A tight lower bound for computing the diameter of a 3D convex polytope

Hervé Fournier * Antoine Vigneron †

April 2, 2007

Abstract

The diameter of a set P of n points in $\mathbb{R}^d$ is the maximum Euclidean distance between any two points in P. If P is the vertex set of a 3-dimensional convex polytope, and if the combinatorial structure of this polytope is given, we prove that, in the worst case, deciding whether the diameter of P is smaller than 1 requires $\Omega(n \log n)$ time in the algebraic computation tree model. It shows that the $O(n \log n)$ time algorithm of Ramos for computing the diameter of a point set in $\mathbb{R}^3$ is optimal for computing the diameter of a 3-polytope. We also give a linear time reduction from Hopcroft's problem of finding an incidence between points and lines in $\mathbb{R}^2$ to the diameter problem for a point set in $\mathbb{R}^7$.

Keywords: Computational geometry; Lower bound; Diameter; Convex polytope; Hopcroft's problem

1 Introduction

The diameter problem for a set P of n points in $\mathbb{R}^d$ is to compute the largest distance between any two points in P. In other words, if we denote by $\mathrm{d}(\cdot,\cdot)$ the Euclidean distance in $\mathbb{R}^d$, it consists in finding $\mathrm{diam}(P) = \max\{\mathrm{d}(x,y) \mid x, y \in P\}$. It is a fundamental problem in computational geometry and has been studied extensively [2, 5, 12, 13, 16, 17]. If $P \subset \mathbb{R}^2$, then its diameter can be computed in $O(n \log n)$ time [16], which is optimal in the algebraic computation tree model [1, 3]. The three dimensional case remained open for a much longer time, but eventually Clarkson and Shor [7] designed an optimal $O(n \log n)$ time randomized algorithm to compute the diameter of a set of n points in $\mathbb{R}^3$, and Ramos [17] found a deterministic counterpart.

The $\Omega(n \log n)$ lower bound for computing the diameter of $P \subset \mathbb{R}^2$ can be broken if P is given as the sequence of the vertices of a convex polygon sorted along its boundary, in which case an $O(n)$ time algorithm is known [16]. Our main result (Theorem 7) is to show that the same speed-up cannot be achieved in $\mathbb{R}^3$, when P is the vertex set of a convex polytope, and the combinatorial structure of this polytope is given. In the worst case $\Omega(n \log n)$ time is required to compute the diameter of P. More precisely, we show that deciding whether the diameter of P is smaller than 1 requires an algebraic computation tree with depth $\Omega(n \log n)$.

*Laboratoire PRiSM, Université de Versailles Saint-Quentin-en-Yvelines, 45 avenue des États-Unis, 78035 Versailles Cedex. Email: herve.fournier@prism.uvsq.fr

†INRA, UR341 Mathématiques et Informatique Appliquées, Domaine de Vilvert, F-78352 Jouy-en-Josas cedex. Email: antoine.vigneron@jouy.inra.fr.

We prove this result by applying Ben-Or's technique [1, 3, 16] to a suitable family of polytopes. Our lower bound implies that the algorithm by Ramos [17] is optimal for computing the diameter of a 3-polytope.

Similar problems of closing the gap between an $\Omega(n)$ lower bound and an $O(n \log n)$ upper bound have been studied recently. Chazelle et al. [6] mention that it is possible to compute the convex hull of two 3-polytopes in linear time, and it is not known whether the convex hull of a subset of the n vertices of a convex polytope can be computed in $O(n)$ time. On the other hand, given the Delaunay triangulation of a set P of n points in $\mathbb{R}^2$ (which is a special case of 3-dimensional convex hulls [8]), it is possible to compute the Delaunay triangulation of any subset of P in $O(n)$ time.

Hopcroft posed the following well known problem [10]. Given n lines and n points in $\mathbb{R}^2$, decide whether there is a point contained in a line. Matoušek [14] gave an $O\left(n^{4/3} 2^{O(\log^* n)}\right)$ time algorithm for this problem, but no $O(n^{4/3})$ time algorithm has been found so far. The only lower bound known for an algebraic computation tree is $\Omega(n \log n)$, and Erickson gave an $\Omega(n^{4/3})$ lower bound in a weaker model of computation [10]. Thus, finding a reduction from Hopcroft's problem to any other problem suggests that this problem is difficult to solve in $o(n^{4/3})$ time. Erickson gave several such reductions to various geometric problems [9], for instance he showed that ray shooting in polyhedral terrains and halfspace emptiness checking in $\mathbb{R}^5$ are harder than Hopcroft's problem. In this paper, we show that the same is true for the diameter problem in $\mathbb{R}^7$. More precisely, we show that there is a linear time reduction from Hopcroft's problem to the diameter problem in $\mathbb{R}^7$ using a real random access machine [16] (real-RAM). We give a similar reduction to the red-blue diameter problem in $\mathbb{R}^6$. Our approach is based on a linearization argument. Using the lifting transformation and advanced data structures for ray shooting [15], the diameter of a set of n points in $\mathbb{R}^d$ can be computed in $O(n^{2-2/(\lceil d/2 \rceil+1)} \log^{O(1)} n)$ time, which is $O(n^{1.6} \log^{O(1)} n))$ for $d = 7$.

2 Notation and preliminaries

We work in a fixed dimension d, so d is an integer such that $d = O(1)$. When $d = 3$, we use an orthonormal coordinate frame $Oxyz$ of $\mathbb{R}^3$. For all $a, b \in \mathbb{R}^d$, we denote by $\mathrm{d}(a, b)$ the Euclidean distance between a and b. For any set P of n points in $\mathbb{R}^d$, the diameter of P, that we denote by $\mathrm{diam}(P)$, is given by

$$\mathrm{diam}(P) = \max_{a,b \in P} \mathrm{d}(a, b).$$

Given two finite point sets $A, B \subset \mathbb{R}^d$, where the points in A are called the *red points* and the points in B are called the *blue points*, the *red-blue* diameter of (A, B) is

$$\mathrm{diam}(A, B) = \max_{a \in A, b \in B} \mathrm{d}(a, b).$$

If $a \in \mathbb{R}^d$ and B is a non-empty subset of $\mathbb{R}^d$, we denote by $\mathrm{d}(a, B)$ the distance between a and B, that is

$$\mathrm{d}(a, B) = \inf_{b \in B} \mathrm{d}(a, b).$$

The convex hull of $A \subset \mathbb{R}^d$ is denoted by $\mathrm{CH}(A)$. For all $a \in \mathbb{R}^d$ and $r > 0$, we denote by $\mathrm{B}(a, r)$ the open Euclidean ball with center a and radius r. We denote by $\mathrm{m}(a, b)$ the midpoint of the line segment ab. We use the notation $\|\cdot\|$ for the L_2 norm. In other words, for all $a, b \in \mathbb{R}^d$, we have $\|a-b\| = \mathrm{d}(a,b)$. We denote by $\langle a, b\rangle$ the inner product of a and b. We use the notation $\bar{u} = (u_1, u_2, \dots, u_m)$ to denote a sequence, and the concatenation of two sequences is written with a coma: $((1,2),(3,4)) = (1,2,3,4)$.

A *3-polytope* is a 3-dimensional convex polytope. The *combinatorial structure* of a 3-polytope P is the set of all inclusion relations between its vertices, edges and facets. In our lower bound arguments, we assume that the combinatorial structure of P is given together with the following information: the coordinates of the vertices of P and, for each facet f of P, the edges of f are given as a sequence ordered along the boundary of f.

2.1 Models of computation

The real-RAM model is the model of computation that is most commonly used to analyze geometric algorithms [16]. It is a random access machine that can store a real number or an integer in each memory cell. It can perform comparisons and arithmetic operations $(+,-,\times,/)$ between real numbers or between integers at unit costs. However, it is not allowed to convert between real variables and integer variables (for instance through a floor function). Integer values can be used as memory addresses to perform indirect addressing, but real numbers cannot serve this purpose. By default, a real-RAM can only use the real and integer constants 0 and 1, but we will also consider a more powerful model of real-RAM that can use arbitrary real constants as operands.

Before we prove our lower bound under the real-RAM model, we will first prove it under the *algebraic computation tree* [3] model. We will only use the algebraic computation tree model for decision problems, so following Ben-Or [1], we use the following definition where leaves are labeled by YES or NO. We denote by $x = (x_1, x_2, \dots, x_n) \in \mathbb{R}^n$ the input to our problem. An algebraic computation tree T is a binary tree where each node is either a computation node (a degree one node, with one son), a branching node (a degree two node, with two sons), or a leaf. A computation node u is associated with an arithmetic operation taken in $\{+,-,\times,/,\sqrt{\cdot}\}$. Each operand is either a real constant, an input number x_i, or a value obtained at a computation node that is an ancestor of u. At each branching node v, we compare with 0 the value obtained at a computation node that is an ancestor of v; each comparison can be taken in $\{>, \geqslant, =\}$. According to the result of this comparison, the program branches to one son of v or the other. So, according to the value of the input point x, the program follows a path in T that leads to a leaf labeled YES or NO. We say that T *decides* the set $W \subset \mathbb{R}^n$ if we reach a leaf labeled YES for all $x \in W$, and we reach a leaf labeled NO for all $x \notin W$.

Ben-Or proved the following result:

Theorem 1 (Ben-Or [1]) *Any algebraic computation tree that decides $W \subset \mathbb{R}^n$ has depth $\Omega(\log(\#W) - n)$, where $\#W$ is the number of connected components of W.*

This lower bound can be extended to real-RAM's that only take real numbers as input. Indeed, suppose that a real-RAM with arbitrary real constants decides a set $W \subset \mathbb{R}^n$. As the input consists of n real numbers, this real-RAM can only perform indirect addressing to memory cells at fixed (integer) memory addresses. The value stored in a fixed memory cell can be used directly by an algebraic computation tree. (Only indirect addressing to a variable address would be impossible to simulate with an algebraic computation tree.) Thus, all the possible branching and algebraic operations of this real-RAM can be unfolded into an algebraic computation tree. Therefore, a lower bound on the depth of all algebraic computation trees that decide W gives a lower bound on the worst case running time of any real-RAM with arbitrary real constants that decides W.

3 Diameter of a 3-polytope

In this section, we show that computing the diameter of a 3-polytope requires $\Omega(n \log n)$ time in the algebraic computation tree model. Our approach is the following. We first construct a family of 3-polytopes that have the same combinatorial structure, but do not all have the same diameter. (In particular, the set of polytopes with diameter smaller than 1 has a large number of connected components.) Then we apply Ben-Or's technique [1, 3, 16].

We will use the inequalities

$$\forall \theta \in \left[-\frac{\pi}{2}, \frac{\pi}{2}\right], \ \frac{\theta^2}{4} \leqslant 1 - \cos\theta \leqslant \frac{\theta^2}{2}, \tag{1}$$

with strict inequalities if $\theta \neq 0$.

Let $n > 0$ be an integer. Let α and φ denote two real numbers such that $0 < \alpha \leqslant \frac{1}{4}$ and $0 < \varphi \leqslant \frac{1}{4}$. Both are to be thought of as small enough, to be chosen later. Then we define $\psi = \frac{\varphi}{n}$, $\gamma = \frac{\alpha}{n}$, $t = \left(1 - \cos\left(\frac{1}{2}\psi\right)\right) / \left(1 + \cos\left(\frac{1}{2}\psi\right)\right)$ and $r = 1 - t$. The length r has the following property (see Figure 1): if e, f, g and h are four points such that $|ef| = |eg| = r$, $\angle feg = \angle feh = \frac{1}{2}\psi$ and $\angle efh = \frac{\pi}{2}$, then the midpoint $\mathrm{m}(g, h)$ is at distance 1 from e.

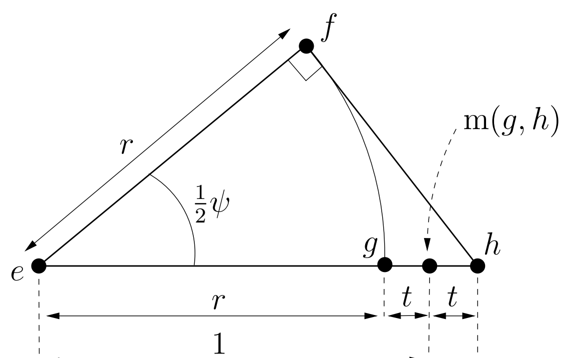

Figure 1: Geometric interpretation of r and t.

Now we define three sets of points in $\mathbb{R}^3$. (See Figure 2.) For all $i \in \{-n, -n+1, \ldots, n\}$, we define

$$a_i = \begin{pmatrix} \frac{1}{2}(1 - \cos(i\gamma)) \\ 0 \\ \frac{1}{2}\sin(i\gamma) \end{pmatrix}$$

and we denote $A = \{a_i \mid -n \leqslant i \leqslant n\}$. For all $i \in \{-n, -n+1, \ldots, n-1\}$ and $s \in \{-1, 1\}$, let

$$c_i^s = \begin{pmatrix} r\cos\left(\left(i + \frac{1}{2}\right)\psi\right) \\ r\sin\left(\left(i + \frac{1}{2}\right)\psi\right) \\ \frac{1}{2}s\alpha \end{pmatrix}$$

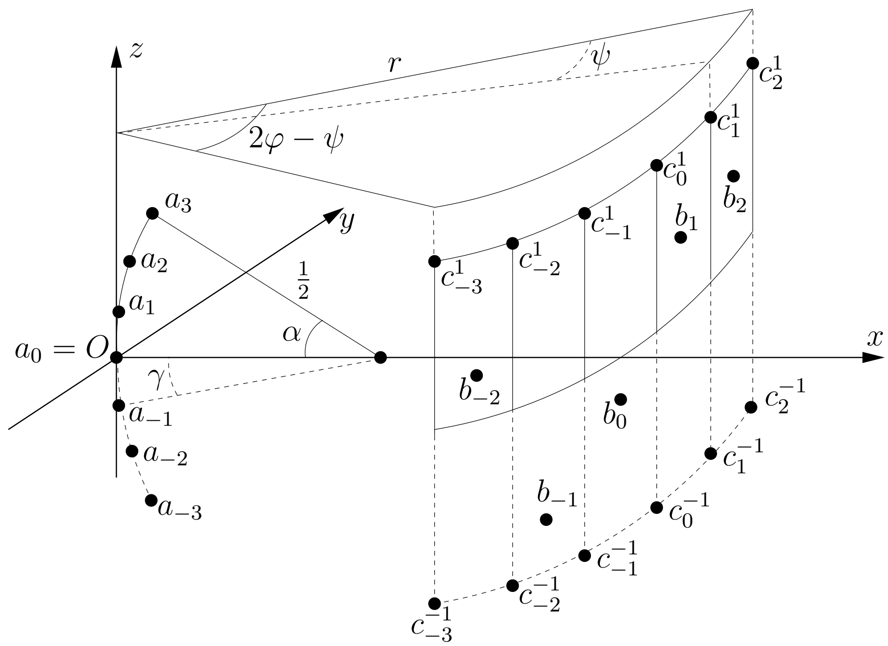

Figure 2: The sets $A, B(\bar{\beta})$ and C when $n = 3$ and $\bar{\beta} \in [-\alpha, \alpha]^{2n-1}$.

and $C = \{c_i^s \mid -n \leqslant i < n,\ s \in \{-1, 1\}\}$. Now for a parameter $\beta \in \mathbb{R}$ and for all $j \in \{-n+1, -n+2, \ldots, n-1\}$, we define

$$b_j(\beta) = \begin{pmatrix} \cos(j\psi) - \frac{1}{2}(1 - \cos\beta) \\ \sin(j\psi) \\ \frac{1}{2}\sin(\beta) \end{pmatrix}$$

For all $\bar{\beta} = (\beta_{-n+1}, \beta_{-n+2}, \ldots, \beta_{n-1}) \in \mathbb{R}^{2n-1}$, we define $B(\bar{\beta}) = \{b_j(\beta_j) \mid -n+1 \leqslant j \leqslant n-1\}$.

The following lemma shows that, for α small enough and $\bar{\beta} \in [-\alpha, \alpha]^{2n-1}$, the graph of $\mathrm{CH}(A \cup B(\bar{\beta}) \cup C)$ does not depend on the angle sequence $\bar{\beta}$.

Lemma 2 *Assume that $\alpha < 2t\cos(\frac{1}{2}\psi)$ and $\bar{\beta} \in [-\alpha, \alpha]^{2n-1}$. Then the graph of $\mathrm{CH}(A \cup B(\bar{\beta}) \cup C)$ is the the union of the graph of $\mathrm{CH}(A \cup C)$ and the set of the edges connecting each $b_j(\beta_j)$ to the points c_{j-1}^1, c_{j-1}^{-1}, c_j^1 and c_j^{-1}. (See Figure 3.)*

Proof: Let H_j be the vertical plane containing $\{c_j^1, c_j^{-1}\}$ and orthogonal to $(O, m(c_j^1, c_j^{-1}))$. (See Figure 4.) Let H'_j be the vertical plane containing the points $c_{j-1}^1, c_{j-1}^{-1}, c_j^1$ and c_j^{-1}. Let H^+ (resp. H^-) be the horizontal plane with equation $z = \frac{1}{2}\alpha$ (resp. $z = -\frac{1}{2}\alpha$). Let Δ_j be the interior of the polytope defined by the planes H_{j-1}, H_j, H'_j, H^+ and H^-.

By elementary trigonometry (see also Figure 1), we can show that

$$\begin{cases} b_j(0) \in \Delta_j, \\ d(b_j(0), H^+) = d(b_j(0), H^-) = \frac{1}{2}\alpha, \\ d(b_j(0), H'_j) > t, \text{ and} \\ d(b_j(0), H_{j-1}) = d(b_j(0), H_j) = t\cos(\frac{1}{2}\psi). \end{cases}$$

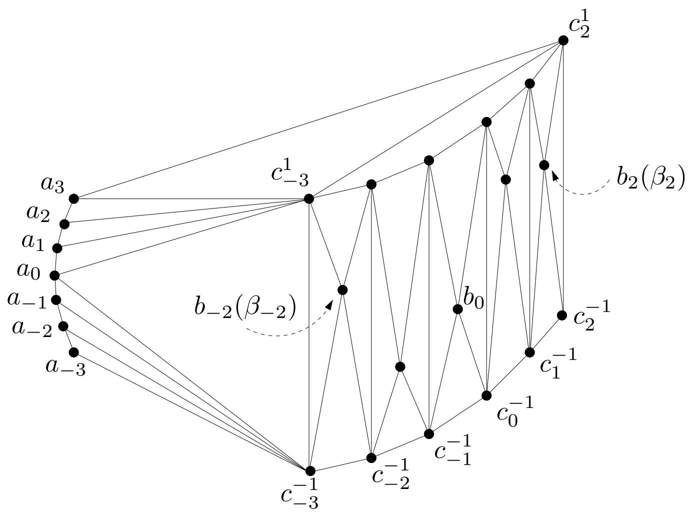

Figure 3: This figure shows $\mathrm{CH}(A \cup B(\bar{\beta}) \cup C)$. The graph of $\mathrm{CH}\,(A \cup C)$ is obtained by removing all the vertices in $B(\bar{\beta})$ and the adjacent edges.

Our assumption that $\frac{1}{2}\alpha < t\cos(\frac{1}{2}\psi)$ implies that $\mathrm{B}(b_j(0), \frac{1}{2}\alpha) \subset \Delta_j$. For all j, we have $\beta_j \in [-\alpha, \alpha]$, so $b_j(\beta_j) \in \mathrm{B}(b_j(0), \frac{1}{2}\alpha)$ and thus $b_j(\beta_j) \in \Delta_j$. Therefore, the only facet of $\mathrm{CH}(A \cup C)$ that is visible from $b_j(\beta_j)$ is the facet $c_{j-1}^{-1} c_{j-1}^{1} c_j^{1} c_j^{-1}$, and no point in $B(\bar{\beta}) \setminus \{b_j(\beta_j)\}$ is visible from $b_j(\beta_j)$; the result follows. $\square$

In order to apply Ben-Or's bound to our problem, we need to show that the special case we consider has a large number of connected components. Intuitively, the following lemma shows that for a given index i, the point $b_i(\beta)$ produces $2n$ connected components when $\beta \in [-\alpha, \alpha]$. Figure 5 shows the simple case where $i = 0$, and thus A and $b_i(\beta) = b_0(\beta)$ are cocircular. Lemma 3 gives the generalization to all values of i.

Lemma 3 *Assume that* $\varphi \leqslant \frac{1}{4n}$ *and* $j \in \{-n+1, -n+2, \ldots, n-1\}$. *Then the set* $\{b_j(\beta) \mid \beta \in [-\alpha, \alpha]$ *and* $\mathrm{diam}(A, \{b_j(\beta)\}) < 1\}$ *has at least* $2n$ *connected components.*

Proof: Let us first compute $\mathrm{d}^2(a_i, b_j(\beta))$. By developing the sum of squares and factoring, we obtain

$$\begin{aligned}
\mathrm{d}^2(a_i, b_j(\beta)) &= \frac{1}{4}\left(2 - (\cos(i\gamma) + \cos\beta) - 2\cos(j\psi)\right)^2 + \sin^2(j\psi) \\
&\quad + \frac{1}{4}(\sin(i\gamma) - \sin\beta)^2 \\
&= 1 + \frac{1}{4}(\cos(i\gamma) + \cos\beta)^2 + \cos^2(j\psi) - (\cos(i\gamma) + \cos\beta) \\
&\quad - 2\cos(j\psi) + (\cos(i\gamma) + \cos\beta)\cos(j\psi) + \sin^2(j\psi) \\
&\quad + \frac{1}{4}(\sin(i\gamma) - \sin\beta)^2 \\
&= \frac{5}{2} + \frac{1}{2}(\cos(i\gamma)\cos\beta - \sin(i\gamma)\sin\beta) - (\cos(i\gamma) + \cos\beta) \\
&\quad - 2\cos(j\psi) + (\cos(i\gamma) + \cos\beta)\cos(j\psi),
\end{aligned}$$

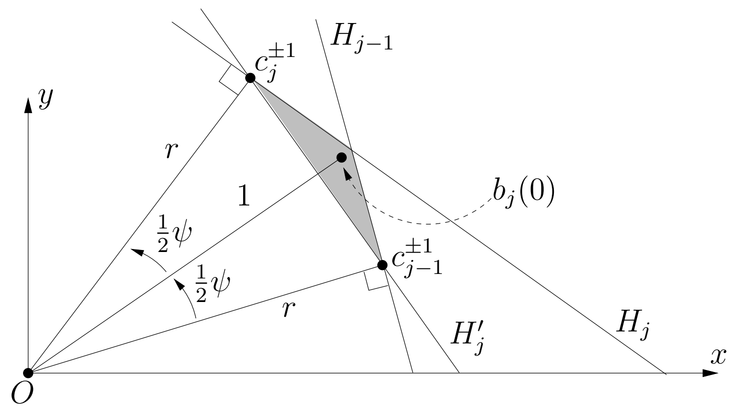

Figure 4: The shaded area is Δ_j, seen from above.

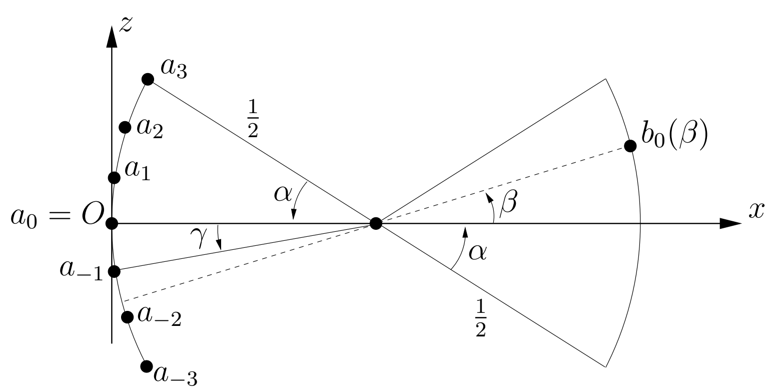

Figure 5: Lemma 3 when $n = 3$ and $i = 0$. The diameter is equal to 1 when β is a multiple of γ, and otherwise it is smaller than 1.

and thus

$$\mathrm{d}^2(a_i, b_j(\beta)) = 1 - \frac{1}{2}(1 - \cos(i\gamma + \beta)) + (1 - \cos(j\psi))(2 - \cos(i\gamma) - \cos\beta). \tag{2}$$

The result follows directly from the following two claims:

Claim 1. Let $i \in \{-n, -n+1, \ldots, n\}$ and $\beta = -i\gamma$. Then $\mathrm{d}(a_i, b_j(\beta)) \geqslant 1$.
This is obvious from Equation (2) since the second term evaluates to 0.

Claim 2. Let $k \in \{-n+1, -n+2, \ldots, n\}$ and $\beta = (k - \frac{1}{2})\gamma$. Then $\mathrm{diam}(A, \{b_j(\beta)\}) < 1$.
Let $i \in \{-n, -n+1, \ldots, n\}$. Let $\nu = \beta + i\gamma$. From Equation (2) we get

$$\mathrm{d}^2(a_i, b_j(\beta)) = 1 - \frac{1}{2}(1 - \cos\nu) + (1 - \cos(j\psi))(2 - \cos(i\gamma) - \cos\beta).$$

Note that $|\nu| \leqslant 2\alpha \leqslant \frac{1}{2} < \frac{\pi}{2}$. Moreover, by the choice of β, we have $|\nu| \geqslant \frac{1}{2}\gamma$. Thus Equation (1) yields $1 - \cos\nu > \frac{1}{4}\nu^2 \geqslant \frac{1}{16}\gamma^2$. Besides we have $1 - \cos(j\psi) \leqslant 1 - \cos\varphi < \frac{1}{2}\varphi^2$

and $2 - \cos(i\gamma) - \cos(\beta) < \frac{1}{2}(i\gamma)^2 + \frac{1}{2}\beta^2 \leqslant \alpha^2$. These inequalities imply that $\mathrm{d}^2(a_i, b_j(\beta)) < 1 - \frac{1}{32}\gamma^2 + \frac{1}{2}\varphi^2\alpha^2$. Remember that $\alpha = n\gamma$, so we obtain $\mathrm{d}^2(a_i, b_j(\beta)) < 1 + \frac{1}{2}\alpha^2\left(\varphi^2 - \frac{1}{16n^2}\right)$. As $\varphi \leqslant \frac{1}{4n}$, we conclude that $\mathrm{d}(a_i, b_j(\beta)) < 1$. $\square$

Lemma 3 only involves pairs of points in $A \times B(\beta)$. The following lemma shows that $A \times B(\bar{\beta})$ contains a diametral pair of $A \cup B(\bar{\beta}) \cup C$. It will allow us to apply Lemma 3 to the problem of finding the diameter of $A \cup B(\bar{\beta}) \cup C$.

Lemma 4 *Assume that $\alpha \leqslant \frac{1}{2}t$. Then for any $\bar{\beta} \in [-\alpha, \alpha]^{2n-1}$, we have*

$$\mathrm{diam}(A \cup B(\bar{\beta}) \cup C) = \mathrm{diam}(A, B(\bar{\beta})).$$

Proof: Clearly we have $\mathrm{d}(a_i, a_j) \leqslant \alpha \leqslant \frac{1}{4}$ and $\mathrm{d}(c_i^s, c_{i'}^{s'}) \leqslant 2r\varphi + \alpha \leqslant \frac{3}{4}$. In the same way, $\mathrm{d}(b_j(\beta), b_{j'}(\beta')) \leqslant 2\varphi + \alpha \leqslant \frac{3}{4}$ and $\mathrm{d}(a_i, c_j^s) \leqslant r + \alpha$. Similarly we have

$$\mathrm{d}(b_j(\beta), c_i^s) \leqslant \mathrm{d}(b_j(0), b_j(\beta)) + \mathrm{d}(b_j(0), c_i^s) \leqslant \frac{\alpha}{2} + 2r\varphi + \frac{\alpha}{2} + t \leqslant 2r\varphi + \frac{3}{2}t.$$

By our assumption that $\varphi \leqslant \frac{1}{4}$ and by Equation (1), we have $t \leqslant \frac{1}{128}$ so $\mathrm{d}(b_j(\beta), c_i^s) \leqslant \frac{3}{4}$. On the other hand, $\mathrm{d}(a_i, b_j(\beta)) \geqslant \mathrm{d}(a_0, b_j(0)) - \mathrm{d}(a_0, a_i) - \mathrm{d}(b_j(0), b_j(\beta)) > 1 - \frac{1}{2}\alpha - \frac{1}{2}\alpha = 1 - \alpha$. The result follows from the facts that $1 - \alpha \geqslant \frac{3}{4}$ and $1 - \alpha \geqslant r + \alpha$. $\square$

In order to be able to apply lemmas 2, 3 and 4, we need to find values of $\alpha \in (0, \frac{1}{4}]$ and $\varphi \in (0, \frac{1}{4}]$ such that the following three conditions hold simultaneously: $\alpha < 2t\cos(\frac{1}{2}\psi)$, $\varphi \leqslant \frac{1}{4n}$ and $\alpha \leqslant \frac{1}{2}t$. We choose $\varphi = \varphi_n = \frac{1}{4n}$ and small enough $\alpha = 2^{-10}n^{-4}$ satisfying these conditions.

From now on, we assume that $\alpha = \alpha_n$ and $\varphi = \varphi_n$ have been chosen as above. We define the sequences $\bar{a} = (a_{-n}, a_{-n+1}, \ldots, a_n)$ and $\bar{c} = (c_{-n}^{-1}, \ldots, c_{n-1}^{-1}, c_{-n}^{1}, \ldots, c_{n-1}^{1})$. For any $\bar{\beta} \in \mathbb{R}^{2n-1}$ we also define the sequence $\bar{b}(\bar{\beta}) = (b_{-n+1}(\beta_{-n+1}), b_{-n+2}(\beta_{-n+2}), \ldots, b_{n-1}(\beta_{n-1}))$. We define the set of sequences

$$\mathcal{S}_n = \{(\bar{a}, \bar{b}(\bar{\beta}), \bar{c}) \ \mid \ \bar{\beta} \in [-\alpha, \alpha]^{2n-1}\} \subset \mathbb{R}^{24n}.$$

Thus, each element of $\mathcal{S}_n$ is a sequence of $8n$ points in $\mathbb{R}^3$.

Lemma 5 *The set $\mathcal{S}_n$ can be decided by an algebraic computation tree with depth $O(n)$.*

Proof: Given three sequences $\bar{u}$, $\bar{v}$ and $\bar{w}$ of respectively $2n+1$, $2n-1$ and $4n$ points in $\mathbb{R}^3$, we want to check in linear time if there exists $\bar{\beta} \in [-\alpha, \alpha]^{2n-1}$ such that $(\bar{u}, \bar{v}, \bar{w}) = (\bar{a}, \bar{b}(\bar{\beta}), \bar{c})$. For fixed n, the coordinates of $\bar{a}$ and $\bar{c}$ are real constants. As this computation tree is allowed to use real constants (even if their expression uses trigonometric functions cos and sin), it is trivial to check that $\bar{u} = \bar{a}$ and $\bar{w} = \bar{c}$. Now it remains to check that there exists $\bar{\beta} \in [-\alpha, \alpha]^{2n-1}$ such that $\bar{v} = \bar{b}(\bar{\beta})$. We denote $\bar{v} = (v_{-n+1}, v_{-n+2}, \ldots, v_{n-1})$. For each integer $j \in \{-n+1, -n+2, \ldots, n-1\}$, we only need to check that v_j belongs to:

- the sphere of center $(\cos(j\psi) - \frac{1}{2}, \sin(j\psi), 0)$ and radius $\frac{1}{2}$,
- the plane $y = \sin(j\psi)$,
- and the halfspace $x \geqslant \cos(j\psi) - \frac{1}{2} + \cos(\alpha)$.

This can obviously be decided by a computation tree of linear depth. $\square$

Now we consider the following subset of $\mathcal{S}_n$:

$$\mathcal{E}_n = \left\{ (\bar{a}, \bar{b}(\bar{\beta}), \bar{c}) \;\mid\; \bar{\beta} \in [-\alpha, \alpha]^{2n-1} \text{ and } \operatorname{diam}(A \cup B(\bar{\beta}) \cup C) < 1 \right\}.$$

Lemma 6 *An algebraic computation tree that, given a sequence $\bar{s}$ of $8n$ points in $\mathbb{R}^3$, decides whether $\bar{s} \in \mathcal{E}_n$, has depth $\Omega(n \log n)$.*

Proof: By Lemma 4 we have

$$\mathcal{E}_n = \left\{ (\bar{a}, \bar{b}(\bar{\beta}), \bar{c}) \;\mid\; \bar{\beta} \in [-\alpha, \alpha]^{2n-1} \text{ and } \operatorname{diam}(A, B(\bar{\beta})) < 1 \right\},$$

and thus

$$\mathcal{E}_n = \{\bar{a}\} \times \prod_{j=-n+1}^{n-1} \{b_j(\beta) \;\mid\; \beta \in [-\alpha, \alpha] \text{ and } \operatorname{diam}(A, \{b_j(\beta)\}) < 1\} \times \{\bar{c}\}.$$

By Lemma 3 we know that $\mathcal{E}_n$ has at least $(2n)^{2n-1}$ connected components. We conclude by applying Ben-Or's bound [1, 3, 16] to $\mathcal{E}_n$. $\square$

The graph of a 3-polytope is planar, so a 3-polytope with n vertices has $O(n)$ edges and facets. Therefore, we can encode the coordinates of its n vertices using $3n$ real numbers, and we can encode its combinatorial structure and the ordering of the edges of each facet around its boundary using $O(n)$ integers—for instance, using a doubly-connected edge-list [8].

In the theorem below, we assume that the input is given using this encoding.

Theorem 7 *Assume that an algebraic computation tree T_n decides whether the diameter of a 3-polytope with n vertices is smaller than 1. Then T_n has depth $\Omega(n \log n)$.*

Proof: We denote by $(\bar{s}, \bar{g})$ the input of the tree T_{8n}, where $\bar{s} = (s_1, s_2, \ldots, s_{8n})$ denotes a sequence of $8n$ points in $\mathbb{R}^3$, and $\bar{g}$ encodes the graph of the convex hull of $S = \{s_1, s_2, \ldots, s_{8n}\}$. By Lemma 5, there is an algebraic computation tree U_n with depth $O(n)$ that decides whether $\bar{s} \in \mathcal{S}_n$. By plugging U_n to each accepting leaf of T_{8n}, we obtain an algebraic computation tree T'_{8n} that accepts 3-polytopes $(\bar{s}, \bar{g})$ such that $\bar{s} \in \mathcal{S}_n$ and $\operatorname{diam}(S) < 1$. In other words, T'_{8n} accepts 3-polytopes $(\bar{s}, \bar{g})$ such that $\bar{s} \in \mathcal{E}_n$. By Lemma 2, all 3-polytopes $(\bar{s}, \bar{g})$ accepted by T'_{8n} have the same graph $\bar{g} = \bar{g}_0$. Therefore, substituting the input part $\bar{g}$ with $\bar{g}_0$ in this tree gives an algebraic computation tree T''_{8n} that decides whether $\bar{s} \in \mathcal{E}_n$. If we denote by d_n the depth of T_n, then the depth of T''_{8n} is $d_{8n} + O(n)$. On the other hand, Lemma 6 tells us that T''_{8n} has depth $\Omega(n \log n)$. It follows that $d_n = \Omega(n \log n)$. $\square$

The lower bound we obtained on the depth of algebraic computation trees computing the diameter of a 3-polytope can be turned into a lower bound for real-RAM's using arbitrary constants. We achieve it by considering a non-uniform model of real-RAM, that is, a sequence of real-RAM's with arbitrary constants, the n-th machine solving the problem for inputs made of n points. This model is of course stronger than the uniform real-RAM model. If a non-uniform real-RAM computes the diameter of a 3-polytope in time $t(n)$, then it can be turned into a non-uniform real-RAM deciding $(\mathcal{E}_n)_{n \in \mathbb{N}}$ in time $t(n) + O(n)$, because the input part that encodes the combinatorial structure is fixed, and thus it can be seen as a set of constants of the real-RAM. As explained in the last paragraph of Section 2.1, since this problem has only real inputs, the lower bound obtained for algebraic computation trees through Lemma 6 holds in the real-RAM model as well.

Theorem 8 *Assume that a real-*RAM *with arbitrary real constants decides whether the diameter of a 3-polytope with* n *vertices is smaller than* 1*. Then it has worst case running time* $\Omega(n \log n)$*.*

Ramos gave an $O(n \log n)$ upper bound on the complexity of computing the diameter of a 3-polytope in the real-RAM model [17]. His algorithm can also be turned into an algebraic computation tree of depth $O(n \log n)$. So theorems 7 and 8 imply that the complexity of computing the diameter of a 3-polytope is $\Theta(n \log n)$, both in the algebraic computation tree model and in the real-RAM model.

4 Diameter is harder than Hopcroft's problem

Hopcroft posed the following problem: given a set L of lines and a set P of points in $\mathbb{R}^2$, decide whether there is a line $\ell \in L$ and a point $p \in P$ such that $p \in \ell$. We will show that the diameter problem for a point set in $\mathbb{R}^7$ is harder than Hopcroft's problem. We first show a reduction to the red-blue diameter problem. In the following two propositions, we deal with a real-RAM that can use the constant $\sqrt{2}$. We will explain at the end of this section how we can avoid using this constant.

Proposition 9 *There is a linear-time reduction from Hopcroft's problem to the red-blue diameter problem in* $\mathbb{R}^6$ *using a real-*RAM *that uses the constant* $\sqrt{2}$*.*

Proof: Let $(a_1, \dots, a_n, b_1, \dots, b_p)$ be an instance of Hopcroft's problem. For all i, the point $a_i = (u_i, v_i, w_i)$ corresponds to the line with equation $u_i x + v_i y + w_i = 0$. Each point b_i is given by its coordinates $(x_i, y_i) \in \mathbb{R}^2$. We denote $c_i = (x_i, y_i, z_i = 1)$. So our instance of Hopcroft's problem has a positive answer if and only if $\langle a_i, c_j \rangle = 0$ for some i and j.

We denote $a'_i = a_i / \|a_i\|$ and $c'_i = c_i / \|c_i\|$. We define the function $\theta : \mathbb{R}^3 \to \mathbb{R}^6$ by

$$\theta(x, y, z) = \frac{1}{x^2 + y^2 + z^2}(x^2, y^2, z^2, \sqrt{2}xy, \sqrt{2}xz, \sqrt{2}yz).$$

Now let the points given by $f_i = \theta(a_i)$ be the red points, and let the points $g_i = \theta(c_i)$ be the blue points. Notice that $\|f_i\|^2 = \|a_i\|^4 / \|a_i\|^4 = 1$, and $\|g_i\|^2 = 1$. It implies that $\|f_i - g_j\|^2 = \|f_i\|^2 + \|g_j\|^2 - 2\langle f_i, g_j \rangle = 2 - 2\|a_i\|^{-2}\|c_j\|^{-2}(u_i x_j + v_i y_j + w_i z_j)^2 = 2 - 2\langle a'_i, c'_j \rangle^2$. Thus, the red-blue diameter of the 6 dimensional point sets $\{f_i \mid 1 \leqslant i \leqslant n\}$ and $\{g_i \mid 1 \leqslant i \leqslant p\}$ is 2 if and only if our instance of Hopcroft's problem is positive. $\square$

A simple modification of the proof of Proposition 9 gives a reduction to the diameter problem in $\mathbb{R}^7$.

Proposition 10 *There is a linear-time reduction from Hopcroft's problem to the diameter problem in* $\mathbb{R}^7$ *using a real-*RAM *that uses the constant* $\sqrt{2}$*.*

Proof: With the notations from the previous proposition, we define $\hat{f}_i = (f_i, 1) \in \mathbb{R}^7$ and $\hat{g}_j = (g_j, -1) \in \mathbb{R}^7$. One have $\|\hat{f}_i - \hat{f}_j\|^2 = \|f_i - f_j\|^2 \leqslant (\|f_i\| + \|f_j\|)^2 \leqslant 4$, and $\|\hat{g}_i - \hat{g}_j\|^2 \leqslant 4$ in the same way. But $\|\hat{f}_i - \hat{g}_j\|^2 = \|f_i - g_j\|^2 + 4 \geqslant 4$. Thus, the diameter of $\{\hat{f}_1, \dots, \hat{f}_n, \hat{g}_1, \dots, \hat{g}_p\}$ is realized by a couple of points of the form $(\hat{f}_i, \hat{g}_j)$. $\square$

In propositions 9 and 10, we allowed the use of the constant $\sqrt{2}$ by the real-RAM machine. It can be avoided at the expense of increasing the dimension if we replace our function $\theta : \mathbb{R}^3 \to \mathbb{R}^6$ in the proof of Proposition 9 by the function $\theta' : \mathbb{R}^3 \to \mathbb{R}^9$ defined as follows:

$$\theta'(x, y, z) = \frac{1}{x^2 + y^2 + z^2}(x^2, y^2, z^2, xy, xy, xz, xz, yz, yz).$$

Thus we obtain the following result:

Proposition 11 *Using a real-*RAM*, there are linear-time reductions from Hopcroft's problem to the red-blue diameter problem in $\mathbb{R}^9$ and to the diameter problem in $\mathbb{R}^{10}$.*

5 Concluding remarks

It would be interesting to extend our lower bound for 3-polytopes to the following randomized setting. A *randomized computation tree* (RCT) is a collection of algebraic computation trees $\{T_i\}_{i \in I}$ together a probability vector $\{p_i\}_{i \in I}$ such that $p_i \geqslant 0$ and $\sum_i p_i = 1$. The depth of this RCT is the maximum depth of $\{T_i\}_{i \in I}$. The RCT model requires that there exists a constant $\varepsilon < 1/2$ such that, for each input, the error probability is bounded above by ε. Given n points in $\mathbb{R}^2$, a RCT of depth $\Omega(n \log n)$ is needed to decide if the diameter of this point set is smaller than 1: it follows from the standard reduction from diameter to set disjointness [16] and the $\Omega(n \log n)$ lower bound on set disjointness for RCTs proved by Grigoriev [11]. Of course this lower bound holds in higher dimension, and is optimal in $\mathbb{R}^2$ and $\mathbb{R}^3$. Is it possible to obtain the same randomized lower bound in $\mathbb{R}^3$ when the combinatorial structure of the convex hull is given? Our lower bound argument does not apply: in the RCT model, the logarithm of the number of connected components of a set does not necessarily provide a lower bound on the depth of a RCT that decides this set [4]. Even a lower bound on the less powerful d-RDT model (randomized algebraic decision trees with test nodes of maximum degree bounded by a constant d) would be interesting.

As we noted earlier, our lower bound for computing the diameter of a convex polytope leaves no room for improvement. Our results on the diameter for point sets in higher dimension, however, are not known to be optimal. First there is no lower bound other than $\Omega(n \log n)$ for Hopcroft's problem in the algebraic computation tree model. Second, even assuming that Hopcroft's problem cannot be solved in $o(n^{4/3})$ time, our result is not entirely satisfactory because the best known algorithm for the red-blue diameter problem [15] in $\mathbb{R}^6$ runs in $O(n^{1.5} \log^{O(1)} n)$ time. On the other hand, the red-blue diameter in $\mathbb{R}^4$ can be computed in $O(n^{4/3} \log^{O(1)} n)$ time, so it would be interesting to prove that this problem is harder than Hopcroft's problem. (Similarly, Erickson [9] asked whether the diameter in $\mathbb{R}^4$ is harder than halfspace emptiness checking in $\mathbb{R}^5$.)

Another intriguing question is the following. In propositions 9 and 10 we find reductions from Hopcroft's problem to diameter problems using a real-RAM that can use the constant $\sqrt{2}$. In proposition 11, we use a real-RAM without constant, and we obtain reductions to diameter problems in 3 dimensions higher. Is it possible to find such a reduction without increasing the dimension?

Acknowledgment

We would like thank Otfried Cheong and the anonymous referees for their helpful comments.

References

[1] M. Ben-Or. Lower bounds for algebraic computation trees. In *Proceedings of the 15th Annual ACM Symposium on Theory of Computing*, pages 80–86, 1983.

[2] S. Bespamyatnikh. An efficient algorithm for the three-dimensional diameter problem. *Discrete and Computational Geometry*, 25(2):235–255, 2000.

[3] P. Bürgisser, M. Clausen, and M. Shokrollahi. *Algebraic Complexity Theory*. Springer, 1997.

[4] P. Bürgisser, M. Karpinski, and T. Lickteig. On randomized semi-algebraic test complexity. *Journal of Complexity*, 9(2):231–251, 1993.

[5] T. Chan. Approximating the diameter, width, smallest enclosing cylinder, and minimum-width annulus. *International Journal of Computational Geometry and Applications*, 12(1–2):67–85, 2002.

[6] B. Chazelle, O. Devillers, F. Hurtado, M. Mora, V. Sacristán, and M. Teillaud. Splitting a delaunay triangulation in linear time. *Algorithmica*, 34(1):39–46, 2002.

[7] K. Clarkson and P. Shor. Applications of random sampling in computational geometry, II. *Discrete and Computational Geometry*, 4:387–421, 1989.

[8] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, Berlin, Germany, 2nd edition, 2000.

[9] J. Erickson. On the relative complexities of some geometric problems. In *Proceedings of the 7th Canadian Conference on Computational Geometry*, pages 85–90, 1995.

[10] J. Erickson. New lower bounds for Hopcroft's problem. *Discrete and Computational Geometry*, 16:389–418, 1996.

[11] D. Grigoriev. Randomized complexity lower bounds. In *Proceedings of the 30th ACM Symposium on Theory of Computing*, pages 219–223, 1998.

[12] S. Har-Peled. A practical approach for computing the diameter of a point set. In *Proceedings of the Seventeenth Annual Symposium on Computational Geometry*, pages 177–186, New York, June 3–5 2001. ACM Press.

[13] G. Malandain and J.-D. Boissonnat. Computing the diameter of a point set. *International Journal of Computational Geometry and Applications*, 12(6):489–510, 2002.

[14] J. Matoušek. Range searching with efficient hierarchical cuttings. *Discrete and Computational Geometry*, 10(2):157–182, 1993.

[15] J. Matoušek and O. Schwarzkopf. On ray shooting in convex polytopes. *Discrete and Computational Geometry*, 10(2):215–232, 1993.

[16] F. Preparata and I. Shamos. *Computational geometry: An introduction*. Texts and Monographs in Computer Science. Springer-Verlag, New York, 2nd edition, 1985.

[17] E. Ramos. An optimal deterministic algorithm for computing the diameter of a three-dimensional point set. *Discrete and Computational Geometry*, 26:233–244, 2001.