 1, whether there are two indices $i \neq j$ such that $0 < a_j - a_i < \varepsilon$. If this is the case, we answer “yes”. Otherwise, we have now established that each pair of numbers in A differs by either zero, or at least ε , and what is left to decide is whether there is any pair with difference zero. Let $A' = a'_1, \dots, a'_n$ be given by $a'_i = a_i + \varepsilon i / (2n)$. Thus, $|a'_j - a'_i|$ and $|a_j - a_i|$ differ by less than $\varepsilon/2$ for any i and j . Observe that, given that there are no i, j such that $0 < a_j - a_i < \varepsilon$, we now have $i \neq j$ and $a_j - a_i = 0$ if and only if $0 < |a'_j - a'_i| < \varepsilon/2$. By Lemma 1, we can decide whether there is any pair i, j for which the latter is the case in $\text{api}(2n)$ time. \square

Since we have an $\Omega(n \log n)$ -time lower bound for ε -closeness, we must now conclude $\text{api}(2n) = \Omega(n \log n)$, so $\text{api}(n) = \Omega(\frac{n}{2} \log \frac{n}{2}) = \Omega(n \log n)$, and we obtain the following theorem:

Theorem 1. *Given a multiset S of n points in the plane, any algebraic decision tree algorithm takes $\Omega(n \log n)$ steps in the worst case to decide whether S contains any point that lies in the interior of the convex hull of S .*

If, in the ε -closeness problem, we replace $<$ by \leq , then an $\Omega(n \log n)$ -time lower bound can be constructed in the same way as for the original ε -closeness problem. Furthermore, we can adapt the proof of Lemma 1 to prove (i) if there are no i and j such that $a_i < a_j \leq a_i + \varepsilon$, then all $2n$ points of $S(A, \varepsilon)$ are extreme points (vertices) of the convex hull of $S(A, \varepsilon)$; (ii) if there are i and j such that $a_i < a_j \leq a_i + \varepsilon$, then $S(A, \varepsilon)$ contains at least one point, namely T_i , that is not an extreme point of the convex hull of $S(A, \varepsilon)$. Thus, we also get the $\Omega(n \log n)$ -time lower bound for the following problem:

Theorem 2. *Given a multiset S of n points in the plane, any algebraic decision tree algorithm takes $\Omega(n \log n)$ steps in the worst case to decide whether S is in convex position (that is, whether S is exactly the set of vertices of the convex hull of S).*

To output the vertices of the convex hull of S , we would need to extend the model of computation so that it admits other output than just “yes” or “no”. No matter how exactly we do that, it is clear that we cannot output the vertices of the convex hull in less than $\Omega(n \log n)$ time, if we cannot even decide whether or not all of S should be output in less than $\Omega(n \log n)$ time.

Other proofs Preparata and Shamos provide another proof that it takes $\Omega(n \log n)$ steps to decide whether a set of points is in convex position (Theorem 3.3 in Section 3.2 in their book [3]). Their proof is based directly on the lower bounds for membership tests for a set W in high-dimensional space with $n!$ connected components—the same lower bounds that also underly the lower bounds for ε -closeness that we used above. The core of their proof is therefore an analysis of the number of connected components of W , the subset of possible inputs that corresponds to sets of points in convex position ([3], p102–103). Their key argument considers inputs in convex position that cannot be continuously transformed into each other without passing through a configuration in which three points are collinear, so that it is not in convex position¹. Thus, the proof by Preparata and Shamos is a bit more specific than ours: they prove our Theorem 2, but not our Theorem 1.

Kirkpatrick and Seidel [1] state the following (Theorem 5.3): given a set S of n distinct points and a natural number $h \leq n$, any fixed-order algebraic decision tree algorithm requires, in the worst case, $\Omega(n \log h)$ steps to verify that the convex hull of S has h vertices. When h is polynomial in n , this bound is equivalent to $\Omega(n \log n)$. Clearly, if we could compute in $o(n \log n)$ time which points of S are vertices of the convex hull, then we could also count them in $o(n \log n)$ time. Thus, the theorem by Kirkpatrick and Seidel also seems to imply an $\Omega(n \log n)$ lower bound for the problem of computing the vertices of the convex hull. Like the proof of our Lemma 2, the proof by Kirkpatrick and Seidel uses a reduction in which multiple points that represent the same number are perturbed, so that all of these points become distinct and only one of them appears on the boundary of the convex hull. However, there is a step in their approach which I do not quite understand and of which I have not been able to verify the correctness yet. In particular, their reduction is only correct if the perturbation is “small enough”, and I do not understand how their reduction algorithm can ensure that this is the case in $o(n \log n)$ time.

References

- [1] D. G. Kirkpatrick and R. Seidel. The ultimate planar convex hull algorithm? *SIAM J. Comput.* 15:287–299, 1986.
- [2] R. Klein. *Algorithmische Geometrie*, 2nd edition, Springer, 2005.
- [3] F. P. Preparata and M. I. Shamos, *Computational Geometry*, Springer, 1985.

¹The reader who wants to verify the details of the proof should beware of minor typing mistakes where permutations on N integers are applied to integers from 0 to $2N - 1$ or from 0 to $N^2 - 1$.