

Congruence Testing of Point Sets in 4 Dimensions

Heuna Kim and Günter Rote
Institut für Informatik, Freie Universität Berlin

March 23, 2016

Abstract

Congruence between two sets of n points in 4 dimension can be checked in $O(n \log n)$ time.

Contents

| | | |
|-----------|---|-----------|
| 1 | Problem Statement | 2 |
| 2 | The Computational Model | 2 |
| 3 | Overview | 3 |
| 4 | Previous Algorithms | 3 |
| 4.1 | Congruence Testing in the Plane | 4 |
| 4.2 | Congruence Testing In Three-Dimensional Space | 4 |
| 4.3 | Congruence Testing in Four and Higher Dimensions | 5 |
| 5 | Four-Dimensional Rotations | 6 |
| 6 | Angles between two Planes in 4-Space | 7 |
| 7 | Plücker Space and Plücker Distance | 10 |
| 7.1 | Other Distances | 12 |
| 8 | The Construction of Hopf Fibrations | 12 |
| 8.1 | Construction by Isoclinic Rotations | 13 |
| 8.2 | Equivalence of an Invariant Family and a Hopf Bundle. | 13 |
| 8.3 | Properties of the Hopf Fibration | 14 |
| 9 | Closest-Pair Graphs | 16 |
| 10 | The Coxeter Classification of Four-Dimensional Point Groups | 17 |
| 10.1 | The Radius of an Inscribed Sphere of a Fundamental Region | 17 |
| 11 | Packing Arguments and Definitions of Constants | 20 |
| 11.1 | Kissing Numbers | 20 |
| 11.2 | The Closeness Threshold | 20 |
| 11.3 | The Icosahedral Threshold | 21 |
| 12 | Canonical Axes | 21 |
| 13 | Congruence Types of Vertex and Edge Figures | 22 |

| | | |
|-----------|--|-----------|
| 14 | 1+3 Dimension Reduction | 23 |
| 15 | The Pruning and Condensing Principles | 23 |
| 16 | Overview of the New Algorithm | 24 |
| 17 | Iterative Condensing Based on the Closest-Pair Graph: Algorithm C | 26 |
| 18 | Generating Orbit-Cycles: Algorithm O | 28 |
| 19 | The Mirror Case: Algorithm R | 30 |
| 20 | Finding Representative Points from a 2-Sphere: Algorithm K | 31 |
| 21 | Marking and Condensing of Great Circles: Algorithm M | 32 |
| 22 | 2+2 Dimension Reduction: Algorithm T | 34 |
| | 22.1 A Canonical Set Procedure | 35 |
| 23 | Concluding Remarks | 38 |
| | 23.1 Practical Implementability | 38 |
| | 23.2 Regularity and Related Open Questions | 38 |
| | Bibliography | 39 |

1 Problem Statement

Problem Statement. We are given two n -point sets $A, B \subset \mathbb{R}^d$. Our aim is to test whether they are *congruent*, i. e., whether there exists an orthogonal matrix and a translation vector t such that $B = RA + t$.

As well as being intuitive, congruence is a fundamental notion in geometry. The well-known congruence criteria for triangles go back to 300 BC, written in Euclid’s Elements. Algorithms for testing congruence for point sets have been developed since the 1970’s [25, 4, 18, 30, 5, 2, 10, 1].

Problem Set-Up. We may eliminate the translation vector t by translating A and B so that their centroids are moved to the origin.

Also, we only consider orthogonal matrices R of determinant $+1$ (“direct congruences”, rotations). If we want to allow orthogonal matrices of determinant -1 as well (“mirror congruence”, rotations and a reflection), we can just repeat the algorithm with a reflected copy of B .

We solve this problem in $O(n \log n)$ time in $d = 4$ dimensions. This paper accompanies our conference article [22], and it contains all details that were omitted there for lack of space.

2 The Computational Model

Because the output of problem is sensitive to numerical errors, it is natural to consider an approximate version of the congruence testing problem. However, congruence testing with error tolerances is known to be NP-hard [14, 20]. We therefore restrict our concern to the exact case. (Error tolerances are discussed again in Chapter 23.)

We assume that we can perform exact arithmetic with real numbers. Thus, we use the *Real Random-Access Machine (Real-RAM) model* (see for example [28, Chapter 1.4]). This assumption is common in Computational Geometry. In this model, we can freely use square roots, sines and cosines, and basic operations from linear algebra such as eigenvalues of 2×2 -matrices (or matrices of any constant size), and we assume that we obtain exact results in

constant time. Sines and cosines and their inverses can be eliminated in favor of purely algebraic operations.

If we would restrict our model to rational inputs, this would severely restrict the problem. For example, in the plane, the only rotational symmetries that a point set with rational coordinates can have are multiples of 90° . A three-fold symmetry is possible in 3 dimensions, but a fivefold symmetry is impossible in any dimension. An input with limited symmetry, however, would not be interesting enough to consider. The previous algorithms for the exact congruence testing problem [25, 4, 18, 30, 5, 2] were developed around 70's–80's under the same assumption. The only papers that explicitly mentioned the assumption are [5, 2].

3 Overview

In Section 4, we survey the previous literature on congruence testing, since we build on ideas from these algorithms.

Our new algorithm uses tools that exploit the geometric structure of 4-space. The first part of the body of the paper (Sections 5–13) develops these tools. We elaborate the necessary background to understand how these tools work.

As a result of symmetries, we will also encounter beautiful mathematical structures. In particular, the structure of Hopf bundles organizes the special case of isoclinic planes in a pleasant way. We describe the interpretation of the Hopf fibration in a very geometric and elementary way by using four-dimensional rotations. This interpretation clarifies the necessary and sufficient condition for a collection of great circles to be in a Hopf bundle.

This part includes preliminaries about four-dimensional rotations, angles between a pair of planes, Plücker coordinates, four-dimensional Coxeter groups, kissing numbers and representing a neighborhood of a vertex or an edge in 4-space.

The second part (Sections 14–22) is devoted to the new algorithm. The techniques of closest pairs, pruning, and dimension reduction have been used for congruence testing before. We extend these ideas and apply them in a novel way. Our algorithm extracts helices around great circles on 3-sphere by taking advantage of the structure of the closest-pair graph.

On top of the previous dimension reduction principle, a new efficient dimension reduction technique, called 2+2 dimension reduction (Section 22) is developed for the case that a given point set has rotational symmetries that make two orthogonal planes invariant. This technique is simple to implement and resolves the tricky case in Brass and Knauer's algorithm [10] as well.

Figure 1 shows the relations between the sections in the two parts.

4 Previous Algorithms

It has been shown that the complexity of the problem is $\Omega(n \log n)$ independently by Atallah (1985) [4], Highnam (1986) [18], and Atkinson (1987) [5]. This bound holds even in one-dimensional space. It can be easily seen by reducing the problem to the set equality problem in the real line.

It has been believed that the congruence testing problem should be fixed-parameter tractable with the dimension parameter d , that is, solvable in time $O(f(d)n^C)$ for an arbitrary function $f(d)$ and a constant C that is independent of d . More precisely, it is believed that $O(n \log n)$ algorithms should exist for any fixed dimension.

The previous best deterministic algorithm by Brass and Knauer (2000) [10] for d -space for $d > 3$ took time $O(n^{\lceil d/3 \rceil} \log n)$. The previous best randomized algorithm by Akutsu (1998) [1] takes time $O(n^{\lfloor d/2 \rfloor / 2} \log n)$ for $d \geq 6$ and time $O(n^{3/2} \log n)$ for $d = 4, 5$. Therefore, the previous best result in 4-space was time $O(n^2 \log n)$ deterministically and time $O(n\sqrt{n} \log n)$ with randomization.



Figure 1: Dependence between the first part (Sections 5–13, in the outer ring) and the second part (Sections 14–22, in the inner disk).

We will survey some previous congruence testing algorithms, because we will use them and build on their ideas. A thorough survey of previous algorithms is given in Kim’s thesis [21].

4.1 Congruence Testing in the Plane

The first efficient algorithm for congruence testing is due to Manacher (1976) [25]. After translating point sets so that the centroid lies on the origin O , Manacher’s algorithm sorts all input points by spherical coordinates (r, θ) , first by angle θ , and secondly by distance r to the origin O in both increasing order. Let $\{p_i = (r_i, \theta_i) | i = 1, \dots, n\}$ be the given point set indexed by the sorted order. Finally, the algorithm generates a cyclic sequence in which r_i and $\angle p_i O p_{i+1} = \theta_{i+1} - \theta_i$ alternate for $1 \leq i \leq n$: $\langle r_1, \theta_2 - \theta_1, r_2, \dots, \theta_n - \theta_{n-1}, r_n \rangle$. Point sets are congruent if and only if the corresponding cyclic sequences are the same up to cyclic shifts. His paper [25] mainly discusses geometric applications of the algorithm by Knuth, Morris and Pratt [24] that determines whether a string y contains a string x as a substring in linear time $O(|x| + |y|)$. It can be determined if two sequences are the same up to cyclic shifts by duplicating one sequence and checking if the original sequence is a subsequence of the duplicated sequence by using the algorithm by Knuth, Morris and Pratt in linear time. Manacher’s algorithm takes time $O(n \log n)$ due to sorting.

4.2 Congruence Testing In Three-Dimensional Space

Two algorithms rely on graph isomorphism for planar graphs: an algorithm of Sugihara [30] for congruence testing between (not necessarily convex) polytopes, and an algorithm of Alt, Mehlhorn, Wagener, and Welzl [2] for point sets. They are based on the fact that the graph

of a convex polytope is planar. Thus, with some preprocessing, they reduce congruence testing to isomorphism testing for labeled planar graphs, which can be solved in time $O(n \log n)$ by applying the partitioning algorithm of Hopcroft and Tarjan [19], see also [32, Section 4.13]. The total time for these algorithms is $O(n \log n)$. We will not use this approach.

Atkinson (1987) [5]. Atkinson’s algorithm determines congruence for two point sets in 3-space. This algorithm begins by reducing input point sets as small as possible while preserving symmetries and then considers a bounded number of all the possible matches of the original point sets, obtained by the reduced sets. The principle of this reduction procedure is different from *pruning* mentioned before (Section 2 and Section 15) but is more related to a *canonical set procedure* (Sections 12 and 22.1). The first part of this algorithm can be considered a canonical set procedure of for a 2-sphere with rotational symmetries. This reduction procedure preserves symmetries, whereas pruning does not. For reduction, Atkinson’s algorithm constructs a closest-pair graph, i. e., a graph whose vertices are points and edges are pairs of vertices that achieve the minimum distance (see Section 9). Then, the algorithm prunes points by the congruence type of a neighborhood of a vertex v , or by a *vertex figure* (see Section 13). Then, either there is a component of the centroid different from the original centroid or the degrees of vertices are bounded since v has at most five closest vertices by the kissing number (see Section 11). The sets can be further reduced by traversing each component if the degree is two and by comparing faces if the degree is three or four or five. At the end, there are only two cases:

- (i) the reduced set is a singleton set or consists of two antipodal points.
- (ii) the reduced set has a bounded cardinality more than two.

For (ii), the algorithm finds all the rotations from a pair of non-antipodal points in one reduced set to any pair of non-antipodal points at the same angle in another reduced set. Then, it checks if at least one of such rotations transforms one input point set to another input point set. For (i), the rotational axis can be identified so Manacher’s algorithm [25] (see the first paragraph of Section 4.1) can be applied after each point is represented by a cylindrical coordinate regarding the common rotational axis as the z -axis.

The variation of the first part of Atkinson’s algorithm will be used in the new algorithm and discussed again in Lemma 20 in Section 20.

4.3 Congruence Testing in Four and Higher Dimensions

Alt, Mehlhorn, Wagener, and Welzl (1988) [2]. This algorithm reduces a d -dimensional congruence testing problem to n subproblems in $(d - 1)$ -space. As in the three-dimensional case of their algorithm, it projects all the points radially to the $(d - 1)$ -dimensional unit sphere S centered at the centroid and labels them with distances to the centroid. Fix one point a in the resulting set from A . The next step is another projection to a $(d - 2)$ -dimensional unit sphere S' which is the intersection of S and the hyperplane that orthogonally bisects the line segment from a to the centroid; project each point x in A except a onto S' along the arc from x to a on the surface of S . Let us denote the resulting set as A' . We can obtain sets B'_1, \dots, B'_n in the same manner by fixing all points of B . Two given point sets A and B are congruent if and only if A' is congruent to at least one of B'_i for $i = 1, \dots, n$ with taking labels into account. As a result, we obtained n subproblems to determine congruence between A' and B'_i for $i = 1, \dots, n$. Thus, the algorithm by Alt, Mehlhorn, Wagener, and Welzl takes time $O(n^{d-2} \log n)$ in d -space.

The following algorithms achieve the best known deterministic and randomized running times in high dimension:

Akutsu (1998) [1]: Akutsu developed a *Monte Carlo algorithm* for congruence testing problem in d -space. The algorithm is randomized and it is based on the birthday paradox. For

two congruent input sets, the algorithm is guaranteed to find the congruence only with high probability. Akutsu stated the running time as $O(n^{(d-1)/2} \log n)$, and he mentioned that it can be reduced to $O(n^{d/4+O(1)})$ by an unpublished idea of J. Matoušek. Matoušek's idea is to match closest *pairs*, i. e., pairs of points that attain the minimum distance (see Section 9) instead of input *points*. Since a closest pair together with the centroid spans a 2-plane, This idea allows to reduce the dimension by two steps at a time. and the orthogonal space of a 2-plane is $(d-2)$ -dimensional if the ambient space is d -dimensional. This idea can be used for the algorithm by Alt, Mehlhorn, Wagener, and Welzl as well, and it improves the time to $O(n^{\lfloor \frac{d}{2} \rfloor} \log n)$. Akutsu's algorithm takes time $O(n^{\lfloor d/2 \rfloor} \log n)$ for $d \geq 6$ and $O(n^{3/2} \log n)$ for $d = 4, 5$. (The accurate analysis was not provided in [1].)

Brass and Knauer (2003) [10]. Brass and Knauer's algorithm extends Matoušek's idea by fixing a *triple* of points that contains two closest pairs instead of fixing a closest pair. This allows to reduce the dimension in steps of 3 unless the triple and the centroid are coplanar. The algorithm achieves time $O(n^{\lfloor d/3 \rfloor} \log n)$. The difficulty of the algorithm is the occurrence of tricky cases, such as a union of two point sets in orthogonal planes. These cases require special handling.

5 Four-Dimensional Rotations

A 4-dimensional rotations can be encoded by a 4×4 rotation matrix R , that it, a 4×4 orthogonal matrix with determinant $+1$. It has four eigenvalues of absolute value 1, whose product is $+1$. They come in conjugate complex pairs, $e^{\pm i\varphi}$ and $e^{\pm i\psi}$. The special case of real eigenvalues 1 or -1 is included. Here, φ and ψ in the eigenvalues correspond to φ and ψ in the angular displacements.

Let us assume first that the four eigenvalues are distinct. This case is called a non-isoclinic rotation. The eigenvectors for each conjugate complex pair, which are conjugate complex, span a real 2-plane. We choose an orthonormal basis v_1, v_2 for the plane corresponding to $e^{\pm i\varphi}$, and an orthonormal basis v_3, v_4 for the plane corresponding to $e^{\pm i\psi}$. In the resulting orthonormal basis v_1, v_2, v_3, v_4 , the matrix R looks as follows:

$$R = R_{\varphi, \psi} = \begin{pmatrix} \cos \varphi & -\sin \varphi & 0 & 0 \\ \sin \varphi & \cos \varphi & 0 & 0 \\ 0 & 0 & \cos \psi & -\sin \psi \\ 0 & 0 & \sin \psi & \cos \psi \end{pmatrix} \quad (1)$$

The pair $\{P, Q\}$ of orthogonal planes $P = \langle v_1, v_2 \rangle$ and $Q = \langle v_3, v_4 \rangle$ is uniquely determined by R , but the basis v_1, v_2 and v_3, v_4 for each plane is not unique. Here, P and Q are the only pair of planes that are invariant under the rotation. We can require that the basis v_1, v_2, v_3, v_4 is positively oriented by flipping v_3 and v_4 and changing from $R_{\varphi, \psi}$ to $R_{\varphi, -\psi}$ if necessary (or by flipping v_1 and v_2 and changing to $R_{-\varphi, \psi}$, if we prefer). The angle pair $\{\varphi, \psi\}$ in the range $-\pi < \varphi, \psi < \pi$ is then unique up to a simultaneous negation to $\{-\varphi, -\psi\}$.

If the four eigenvalues are not distinct, they come in equal pairs $e^{i\varphi}, e^{i\varphi}, e^{-i\varphi}, e^{-i\varphi}$. Such rotations are called *isoclinic rotations*. We exclude the special cases $1, 1, 1, 1$ and $-1, -1, -1, -1$ (the identity and the inversion) and assume $0 < \varphi < \pi$. We can still find a basis v_1, v_2, v_3, v_4 for which the matrix has the form $R = R_{\varphi, \varphi}$, but the decomposition into two orthogonal planes is not unique. There are infinitely many pairs of orthogonal planes that are invariant under the rotation. If we insist that the basis is positively oriented, the rotation can either be written as $R_{\varphi, \varphi}$ or $R_{\varphi, -\varphi}$. A rotation of the first class is called a *right rotation*, and the second class a *left rotation*.




Figure 2: Angles between two planes P and Q in 4-space.

Axis Planes. For a rotation R that is not isoclinic, there is a pair of orthogonal planes $P = \langle v_1, v_2 \rangle, Q = \langle v_3, v_4 \rangle$ that is invariant under the rotation R . We call these two planes P, Q axis planes of the rotation R . Given a rotation matrix R , we can find axis planes P and Q by computing v_1, v_2, v_3, v_4 as follows. Let $u_\varphi, u_{-\varphi}, v_\psi, v_{-\psi}$ be eigenvectors of R such that u_θ corresponds to the eigenvalue $e^{i\theta}$ where $\theta \in \{\varphi, -\varphi, \psi, -\psi\}$. Then, $u_\varphi = v_1 + iv_2, u_{-\varphi} = v_1 - iv_2$ and $u_\psi = v_3 + iv_4, u_{-\psi} = v_3 - iv_4$ are conjugate pairs. Then,

$$v_1 = \frac{u_\varphi + u_{-\varphi}}{2}, \quad v_2 = \frac{u_\varphi - u_{-\varphi}}{2i},$$

$$v_3 = \frac{u_\psi + u_{-\psi}}{2}, \quad v_4 = \frac{u_\psi - u_{-\psi}}{2i}.$$

6 Angles between two Planes in 4-Space

In classical two- or three-dimensional geometry, an angle (between two lines, or between two planes in space, or between a line and a plane) fixes the relative position between the involved object up to congruence. In four (and higher) dimensions, we need *two* angles to fix the relative position between two planes. (More generally, k angles are defined between k -dimensional subspaces.)

Geometric Definition Using Half Rays. We first define an angle between two planes in 4-space geometrically and then show the computation method. We assume that the planes go through the origin, thus forming two-dimensional linear subspaces. The angle between two planes P, Q is defined as a pair α, β of two real numbers in the range $[0, \pi/2]$: We define α as the minimum angle between any two half lines Op in P and Oq in Q . Such an α exists [26]. We denote such half lines that make the minimum angle α as Op^* and Oq^* . Let U be the plane that contains Op^* and Oq^* . The plane U is perpendicular to the planes P and Q , since otherwise Op^* and the projection of Op^* to Q would make an angle smaller than α . Let U' be the orthogonal plane to U . Then U' is also perpendicular to P and Q . Let ℓ_p be the intersection of P and U' . Similarly, let ℓ_q be the intersection of Q and U' . We define the acute angle made by the lines ℓ_p and ℓ_q as β . See Figure 2.

For later reference we mention the following proposition:

Proposition 1. *There is an orthonormal basis u_1, u_2 of P and an orthonormal basis v_1, v_2 of Q such that*

1. *the angle between u_1, v_1 is α ,*
2. *the angle between u_2, v_2 is β ,*
3. *If $\alpha > 0$, the plane spanned by u_1, v_1 is perpendicular both to P and to Q ,*
4. *If $\beta > 0$, the plane spanned by u_2, v_2 is perpendicular both to P and to Q ,*

Isoclinic Planes. The pair of half lines Op^* and Oq^* that defines the angle α may not be unique but an infinite number of such pairs can exist. Accordingly, there are infinitely many perpendicular planes U to P and Q . In this case, an infinite number of planes U cut out P and Q by making all the equal angles [26]. Thus, $\alpha = \beta$. We say that a pair of planes of angle α , α where $0 \leq \alpha \leq \pi/2$ is *isoclinic*.

Computation. This can be formulated as follows, see [9].

$$\begin{aligned} \cos \alpha &= \max_{\substack{u_1 \in P \\ |u_1|=1}} \max_{\substack{v_1 \in Q \\ |v_1|=1}} u_1^T v_1 = (u_1^*)^T (v_1^*), \\ \cos \beta &= \max_{\substack{u_2 \in P \\ |u_2|=1 \\ u_2^T u_1^* = 0}} \max_{\substack{v_2 \in Q \\ |v_2|=1 \\ v_2^T v_1^* = 0}} u_2^T v_2. \end{aligned}$$

Let M_P and M_Q be the orthonormal basis of P and Q . The previous equation boils down to

$$\max_{\substack{u \in P \\ |u|=1}} \max_{\substack{v \in Q \\ |v|=1}} u^T v = \max_{\substack{y \in \mathbb{R}^2 \\ |y|=1}} \max_{\substack{z \in \mathbb{R}^2 \\ |z|=1}} y^T (M_P^T M_Q) z.$$

By the minimax characterization of the singular value decomposition (SVD), the angle can be computed as the SVD of the 2×2 matrix $M_P^T M_Q$:

$$Y^T (M_P^T M_Q) Z = \begin{pmatrix} \cos \alpha & 0 \\ 0 & \cos \beta \end{pmatrix},$$

where Y and Z are orthogonal 2×2 matrices.

Equivalent Definition by Orthogonal Projections. Let P, Q be a pair of planes and let C' be an orthogonal projection of a unit circle C in P to Q . The maximum and the minimum distances $\cos \alpha, \cos \beta$ from a point on C' to the origin define the same angle α and β . According to this definition, generally the orthogonal projection C' is an ellipse with the following reasoning. If P and Q are not isoclinic, there is a pair of planes $\{U, U'\}$ that are perpendicular to both P and Q and cut out P and Q at the angle α and β respectively. Then, the intersection of Q and U , (or the line containing the half line Oq^*) is the major axis (longer axis) of the ellipse because U defines the closest distance between any half lines in P and Q . Similarly, the intersection of Q and U' , (or ℓ_q) is the minor axis (shorter axis) of the ellipse. However, for a pair of isoclinic planes, the orthogonal projection C' becomes a circle. The argument is shown in Proposition 2.

Left and Right Pairs of Isoclinic Planes. Let P, Q be a pair of 2-planes in 4-space. Let v_1, v_2 be orthonormal vectors that span P . Let v'_1, v'_2 be the projections of v_1, v_2 to Q , and consider the projections v''_1 and v''_2 of v'_1 and v'_2 to the plane P^\perp orthogonal to P . If v''_1 and v''_2 are positively oriented together with v_1, v_2 , in other words, if the determinant of (v_1, v_2, v''_1, v''_2) is positive, we say that $\{P, Q\}$ is a *right pair*, otherwise a *left pair*. Whether P, Q is a right or left pair does not dependent on the order of P and Q . This classification as right or left is called the *chirality* of a pair of planes. We will use this classification only for isoclinic planes. When $\alpha = \beta = \pi/2$ (completely orthogonal planes) or $\alpha = \beta = 0$ (identical planes), the pair of planes is both a left pair and right pair.

Relations to Clifford Parallelism. A pair (C, D) of great circles in a 3-sphere \mathbb{S}^3 is called Clifford parallel if for any point p in C , the shortest distance from p to D is the same. A pair (C, D) of great circles is Clifford parallel if and only if the pair (P, Q) of planes spanned by C and D is isoclinic. This relation is stated as the following proposition.

Proposition 2. *For a given pair of planes P and Q , the following three statements are equivalent.*

1. P and Q are isoclinic.
2. The unit circles C on P and D on Q are Clifford parallel.
3. The projection C' of a unit circle C on P to Q is a circle and so is the projection D' of a unit circle D on Q to P .

Proof. [1 \Leftrightarrow 2.] The planes P and Q are at angle α, α if and only if there are an infinite number of pairs of half lines, each pair of which consists of a half line Op on P and a half line Oq on Q that form an angle α . Also, these pairs of lines $\{Op, Oq\}$ define the pairs of points in C and D with the shortest distance. For any pair $\{Op, Oq\}$, the intersection p' of Op and C , the intersection q' of Oq and D , and the origin form congruence triangles. Then, the distance between p' and q' is constant for any points p' on C and q' on D . Thus, the planes P and Q are at angle at α, α if and only if unit circles C, D on P and Q are Clifford parallel.

[1 \Leftrightarrow 3.] In the previous paragraph, it is explained why the projections C' and D' generally form ellipses. Similarly to the previous argument, points in Op are orthogonally projected to points in Oq for the previously defined pairs $\{Op, Oq\}$ since any plane perpendicular to both P and Q cuts out P and Q with the same angle α . Therefore, the projections of any points in C to Q are equidistant from the origin and so are the projections of any points in D to P . Hence, P and Q are isoclinic if and only if the projection C' and D' are circles. \square

The terminology for a pair of planes carries over to a pair of great circles, such as an angle of great circles, a right pair of great circles, and a left pair of great circles. For example a pair of great circles in a 3-sphere \mathbb{S}^3 is at angle (α, β) if the pair of planes spanned by them is at angle (α, β) .

A right/left pair of great circles are also known as Clifford parallel of the first/second kind [7, 8, 23]. Clifford parallelism is originally defined for lines in *elliptic geometry*. Three-dimensional elliptic geometry can be modeled as the space $\mathbb{S}^3/\mathbb{Z}_2$, i.e., the 3-sphere with opposite points identified. The great circles on the 3-sphere become the *lines* of elliptic 3-space. We are primarily concerned with the geometry of the 3-sphere, and we will apply the notion of Clifford parallelism to great circles.

Generating Clifford-parallel Circles. The following useful lemma shows how to parametrize Clifford-parallel circles. This lemma is used in Section 8.3.

Lemma 3. *Let v_1, v_2, v_3, v_4 be a positively-oriented orthonormal basis and let C be the great circle in a 3-sphere \mathbb{S}^3 spanned by two orthonormal vectors v_1 and v_2 . If C, D forms a right pair of great circles in \mathbb{S}^3 at angle (α, α) , then D is spanned by the two orthonormal vectors*

$$\begin{aligned} u_1(\delta) &= v_1 \cos \alpha + (v_3 \cos \delta + v_4 \sin \delta) \sin \alpha \\ u_2(\delta) &= v_2 \cos \alpha + (v_4 \cos \delta - v_3 \sin \delta) \sin \alpha \end{aligned} \tag{2}$$

for some $\delta \in [0, 2\pi)$, and it can be parameterized as

$$\begin{aligned} \cos \gamma \cdot u_1(\delta) + \sin \gamma \cdot u_2(\delta) &= \\ v_1 \cos \alpha \cdot \cos \gamma + v_2 \cos \alpha \cdot \sin \gamma + v_3 \sin \alpha \cdot \cos(\delta + \gamma) + v_4 \sin \alpha \cdot \sin(\delta + \gamma) \end{aligned}$$

with parameter $0 \leq \gamma < 2\pi$. Similarly, for a left pair, D is spanned by

$$\begin{aligned} u'_1(\delta) &= v_1 \cos \alpha + (v_3 \cos \delta + v_4 \sin \delta) \sin \alpha \\ u'_2(\delta) &= v_2 \cos \alpha + (v_3 \sin \delta - v_4 \cos \delta) \sin \alpha \end{aligned}$$

for some $\delta \in [0, 2\pi)$.

Proof. The circle D can be written as $\{ \cos \gamma \cdot u_1 + \sin \gamma \cdot u_2 \mid 0 \leq \gamma < 2\pi \}$, for any choice of two orthonormal vectors $u_1, u_2 \in D$. In terms of the basis v_1, v_2, v_3, v_4 , we may write $u_1 = \sum_{i=1}^4 a_i v_i$ and $u_2 = \sum_{i=1}^4 b_i v_i$ with

$$\sum_{i=1}^4 a_i^2 = \sum_{i=1}^4 b_i^2 = 1, \quad \sum_{i=1}^4 a_i b_i = 0. \quad (3)$$

The projection of D on the $v_1 v_2$ -plane containing C is then the curve $\cos \gamma \cdot (a_1 v_1 + a_2 v_2) + \sin \gamma \cdot (b_1 v_1 + b_2 v_2)$. By the definition isoclinic planes of angle α, α , this projection must be a circle of radius $\cos \alpha$. Thus, the two vectors $a_1 v_1 + a_2 v_2$ and $b_1 v_1 + b_2 v_2$ must be perpendicular vectors of length $\cos \alpha$. We may choose u_1 in such a way that its projection $a_1 v_1 + a_2 v_2$ becomes the vector $v_1 \cos \alpha$. This implies $a_1 = \cos \alpha, a_2 = 0, b_1 = 0, b_2 = \pm \cos \alpha$. By flipping u_2 if necessary, we can achieve that $b_2 = +\cos \alpha$.

We have $a_1^2 + a_2^2 = b_1^2 + b_2^2 = \cos^2 \alpha$ and $a_1 b_1 + a_2 b_2 = 0$. In view of (3), this implies $a_3^2 + a_4^2 = b_3^2 + b_4^2 = \sin^2 \alpha$ and $a_3 b_3 + a_4 b_4 = 0$. Thus, the two vectors (a_3, a_4) and (b_3, b_4) are orthogonal vectors of norm $\sin \alpha$. The general form of such vectors is $(a_3, a_4) = \sin \alpha \cdot (\cos \delta, \sin \delta)$ and $(b_3, b_4) = \pm \sin \alpha \cdot (-\sin \delta, \cos \delta)$. It can be checked that the choice of the positive sign leads to a right pair, and the other choice leads to a left pair. \square

7 Plücker Space and Plücker Distance

A classical and effective tool of dealing with linear subspaces in a vector space is the Grassmannian and the Plücker embedding. The *Grassmannian* $\mathbb{G}(k, V)$ is the collection of all k -dimensional linear subspaces of a vector space V . The *Plücker embedding* embeds Grassmannians into higher-dimensional projective space. In this paper, we only need the Grassmannian $\mathbb{G}(2, \mathbb{R}^4)$. Its Plücker coordinates are six-dimensional homogeneous coordinates in real projective 5-space.

The Grassmannian $\mathbb{G}(2, \mathbb{R}^4)$ is the set of all the planes going through the origin in Euclidean 4-space. This is equivalent to lines in real projective 3-space. Let $x = (x_0, x_1, x_2, x_3), y = (y_0, y_1, y_2, y_3)$ be homogeneous coordinates of two distinct points in real projective 3-space. Then the Plücker vector for a line going through these two points is the sextuple $(p_{ij})_{0 \leq i < j \leq 3}$ where $p_{ij} = x_i y_j - x_j y_i$; in other words, the sextuple of the 2×2 determinants of the 2×4 matrix with rows x and y . These coordinates are determined only up to scaling, and they are regarded as homogeneous coordinates in projective 5-space. For convenience, we call real projective 5-space with Plücker coordinates *Plücker space*. Each plane in Euclidean 4-space corresponds to a point in Plücker space.

To get a metric for 2-planes in Euclidean 4-space, we normalize the Plücker coordinates so that the norm of the Plücker vector becomes 1. This represents each plane by two antipodal points on the unit 5-sphere \mathbb{S}^5 , since the normalized Plücker vector is determined only up to a sign change. The *Plücker distance* between two planes P, Q in 4-space is defined as the Euclidean distance between normalized Plücker coordinates of P and Q , choosing representative points from each antipodal pair such that the distance becomes smallest.

The Plücker distance is a metric since it is basically a Euclidean distance. The next lemma shows that the Plücker distance is geometrically meaningful and does not depend on the choice of a coordinate system.

Lemma 4. *If a pair of planes P, Q is at angle (α, β) , with $0 \leq \alpha, \beta \leq \pi/2$, their Plücker distance is $\sqrt{2(1 - \cos \alpha \cos \beta)}$.*

Proof. By Proposition 1, we can choose an orthonormal basis $u = (u_1, u_2, u_3, u_4), v = (v_1, v_2, v_3, v_4)$ for P and an orthonormal basis $u' = (u'_1, u'_2, u'_3, u'_4), v' = (v'_1, v'_2, v'_3, v'_4)$ for Q such that u', v' are parallel to the projections of u, v to Q . (If $\alpha = \pi/2$, the projection of u to Q is the zero vector, and in this case we consider any vector u' as parallel to it; the same holds for $\beta = \pi/2$.)

The Plücker coordinates are then the six values $(k_{ij})_{0 \leq i < j \leq 3}$, where $k_{ij} = u_i v_j - v_i u_j$. Since u, v is an orthonormal basis, we can check that these Plücker coordinates are already normalized, using the relations $k_{ij} = -k_{ji}$ and $k_{ii} = 0$:

$$\begin{aligned} \sum_{0 \leq i < j \leq 3} k_{ij}^2 &= \frac{1}{2} \sum_{0 \leq i < j \leq 3} (k_{ij}^2 + k_{ji}^2) = \frac{1}{2} \sum_{i=0}^3 \sum_{j=0}^3 k_{ij}^2 \\ &= \frac{1}{2} \sum_{i=0}^3 \sum_{j=0}^3 (u_i v_j - v_i u_j)^2 = \frac{1}{2} \sum_{i=0}^3 \sum_{j=0}^3 (u_i^2 v_j^2 + v_i^2 u_j^2 - 2u_i u_j v_i v_j) \\ &= \frac{1}{2} \left(\sum_{i=0}^3 u_i^2 \sum_{j=0}^3 v_j^2 - 2 \sum_{i=0}^3 u_i v_i \sum_{j=0}^3 u_j v_j + \sum_{j=0}^3 u_j^2 \sum_{i=0}^3 v_i^2 \right) \\ &= \frac{1}{2} (1 - 2 \cdot 0 + 1) = 1 \end{aligned}$$

Similarly, the Plücker coordinates $\ell_{ij} = u'_i v'_j - v'_i u'_j$ of Q are normalized: $\sum_{0 \leq i < j \leq 3} \ell_{ij}^2 = 1$. The squared Euclidean distance between these points is

$$\sum_{0 \leq i < j \leq 3} (k_{ij} - \ell_{ij})^2 = \sum_{0 \leq i < j \leq 3} k_{ij}^2 + \sum_{0 \leq i < j \leq 3} \ell_{ij}^2 - 2 \sum_{0 \leq i < j \leq 3} k_{ij} \ell_{ij} = 2 \left(1 - \sum_{0 \leq i < j \leq 3} k_{ij} \ell_{ij} \right)$$

Let us evaluate the last term. From the angle between P and Q , we get $u^T u' = \cos \alpha$, $v^T v' = \cos \beta$, and $u^T v' = v^T u' = 0$, and we have $k_{ij} = -k_{ji}$ and $\ell_{ij} = -\ell_{ji}$.

$$\begin{aligned} \sum_{0 \leq i < j \leq 3} k_{ij} \ell_{ij} &= \frac{1}{2} \sum_{0 \leq i < j \leq 3} (k_{ij} \ell_{ij} + k_{ji} \ell_{ji}) = \frac{1}{2} \sum_{i=0}^3 \sum_{j=0}^3 k_{ij} \ell_{ij} \\ &= \frac{1}{2} \sum_{i=0}^3 \sum_{j=0}^3 (u_i v_j - v_i u_j)(u'_i v'_j - v'_i u'_j) \\ &= \frac{1}{2} \sum_{i=0}^3 \sum_{j=0}^3 (u_i u'_i v_j v'_j + v_i v'_i u_j u'_j - u_i v'_i u'_j v_j - v_i u'_i u_j v'_j) \\ &= \frac{1}{2} \left(\sum_{i=0}^3 u_i u'_i \sum_{j=0}^3 v_j v'_j + \sum_{i=0}^3 v_i v'_i \sum_{j=0}^3 u_j u'_j \right. \\ &\quad \left. - \sum_{i=0}^3 u_i v'_i \sum_{j=0}^3 u'_j v_j - \sum_{i=0}^3 u'_i v_i \sum_{j=0}^3 u_j v'_j \right) \\ &= \frac{1}{2} ((u^T u')(v^T v') + (v^T v')(u^T u') - (u^T v')(u'^T v) - (u'^T v)(u^T v')) \\ &= \frac{1}{2} (\cos \alpha \cos \beta + \cos \beta \cos \alpha - 0 - 0) = \cos \alpha \cos \beta \end{aligned}$$

Thus, the Euclidean distance is $\sqrt{2(1 - \cos \alpha \cos \beta)}$. Each circle is represented by two antipodal points. Thus we should consider also negated vectors $(-k_{ij})_{0 \leq i < j \leq 3}$ and $(-\ell_{ij})_{0 \leq i < j \leq 3}$. Negation can be achieved by negating one of the vectors u, v or one of the vectors u', v' . In the formulas, this changes $\cos \alpha$ to $\cos(\pi - \alpha) = -\cos \alpha$ or $\cos \beta$ to $\cos(\pi - \beta) = -\cos \beta$. The minimum distance is obviously achieved when choosing the positive cosine values, i.e., when choosing the angles in the range $0 \leq \alpha, \beta \leq \pi/2$. This finishes the proof.

This smaller distance value together with the larger value, $\sqrt{2(1 + \cos \alpha \cos \beta)}$, form the sides of a rectangle inscribed in the unit circle, as expected for two pairs of antipodal points on a sphere. \square

Since the angle between two planes is invariant under rotations, we get the following direct consequence:

Corollary 5. *The Plücker distance is invariant under rotations and reflections.* \square

7.1 Other Distances

Conway, Hardin and Sloan [11] considered different distances in Grassmannian space. One of the distances that they considered is called the *chordal distance*. For the Grassmannian $\mathbb{G}(2, \mathbb{R}^4)$, the chordal distance is defined as $\sqrt{\sin^2 \alpha + \sin^2 \beta}$ for a pair of planes at angle α, β . The chordal distance coincides with Plücker distance if a pair of planes is isoclinic. They also considered the *geodesic distance* [33], and *Asimov's distance* [3, 17].

Those three distances are related to the angles between k -dimensional subspaces. As mentioned above, two k -dimensional subspaces define k angles $(\alpha_1, \dots, \alpha_k)$. (If $k > d/2$, some of these angles are necessarily 0.) the geodesic distance, the chordal distance, and Asimov's distance are defined as the Euclidean norm of this angle vector, the Euclidean norm of vector $(\sin \alpha_1, \dots, \sin \alpha_k)$, and the largest angle, respectively. The geodesic distance has its name because it represents the true geodesic distance in the manifold $\mathbb{G}(k, \mathbb{R}^d)$.

Dense packings of subspaces in $\mathbb{G}(k, \mathbb{R}^d)$ for these three distances were experimentally determined in [11].

For our algorithm, it is convenient if we can embed $\mathbb{G}(2, \mathbb{R}^4)$ into some Euclidean space. Apart from this, the precise characteristics of the distance don't matter. The Plücker distance is measured in an ambient space of dimension $\binom{d}{k}$. On the other hand, the chordal distance embeds the Grassmannian $\mathbb{G}(k, \mathbb{R}^d)$ to an m -dimensional sphere, where $m = \binom{d+1}{2} - 1$. (It represents a subspace by the $d \times d$ symmetric projection matrix onto that subspace.) Thus, for large k and d , this distance is preferable over the Plücker distance, for which the dimension grows exponentially. However, for our parameters $d = 4$ and $k = 2$, the chordal distance would require a 9-sphere, whereas the Plücker distance requires only a 5-sphere. Since the ambient dimension affects the complexity of the algorithm, the Plücker distance is preferable in our case.

8 The Construction of Hopf Fibrations

For the well-known properties of the Hopf fibration of the 3-sphere, see for example [7, Section 4.3.7] and [8, Chapter 18.8]. The goal of this section is to understand the relations among isoclinic rotations (Section 5), a left/right pair of great circles (Section 6) and Hopf fibrations.

Before starting, we sketch the flow of this section without the definitions of terms:

1. We first construct *the invariant family* for a given great circle C by using only isoclinic rotations (Lemmas 6 and 7).
2. Then, we define a *Hopf map* h with respect to C . This shows that the invariant family for C is equivalent to the *Hopf bundle* for h (Lemma 8). This gives a geometric perspective of constructing a *Hopf fibration*.
3. In Section 8.3, we derive the following properties.
 - (a) From 1, we know there is a unique bundle for each great circle C (Corollary 9). Also, a pair of Clifford-parallel great circles belong to a common Hopf bundle (Lemma 10).
 - (b) The relation between two great circles defined by a right pair (or a left pair) is transitive; thus, it is an equivalence relation. From 3a, an equivalence class is actually a right Hopf bundle (or a left Hopf bundle) (Corollary 11). The converse is also true.
 - (c) The Hopf fibration maps two great circles at angle α, α to two points at geodesic distance 2α in a 2-sphere (Lemma 12).

To minimize confusion, we provide proofs and explanations only for a right Hopf fibration but parallel lemmas are also valid for a left Hopf fibration.

8.1 Construction by Isoclinic Rotations

This construction of a right Hopf fibration provides the relation between a right Hopf fibration and a right rotation. We will now construct a right Hopf bundle as the set of orbits of all right rotations that map a given circle to itself.

We mention that the convention of assigning the label left or right to versions of a structure that come in pairs is somewhat arbitrary and not uniform in the literature. In fact, right Hopf bundles are also connected to *left* rotations: a right Hopf bundle is the set of images of a given circle under all left rotations. In the quaternion representation, right rotations according to the convention in this paper are carried out by left multiplication with unit quaternions. Thus, there are good reasons also for alternative choices.

Lemma 6. *For every great circle C in a 3-sphere and two points $p, q \in C$, there exists a unique right rotation R which rotates C in itself and rotates p to q .*

Proof. Choose an orthonormal basis v_1, v_2 of the plane P spanned by C . There exists a unique rotation $\begin{pmatrix} \cos \gamma & -\sin \gamma \\ \sin \gamma & \cos \gamma \end{pmatrix}$ in the $v_1 v_2$ -plane that rotates p to q . Extend v_1, v_2 to a positively-oriented basis v_1, v_2, v_3, v_4 . It follows that, in this basis, the rotation matrix R must have the form

$$\rho(\gamma) = R_{\gamma, \gamma}. \quad (4)$$

The reason is as follows: The first two columns are fixed by the requirement that P maps p to q and leaves the plane through C fixed. The last two columns are then fixed by the requirement that R is a right rotation. \square

Actually, if $p \neq \pm q$, the requirement that R rotates C in itself is redundant.

The family of right rotations $\mathcal{F} = \{\rho(\gamma) | \gamma \in [0, 2\pi)\}$ rotates a great circle C to itself where $\rho(\gamma)$ is defined in (4). Note that \mathcal{F} is a one-dimensional group isomorphic to the special orthogonal group $\text{SO}(2)$. We call \mathcal{F} the *family of right rotations* for C .

Lemma 7. *Let \mathcal{F} be the family of right rotations for some great circle in a 3-sphere \mathbb{S}^3 . For every point p in \mathbb{S}^3 , the orbit of p under rotations in \mathcal{F} is a great circle of \mathbb{S}^3 .*

Proof. If $p = (x, y, z, w)$, the orbit of p generated by a right rotation is

$$\mu(\gamma) = \rho(\gamma)p = u_1 \cos \gamma + u_2 \sin \gamma \quad (5)$$

where

$$u_1 = p = (x, y, z, w)^T, u_2 = (-y, x, -w, z)^T, \rho(\gamma) \in \mathcal{F}$$

for $\gamma \in [0, 2\pi)$. Here, u_1 and u_2 are orthogonal and $\mu(\gamma)$ is a parametrization of a great circle, so p generates a great circle as an orbit. \square

From Lemma 7, \mathcal{F} partitions \mathbb{S}^3 into the set of great circles that are invariant under rotations in \mathcal{F} . Let Γ be the set of these great circles and call it the *invariant family* for C . Eventually, the invariant family Γ forms a right *Hopf bundle*. In the remainder of this section, we illustrate what this statement means more precisely.

8.2 Equivalence of an Invariant Family and a Hopf Bundle.

A *fibration* is a map that projects a *fiber bundle* to a *base*, by identifying a subspace of the fiber bundle, called a *fiber*, to a point in a base in a way that any point in the fiber bundle is contained in exactly one fiber.

Given a great circle C , by choosing a Cartesian coordinate system (x, y, z, w) such that C lies in the x, y -plane, we define the right Hopf map with respect to C as

$$h(x, y, z, w) = (2(xw - yz), 2(yw + xz), 1 - 2(z^2 + w^2)) \quad \text{for } (x, y, z, w) \in \mathbb{S}^3 \subset \mathbb{R}^4$$

For reference, the left Hopf map is defined as $h'(x, y, z, w) = (2(xw + yz), 2(yw - xz), 1 - 2(z^2 + w^2))$.

The following lemma shows that a right Hopf map is a fibration.

Lemma 8. *Let C be a great circle in a 3-sphere \mathbb{S}^3 . Let Γ be the invariant family for C . Let h be the right Hopf map with respect to C . Then, for any $D \in \Gamma$, h maps all points in D to one point in the 2-sphere \mathbb{S}^2 .*

Proof. The map h indeed maps the 3-sphere to the 2-sphere, since $(1 - 2z^2 - 2w^2)^2 + 4(xw - yz)^2 + 4(yw + xz)^2 = (x^2 + y^2 + z^2 + w^2)^2 = 1$.

As we have seen in Lemma 7, for any $D \in \Gamma$, if $p = (x, y, z, w) \in D$, we can parametrize D as $\mu(\gamma)$ in (5) for $\gamma \in [0, 2\pi)$. Then,

$$\begin{aligned} h(\mu(\gamma)) &= h((x, y, z, w) \cos \gamma + (-y, x, -w, z) \sin \gamma) \\ &= h((x \cos \gamma - y \sin \gamma, y \cos \gamma + x \sin \gamma, z \cos \gamma - w \sin \gamma, w \cos \gamma + z \sin \gamma)) \\ &= (2[(x \cos \gamma - y \sin \gamma)(w \cos \gamma + z \sin \gamma) - (y \cos \gamma + x \sin \gamma)(z \cos \gamma - w \sin \gamma)], \\ &\quad 2[(y \cos \gamma + x \sin \gamma)(w \cos \gamma + z \sin \gamma) + (x \cos \gamma - y \sin \gamma)(z \cos \gamma - w \sin \gamma)], \\ &\quad 1 - 2[(z \cos \gamma - w \sin \gamma)^2 + (w \cos \gamma + z \sin \gamma)^2]) \\ &= (2[(xw \cos^2 \gamma - yz \sin^2 \gamma + (xz - yw) \cos \gamma \sin \gamma) \\ &\quad - (yz \cos^2 \gamma - xw \sin^2 \gamma + (xz - yw) \cos \gamma \sin \gamma)], \\ &\quad 2[(yw \cos^2 \gamma + xz \sin^2 \gamma + (xw + yz) \cos \gamma \sin \gamma) \\ &\quad + (xz \cos^2 \gamma + yw \sin^2 \gamma - (xw + yz) \cos \gamma \sin \gamma)], \\ &\quad 1 - 2[(z^2 \cos^2 \gamma + w^2 \sin^2 \gamma - zw \cos \gamma \sin \gamma) \\ &\quad + (w^2 \cos^2 \gamma + z^2 \sin^2 \gamma + zw \cos \gamma \sin \gamma)]) \\ &= (2(xw - yz), 2(yw + xz), 1 - 2(z^2 + w^2)) \end{aligned}$$

This is independent of γ , which shows that D is a fiber.

This also implies that a right Hopf map is a fibration, since elements in Γ , or fibers, are mutually disjoint and cover the whole \mathbb{S}^3 . Remember that in Lemma 7 every point in \mathbb{S}^3 is in a unique orbit, so in a unique element of Γ , under rotations in \mathcal{F} . \square

We have shown that the right Hopf map with respect to a great circle C is a fibration that maps an invariant family Γ for C to a 2-sphere \mathbb{S}^2 in a way that each great circle in Γ is mapped to a point in \mathbb{S}^2 . Hereafter, we are allowed to call a Hopf map a Hopf fibration, the fiber bundle Γ the Hopf bundle, and each great circle in Γ a Hopf fiber. Also, the 2-sphere which is a base for a Hopf fibration is called a base sphere. We speak of a right Hopf fibration/bundle/fiber if a Hopf fibration/bundle/fiber was induced by a right Hopf map.

8.3 Properties of the Hopf Fibration

Because the invariant family for a great circle C is equivalent to a right Hopf bundle, we may as well state Lemma 6 as follows:

Corollary 9. *Every great circle belongs to a unique right Hopf bundle.*

Next, we show the relation between great circles and a Hopf fibration.

Lemma 10. *Any right pair of great circles belongs to a common right Hopf bundle.*

Proof. Let C, D be a right pair of great circles at angle α, α . Let h be a right Hopf map for C . With the same basis for h , C is represented by $\{(\cos \gamma, \sin \gamma, 0, 0) | \gamma \in [0, 2\pi)\}$. Then h maps C to the north pole since $h(\cos \gamma, \sin \gamma, 0, 0) = (0, 0, 1)$. The circle D can be represented by $\{(\cos \alpha \cdot \cos \gamma, \cos \alpha \cdot \sin \gamma, \sin \alpha \cdot \cos(\delta + \gamma), \sin \alpha \cdot \sin(\delta + \gamma)) | \gamma \in [0, 2\pi)\}$ for some $\delta \in [0, 2\pi)$ by Lemma 3. Thus, h maps D to a point $(\sin 2\alpha \sin \delta, \sin 2\alpha \cos \delta, \cos 2\alpha)$, since

$$\begin{aligned} & h(\cos \alpha \cos \gamma, \cos \alpha \sin \gamma, \sin \alpha \cos(\delta + \gamma), \sin \alpha \sin(\delta + \gamma)) \\ &= \left(2(\cos \alpha \sin \alpha \cos \gamma \sin(\delta + \gamma) - \cos \alpha \sin \alpha \sin \gamma \cos(\delta + \gamma)), \right. \\ & \quad \left. 2(\cos \alpha \sin \alpha \sin \gamma \sin(\delta + \gamma) + \cos \alpha \sin \alpha \cos \gamma \cos(\delta + \gamma)), \right. \\ & \quad \left. 1 - 2(\sin^2 \alpha \cos^2(\gamma + \gamma_0) + \sin^2 \alpha \sin^2(\delta + \gamma)) \right) \\ &= (\sin 2\alpha \sin \delta, \sin 2\alpha \cos \delta, \cos 2\alpha). \end{aligned} \tag{6}$$

This is independent of the parameter γ . Therefore, the right Hopf bundle given by h contains both C and D as Hopf fibers. \square

This relation of being a right pair is actually transitive.

Corollary 11. *Let us denote $C \parallel_+ D$, when two great circles C and D in \mathbb{S}^3 form a right pair. Then \parallel_+ is transitive. Each equivalence class of \parallel_+ is a right Hopf bundle.*

Proof. If $C \parallel_+ D$ and $D \parallel_+ E$, there is a unique right Hopf fibration containing C, D and E by Corollary 9 and Lemma 10. Then, C, D , and E are contained in the same Hopf fibration, so $C \parallel_+ E$. Observe that \parallel_+ is also reflexive and symmetric. The equivalence class is exactly the invariant family for C , or a right Hopf bundle. \square

In our construction, the Hopf map h depends on the choice of a great circle C and a Cartesian coordinate system (x, y, z, w) . However, the Hopf fibration resulting from the map is independent of this choice. To show this, we show how a pair of Hopf fibers is mapped to a pair of points in a base sphere. A similar statement can be found in [8, Exercise 18.11.18].

Lemma 12. *Two right Hopf fibers at angle α, α are mapped to a pair of points of geodesic distance 2α by the right Hopf map.*

Proof. By (6) in Lemma 10, the great circle that is mapped to the point with spherical coordinates (γ, δ) , $\begin{pmatrix} \sin 2\gamma \sin \delta \\ \sin 2\gamma \cos \delta \\ \cos 2\gamma \end{pmatrix}$ by the right Hopf map is spanned by the two orthonormal vectors

$$u(\gamma, \delta) = \begin{pmatrix} \cos \gamma \\ 0 \\ \cos \delta \sin \gamma \\ \sin \delta \sin \gamma \end{pmatrix}, \quad v(\gamma, \delta) = \begin{pmatrix} 0 \\ \cos \gamma \\ -\sin \delta \sin \gamma \\ \cos \delta \sin \gamma \end{pmatrix}.$$

Let us calculate the angle α between two great circles C and D with spherical coordinates (γ_1, δ_1) and (γ_2, δ_2) . Rotating the 2-sphere by varying the parameter δ corresponds to a rotation of the v_3v_4 -plane in \mathbb{R}^4 , which leaves the angles between C and D unchanged. Thus, by an appropriate rotation, we may assume that $\delta_1 = 0$.

The length of the projection of any point of C onto the plane spanned by D is $\cos \alpha$. We project $u(\gamma_1, 0)$ on the plane with orthonormal basis $u(\gamma_2, \delta_2)$ and $v(\gamma_2, \delta_2)$, and get

$$\begin{aligned}
\cos^2 \alpha &= \langle u(\gamma_1, 0), u(\gamma_2, \delta_2) \rangle^2 + \langle u(\gamma_1, 0), v(\gamma_2, \delta_2) \rangle^2 \\
&= (\cos \gamma_1 \cos \gamma_2 + \cos \delta_2 \sin \gamma_1 \sin \gamma_2)^2 + (-\sin \delta_2 \sin \gamma_1 \sin \gamma_2)^2 \\
&= \cos^2 \gamma_1 \cos^2 \gamma_2 + \sin^2 \gamma_1 \sin^2 \gamma_2 + 2 \cos \delta_2 \cos \gamma_1 \cos \gamma_2 \sin \gamma_1 \sin \gamma_2 \\
&= \cos^2 \gamma_1 \cos^2 \gamma_2 + (1 - \cos^2 \gamma_1)(1 - \cos^2 \gamma_2) + \frac{1}{2} \cos \delta_2 \sin 2\gamma_1 \sin 2\gamma_2 \\
&= 2 \cos^2 \gamma_1 \cos^2 \gamma_2 + 1 - \cos^2 \gamma_1 - \cos^2 \gamma_2 + \frac{1}{2} \cos \delta_2 \sin 2\gamma_1 \sin 2\gamma_2 \\
&= \frac{1}{2}(2 \cos^2 \gamma_1 - 1)(2 \cos^2 \gamma_2 - 1) + \frac{1}{2} + \frac{1}{2} \cos \delta_2 \sin 2\gamma_1 \sin 2\gamma_2 \\
&= \frac{1}{2}(\cos 2\gamma_1 \cos 2\gamma_2 + 1 + \cos \delta_2 \sin 2\gamma_1 \sin 2\gamma_2)
\end{aligned} \tag{7}$$

We get

$$\cos 2\alpha = 2 \cos^2 \alpha - 1 = \cos 2\gamma_1 \cos 2\gamma_2 + \sin 2\gamma_1 \sin 2\gamma_2 \cos \delta_2 \tag{8}$$

from (7).

On the other hand, we compute the spherical distance c between the two points in the base sphere corresponding to $(\gamma_1, 0)$ and (γ_2, δ_2) . It is the third side of a spherical triangle with angle $\delta_2 - 0 = \delta_2$ at the north pole and sides $a = 2\gamma_1$ and $b = 2\gamma_2$. By the spherical cosine law,

$$\begin{aligned}
\cos c &= \cos a \cos b + \sin a \sin b \cos \delta_2 \\
&= \cos 2\gamma_1 \cos 2\gamma_2 + \sin 2\gamma_1 \sin 2\gamma_2 \cos \delta_2
\end{aligned}$$

From this and (8), we conclude that $c = 2\alpha$.

Therefore, the angular distance between C and D after applying the Hopf map is 2α . \square

We have defined the Hopf map with respect to an arbitrarily chosen great circle C from an invariant family Γ . The previous lemma shows that this choice is not essential. Choosing a different circle $C' \in \Gamma$ will lead to a different Hopf map, but the two images are related by an isometry of the 2-sphere.

The following theorem summarizes the relations of right pairs and isoclinic rotations. The analogous statement is valid for left pairs as well.

Theorem 13. *For a pair of great circles C and D , C and D is a right-isoclinic pair if and only if D is an orbit under all right rotations that leave C invariant.*

Proof. By Lemma 6 and Lemma 7, there exists the family \mathcal{F} of right rotations that leave C invariant, and orbits of \mathcal{F} are great circles of a 3-sphere. By Lemma 8, the set of these orbits is equivalent to the set of fibers in the (unique) right Hopf bundle containing C . Thus, a circle D is an orbit of \mathcal{F} if and only if C and D are in the common right Hopf bundle. Lemma 10 implies that C and D are indeed in the common right Hopf bundle. \square

9 Closest-Pair Graphs

The closest-pair graph G of a given point set P is a graph G such that each point in P is a vertex of G and a pair of vertices p and q in P forms an edge if and only if the distance between p and q achieves the minimum distance among all the pairs in P . The closest-pair graph should not be confused with the nearest-neighbor graph, a directed graph where an edge from p to q indicates that q is a nearest-neighbor of p . The closest-pair graph may have isolated vertices, but it must contain at least one edge.

The closest-pair graph has the following properties:

1. It can be constructed in time $O(n \log n)$ by divide and conquer [6] in any fixed dimension,

2. the degree of each vertex is bounded by the “kissing number”, as described earlier in Section 11, and consequently
3. the number of edges is at most linear in the number of vertices.

By taking advantage of these characteristics, closest pairs were already used in previous work [5, 10], see Section 4.2. We will compute closest-pair graphs not only for the input point sets in 4-space, but also for sets of planes in Plücker space, see Section 7.

10 The Coxeter Classification of Four-Dimensional Point Groups

The new algorithm treats the case that a given point set can be generated by a series of reflections as a special case. We use the classification of Coxeter groups to argue that the cardinalities of point sets that have a high degree of symmetry are bounded. To this end, this section provides the classification of such groups in 4-space and shows the computation of inradii of fundamental regions of such groups.

The discrete groups generated by mirror reflections have been classified in arbitrary dimension, cf. [13, Table IV on p. 297]. Each of these groups of rank d can be represented by a collection of mirrors r_1, \dots, r_d such that $r_i \cdot r_i = 1$ and $(r_i r_j)^{m_{ij}} = 1$ for $i \neq j$ and $m_{ij} \geq 2$; such a group is called a Coxeter group. The value of m_{ij} specifies the dihedral angle between mirrors r_i and r_j as π/m_{ij} . The Coxeter diagram, or the Coxeter graph, explicitly encodes these values of m_{ij} .

In four dimensions, Coxeter groups were first enumerated by Édouard Goursat in 1899: There are five *irreducible groups*, which are the symmetry groups of the five regular 4-dimensional polytopes, called A_4, C_4, B_4, F_4 and G_4 according to [13] (or alternatively A_4, BC_4, D_4, F_4 , and H_4 in today’s terms). Here, A_m denotes a group related to reflectional symmetries of an m -simplex. Similarly, C_m is for an m -hypercube and an m -cross-polytope, B_m is for an m -demihypercube. F_4 represents the reflectional symmetry group for a 24-cell and G_4 is for a 120-cell and a 600-cell.

In addition, there are the *reducible* groups, direct products of lower-dimensional reflection groups. They are the groups $A_3 \times A_1, C_3 \times A_1, G_3 \times A_1$, and $D_2^p \times D_2^q$, for $p, q \geq 2$, where G_3 is the symmetry group of the icosahedron and the dodecahedron, C_3 is the symmetry group of the cube and the octahedron, and D_2^p is the dihedral group of order $2p$ (alternatively denoted as $I_2(p)$), the symmetry group of the regular p -gon. Some groups have alternative representations: $D_3 = C_3, D_2^2 = A_1 \times A_1, D_2^3 = A_2$, and $D_2^4 = C_2$. For reducible groups, their mirrors fall into two or more classes such that each mirror in one class is perpendicular to the mirrors in the other classes. Since reflections at perpendicular mirrors commute, this is how reducible groups yield the decomposition into a direct product of smaller groups.

The arrangement of all mirror hyperplanes of a reflection group cuts the 3-sphere into equal cells, which can be taken as the *fundamental regions* of the group. These fundamental regions are not necessarily equal to the Voronoi regions of the point set; the Voronoi regions are usually cut into smaller cells by mirrors passing through the centers of the regions. It is known that the fundamental region of a reflection group is a spherical *simplex* [13, Theorem 11.23]. Thus, in 4-space, the fundamental region is a spherical tetrahedron T , and the group is generated by four independent reflections.

The inradii of fundamental regions of Coxeter groups are of our particular interest for Section 19.

10.1 The Radius of an Inscribed Sphere of a Fundamental Region

By looking at the Coxeter diagrams (Table 1 and [13, Table IV on p. 297]), we can read off the dihedral angles between bounding mirrors of fundamental regions of Coxeter groups by the




Figure 3: Finding the radius R_0 of an inscribed 0-sphere in a spherical 1-simplex T . This two-dimensional sketch of the situation holds analogously in higher dimensions.

following conventions. Each vertex in a Coxeter diagram represents a mirror. A pair of vertices with no edge means a pair of perpendicular mirrors. An edge with no number connects a pair of mirrors whose dihedral angle is $\pi/3$. Edges with number p represent a dihedral angle π/p .

Among all four-dimensional Coxeter groups, we are interested in the finite list, excluding the infinite family $D_2^p \times D_2^q$; that is, $A_4, C_4, B_4, F_4, G_4, A_3 \times A_1, C_3 \times A_1$, and $G_3 \times A_1$. Given the dihedral angles, normal vectors of the vertices of a spherical tetrahedron T can be computed: when the dihedral angle between two mirrors i and j is π/m_{ij} ,

$$\hat{n}_i \cdot \hat{n}_j = \cos(\pi - \pi/m_{ij}) = -\cos(\pi/m_{ij})$$

where \hat{n}_i, \hat{n}_j are unit normal vectors to mirrors i and j and \cdot is an inner product.

Table 1 enumerates such normal vectors for the above groups. The unit normal vector of the i -th bounding mirror in a diagram is in the i -th row. For example, for A_4 , the first two bounding mirrors form an angle $\pi/3$ and the corresponding unit normal vectors in the outward directions from the fundamental region form an angle $\pi - \pi/3$. That is why the inner product of the first two normal vectors for A_4 is $\cos(2\pi/3) = -\cos(\pi/3) = -1/2$.

We call the distance from the center of the inscribed sphere of a spherical tetrahedron T to the mirrors *the Euclidean radius* of T . We compute the Euclidean radius R_0 of each spherical tetrahedron T that is a fundamental region of one of above groups as follows.

Each bounding mirror r_i of T spans a hyperplane H_i going through the origin for $i = 1, 2, 3, 4$. After translating each H_i in the orthogonal direction of H_i toward the interior of the fundamental region by distance 1, let p be the intersection point of these translated hyperplanes. This point has equal distance 1 from all hyperplanes, but it does not lie on the unit sphere. Let δ_h be the distance of p from the origin O . Then, by rescaling the segment \overline{Op} by $1/\delta_h$ we get the center c of the inscribed sphere, and the Euclidean radius of T is $R_0 = 1/\delta_h$, see Figure 3. The approximated values of the Euclidean radii R_0 of fundamental regions of 4-dimensional finite Coxeter groups are computed in this way. These values are given in Table 1.

The smallest Euclidean radius occurs for G_4 , and its value is $R_0 \geq 0.039102328$.

We are interested in this radius for the following reason. Our algorithm will build the closest-pair graph G of a set of points in \mathbb{S}^3 . One special case arises when every edge uv of G acts as a mirror, exchanging u and v together with all their neighbors in G . From this it follows that every component of G is an orbit of some point u_0 under the reflection group *generated* by the mirrors perpendicular to the edges incident to u_0 . The above considerations allow us to bound the closest-pair distance of G in such cases:

Thus, the minimum distance $2R_0 \geq 0.0782$ as it is claimed in the following lemma.

Lemma 14. *Consider the orbit I of a point u_0 on the unit sphere \mathbb{S}^3 under a finite four-dimensional symmetry group H which is not $D_2^p \times D_2^q$. Let δ be the minimum distance δ between*


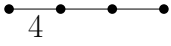
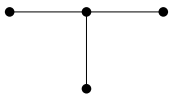
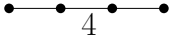
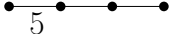

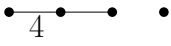
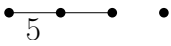
| Groups | Coxeter Diagrams | Normal Vectors | R_0 |
|------------------|---|---|----------------|
| A_4 |  | $(1, 0, 0, 0)$ $(-\frac{1}{2}, \frac{\sqrt{3}}{2}, 0, 0)$ $(0, -\frac{1}{\sqrt{3}}, \frac{\sqrt{2}}{\sqrt{3}}, 0)$ $(0, 0, -\frac{\sqrt{3}}{2\sqrt{2}}, \frac{\sqrt{5}}{2\sqrt{2}})$ | 0.2236067977 |
| C_4 |  | $(1, 0, 0, 0)$ $(-\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}, 0, 0)$ $(0, -\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}, 0)$ $(0, 0, -\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}})$ | 0.1429000737 |
| B_4 |  | $(1, 0, 0, 0)$ $(-\frac{1}{2}, \frac{\sqrt{3}}{2}, 0, 0)$ $(0, -\frac{1}{\sqrt{3}}, \frac{\sqrt{2}}{\sqrt{3}}, 0)$ $(0, -\frac{1}{\sqrt{3}}, -\frac{1}{\sqrt{6}}, \frac{1}{\sqrt{2}})$ | 0.1889822365 |
| F_4 |  | $(1, 0, 0, 0)$ $(-\frac{1}{2}, \frac{\sqrt{3}}{2}, 0, 0)$ $(0, -\frac{\sqrt{2}}{\sqrt{3}}, \frac{1}{\sqrt{3}}, 0)$ $(0, 0, -\frac{\sqrt{3}}{2}, \frac{1}{2})$ | 0.009671356812 |
| G_4 |  | $(1, 0, 0, 0)$ $(-\frac{1+\sqrt{5}}{4}, \frac{\sqrt{10-2\sqrt{5}}}{4}, 0, 0)$ $(0, -\frac{2}{\sqrt{10-2\sqrt{5}}}, \frac{\sqrt{6-2\sqrt{5}}}{\sqrt{10-2\sqrt{5}}}, 0)$ $(0, 0, -\frac{\sqrt{10-2\sqrt{5}}}{2\sqrt{6-2\sqrt{5}}}, \frac{\sqrt{14-6\sqrt{5}}}{2\sqrt{6-2\sqrt{5}}})$ | 0.03910328003 |
| $A_3 \times A_1$ |  | $(1, 0, 0, 0)$ $(-\frac{1}{2}, \frac{\sqrt{3}}{2}, 0, 0)$ $(0, -\frac{1}{\sqrt{3}}, \frac{\sqrt{2}}{\sqrt{3}}, 0)$ $(0, 0, 0, 1)$ | 0.3015113445 |
| $C_3 \times A_1$ |  | $(1, 0, 0, 0)$ $(-\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}, 0, 0)$ $(0, -\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}, 0)$ $(0, 0, 0, 1)$ | 0.2108874992 |
| $G_3 \times A_1$ |  | $(1, 0, 0, 0)$ $(-\frac{1+\sqrt{5}}{4}, \frac{\sqrt{10-2\sqrt{5}}}{4}, 0, 0)$ $(0, -\frac{2}{\sqrt{10-2\sqrt{5}}}, \frac{\sqrt{6-2\sqrt{5}}}{\sqrt{10-2\sqrt{5}}}, 0)$ $(0, 0, 0, 1)$ | 0.1303737577 |

Table 1: Coxeter diagrams, outward unit normal vectors of fundamental regions, and inradii R_0 of fundamental regions of four-dimensional Coxeter groups.

two points of I , and suppose that the group H is generated by the mirror reflections between u_0 and each point $v \in I$ at distance δ from u_0 . Then δ is at least $\delta_1 := 0.07$.

Proof. We have a finite list of possible groups H , and for each such group, we know the shape of its fundamental region. For each fundamental region T , we consider all possibilities how the orbit of a point $u_0 \in T$ might give rise the orbit I . We have to place u_0 on some facets F of T and equidistant from the remaining facets of T . Otherwise, if the point were not chosen equidistant from the other faces, the closest-pair graph would not contain edges that generate all four mirrors. The minimum distance is achieved when $F = \emptyset$ and u_0 lies in the interior of T . We have analyzed each of the eight groups case by case in Table 1, and we have seen that the minimum distance δ is $2R_0 \approx 0.0782$. \square

11 Packing Arguments and Definitions of Constants

We often rely on packing arguments to bound the complexity of certain configurations in situations where we know that the minimum distance between points or the minimum Plücker distance between planes is bounded from below and we want to show that the number of points or planes is bounded from above.

Lemma 15. *If there are m points on a d -sphere \mathbb{S}^d with minimum Euclidean distance δ , then*

$$m \cdot \kappa_d(\delta/2)^d \leq \omega_d$$

where κ_d is the volume of the Euclidean unit d -ball and ω_d is the surface area of the unit d -sphere.

Proof. The geodesic distance between the points is at least $\alpha := 2 \arcsin \frac{\delta}{2}$. Thus, if we draw a spherical d -ball of radius $\alpha/2$ around the m points, these balls form a packing on the surface of the d -sphere. The boundary of such a ball is a $(d-1)$ -sphere of Euclidean radius $\sin \frac{\alpha}{2} = \frac{\delta}{2}$ (see Figure 3 for a similar situation). The d -dimensional volume of the spherical ball must be at least the volume of the (Euclidean) d -ball with the same boundary, which is $\kappa_d(\delta/2)^d$. \square

We will now define a few constants that are used in the new algorithm:

11.1 Kissing Numbers

We denote by K_d the kissing number on the d -dimensional sphere \mathbb{S}^d : the maximum number of equal interior-disjoint balls on \mathbb{S}^d that can simultaneously touch a ball of the same size. We use the known bounds $K_2 = 5$, $K_3 = 12$, and $40 \leq K_5 \leq 44$.

These bounds on K_d are derived from the corresponding kissing numbers in Euclidean space: These kissing number are known to be 6 in the plane, 12 in 3-space, 24 in 4-space, and between 40 and 44 in 5-space. (On the 2-sphere, the kissing number is smaller than in the plane.) In the closest-pair graph, the degree of every vertex is obviously bounded by these kissing numbers.

11.2 The Closeness Threshold

We define the constant $\delta_0 := 0.0005$ for applying 1+3 dimension reduction. If the minimum distance of a point set is greater than δ_0 , the size of the point set is at most n_0 , where n_0 is the maximum packing of 3-balls of radius $\delta_0/2$ on a unit 3-sphere \mathbb{S}^3 . Then, we can apply the dimension reduction principle in Section 14 without affecting the time complexity. Otherwise, we need to go through other condensing procedures. The value of δ_0 will be justified in Lemma 18 in Section 18.

Since the surface area ω_3 of \mathbb{S}^3 is $2\pi^2$ and the volume κ_3 of the unit 3-ball is $\frac{4}{3}\pi$, by Lemma 15, a rough estimate of the upper bound n_0 is

$$n_0 := \lfloor 2\pi^2 / (\frac{4}{3}\pi(\delta_0/2)^3) \rfloor < 3.016 \times 10^{11}. \quad (9)$$

11.3 The Icosahedral Threshold

In Algorithm M, we will need the maximum number C_1 of planes at pairwise Plücker distance greater than or equal to a certain distance δ_{\min} . The distance δ_{\min} is the Plücker distance of two isoclinic planes P, Q that are mapped to two adjacent vertices of a regular icosahedron by the Hopf map.

By Lemma 4, the Plücker distance between two planes of angle α , α is $\sqrt{2} \sin \alpha$ and by Lemma 12, such two planes are mapped to a pair of points of geodesic distance 2α by the Hopf fibration. The edge length of an icosahedron with a circumscribed unit sphere is $\delta_{\text{ico}} = \sqrt{50 - 10\sqrt{5}}/5$, so the geodesic distance of two vertices in a sphere is

$$2 \arcsin(\delta_{\text{ico}}/2) \approx 1.107148718,$$

and the angle between two such planes is

$$\alpha_{\min} = \arcsin(\delta_{\text{ico}}/2) \approx 0.5535743590.$$

The Plücker distance of such pair of planes is

$$\delta_{\min} = \sqrt{2} \sin(\arcsin(\delta_{\text{ico}}/2)) \approx 0.7434960688.$$

Each plane in $\mathbb{G}(2, \mathbb{R}^4)$ is represented by two antipodal points on the 5-sphere. By Lemma 15, the upper bound can thus be computed as

$$2C_1 \leq 2\pi^3 / \left(\frac{8}{15} \pi^2 (\delta_{\min}/2)^5 \right),$$

since the surface area ω_5 of the unit 5-sphere is $2\pi^3$ and the volume κ_5 of a unit 5-ball is $\frac{8}{15}\pi^2$. This gives

$$C_1 \leq 829. \tag{10}$$

This bound is a rough estimate, and the correct bound is likely to be much smaller, for two reasons: (i) Lemma 15 uses only a crude volume argument. (ii) The Plücker coordinates cannot lie anywhere on the 5-sphere, but they are restricted to a 4-dimensional manifold, the Plücker quadric. Unfortunately, this does not directly allow us to apply Lemma 15 for $d = 4$ dimensions, because the Plücker surface is negatively curved, and moreover, we don't even know its volume.

Packings in Grassmannian spaces were considered by Conway, Hardin and Sloane [11]. They used the *chordal distance*, defined as $\sqrt{\sin^2 \alpha + \sin^2 \beta}$ for a pair of planes at angle $\{\alpha, \beta\}$ (see Section 7.1). According to their experimental results, for planes in 4-space, when there are 37 planes, the maximum chordal distance that they achieve is about 0.728633689875024. For isoclinic planes, in which we are interested here, the chordal distance coincides with the Plücker distance. This suggests that the true bound C_1 is near 37, much smaller than the bound (10).

12 Canonical Axes

The canonical-axes construction is a well-known procedure for detecting the rotational symmetries in a planar point configuration [25, 4, 18]. We use this construction in Step C11 of Algorithm C in Section 17. We can encode labeled-point sets on a circle as a cyclic string. This can be done by alternating between labels for a point and angular distances between two adjacent points [25].

In addition, by making cyclic shifts in this string so that the string becomes lexicographically smallest, but still starts with labels (not angles), we can get a unique representation of points on a circle with labels. These starting points that yield the same string give rise to a set of p equidistant rays starting from the origin. We call a collection of these points *canonical axes*. Canonical axes can be found in $O(n \log n)$ time when n is the number of labeled points by

standard string-processing techniques [24]. Canonical axes have the same rotational symmetries as the original configuration.

We assume labels are preserved under rotations. Then, if a circle K is mapped to a circle K' , canonical axes of K are mapped to canonical axes of K' , so canonical axes are equivariant under rotations.

Moreover, canonical axes are more special than most other equivariant criteria. If canonical axes of a circle K are mapped to canonical axes of another circle K' by a rotation R , K' itself is also mapped to K by R , so the converse of equivariance is also true.

In this context, we define a canonical set and a canonical set procedure as a special condensing. Refer to Section 22 about a canonical set procedure. Canonical axes are the result of a canonical set procedure for a circle and rotational symmetries on a circle. The new algorithm employs a canonical set procedure for a square torus and translational symmetries on a torus. See Section 22.

13 Congruence Types of Vertex and Edge Figures

Assume that a geometric graph G is given. A vertex figure of a vertex v , that is, v together with its neighbors, is first introduced in Atkinson's algorithm [5] (see Section 4.2). Atkinson argued that a vertex figure can be encoded in a string of length $O(1)$ if the vertex is of bounded degree, but he omitted the details about representing a vertex figure. Our new algorithm uses vertex figures and also introduces an extended concept called an edge figure of an edge uv , that is, all neighbors of u and v together with the edge uv .

We now describe the representation details of a vertex figure, and of an edge figure. We assume that a given directed graph $G = (V, E)$ is embedded in the 3-sphere \mathbb{S}^3 such that all the edge lengths are the same and G has its maximum degree at most 12. This bound follows from the fact that G is a (directed) subgraph of a closest-pair graph on the 3-sphere, and the kissing number on the 3-sphere is 12.

Let us begin with a vertex figure. If a vertex v is of degree 0 or 1, there is only one congruence type of the neighborhood of v . If a vertex v is of degree 2, the angle between the two incident edges determines its congruence type.

If a vertex v has degree greater than 2, we can choose two distinct incident edges va and vb such that the angle $\angle avb$ attains the minimum among all pairs of such edges. We order a and b and call these two edges a base pair. Observe that the vectors vO (O is the origin), va , vb are not coplanar. Otherwise, v, a, b are in a great circle, v has only a and b as its neighbors, and v is of degree two.

Let n_1, n_2, n_3 be the vectors obtained by applying the Gram–Schmidt orthogonalization process to the ordered vectors vO, va, vb . Extend these vectors to a positively oriented orthonormal basis n_1, n_2, n_3, n_4 , and use this basis as the basis vectors of a Cartesian coordinate system with origin v . Sort the neighbors of v lexicographically by coordinates. Tag each coordinate with direction labels, $+1$ for outgoing edges, -1 for incoming edges, and ± 1 for bidirected edges. The concatenated sorted string of coordinates with tags yields a string corresponding to the base pair va, vb .

We can obtain such strings for all possible base pairs. We use the lexicographically smallest string to represent the congruence type of the vertex figure of the vertex v .

The congruence type of a directed edge uv is defined similarly by using a Cartesian coordinate system defined by some base pair. In this case, we choose the base pair in a way that it includes uv and another incident edge vb which forms the minimum angle to uv .

These representations of vertex figures and edge figures can be constructed in time $O(c^2 \log c)$ for degree c , and the length of the representations is $O(c)$. Since the maximum degree $c \leq 12$, all these bounds are constant.

14 1+3 Dimension Reduction

In 4- or higher-dimensional space, Alt et al. [2] were the first to use *dimension reduction*. The corresponding dimension reduction method for our algorithm is more a variant of one in Akutsu [1].

We pick an arbitrary point $a_0 \in A$. This point must be mapped to one of the points in B . For any $b \in B$, we try to rotate A so that a_0 lies on b . To look for rotations R that leave this point fixed, we project A and B on the hyperplane H perpendicular to $Oa_0 = Ob$. To each projected point, we attach the signed distance from H as a label. A rotation in H that preserves labels can be extended to a rotation R that fixes the point a_0 . For each $b \in B$, the problem is thus reduced to the one-lower dimensional congruence testing problem for labeled point sets.

The new algorithm uses this method for A and B in 4-space in the following situation. Suppose that we know pruned sets A_0, B_0 obtained from A and B such that $|A_0| \leq n_0$ and A_0 is mapped to B_0 by the pruning principle where n_0 is the constant defined in (9). Then by limiting the choice of a_0 in A_0 and the choice of b in B_0 , we can check the congruence of A and B in time $O(n_0 n \log n) = O(n \log n)$, by known methods for 3-space.

15 The Pruning and Condensing Principles

The pruning principle plays a very important role in our algorithm. The general scheme of *conventional pruning* is as follows.

1. Classify a set of points by the value of some function f .
2. Take only one class of points of the same value of a criterion, preferably the smallest class, into consideration and *temporarily* ignore other sets.

We need to make sure that this pruning is canonical, i. e., a criterion f is *equivariant under rotations*; this means that $f(R \cdot x) = R \cdot f(x)$ for all rotations R .

For example, given point sets A and B , if there is a point of non-zero distance to the origin in A , we can prune A and B by their distances to the origin as follows. First, classify points by their distances to the origin, count the number of points in each class, and temporarily focus only on the classes of the smallest cardinality A_s, B_s . After pruning by distances to the origin, we may assume that all the points are in the same distance to the origin, which means that they are on the same sphere. This is one of the main advantages of pruning. We can assume more structure on the point set after pruning.

Another advantage of pruning is that it reduces the number of points that we need to consider. If we can guarantee a size reduction by a factor $c < 1$, we say that we reduced the number of points successfully. Then, pruning can be repeated until it gets stuck, without affecting the time and space complexity. After pruning, we can restart the algorithm from the beginning with the pruned set without affecting the complexity of the algorithm, since $O(n \log n) + O(cn \log cn) + O(c^2 n \log c^2 n) + \dots = O(n \log n)$.

To assure that each pruning reduces the number of points by at least a half, we choose the set of the smallest cardinality. To ensure that the results for A and B are identical, we use some lexicographic rule for tie-breaking. We will only say that we “prune by criterion f ”, without explicitly mentioning a tie-breaking rule.

We generalize this principle to a more general method, called *condensing* while maintaining the advantages of pruning: We *condense* a finite set A to a nonempty set $A' = F(A)$ of smaller size, not necessarily a subset of A , by using an equivariant function F , that is, $F(R \cdot A) = R \cdot F(A)$ for any rotation matrix R . The condensing principle is not restricted to picking a subset according to some criterion, but it can define the function F in a more general way. For example, when we have a perfect matching of the points A that has been constructed by geometric criteria, we can replace A by the midpoints of the edges. Afterwards we only consider these midpoints instead of the original set A .

Any rotational invariants, such as distances and angles, can be used for designing a function F for condensing. For example, the new algorithm constructs canonical axes (Section 12) and edge figures (Section 13).

At first glance, condensing and pruning looks dangerous because it *throws away information*. This could introduce new symmetries, and it might happen that the condensed/pruned sets are congruent, whereas the original sets are not. Therefore, the condensed/pruned set should be kept only temporarily. The prime goal of iterative condensing steps is to eventually reduce point sets to small enough sets A'' and B'' so that we can afford *dimension reduction*.

Pruning and condensing are very powerful and versatile methods, because we can use them with any equivariant construction that one might think of as long as it is not too hard to compute. They allow us to concentrate on the cases where condensing makes no progress, and these cases are highly structured and symmetric. The difficulty is to pick the right pruning/condensing criteria, and to decide how to proceed when pruning and condensing gets stuck.

16 Overview of the New Algorithm

Finally, we introduce the new optimal algorithm in 4-space. We first describe the flow of the new algorithm and the relation between the modules.

As can be seen in Figure 4, the new algorithm consists of six main modules; iterative pruning, generating orbit cycles, marking and pruning great circles, the mirror case, 2+2 dimension reduction and 1+3 dimension reduction. The module called “1+3 dimension reduction” has already been presented in Section 14. We explain the other five modules in Sections 17–22.

Preliminaries. We will declare the problem to be trivially solvable if the closest-pair distance δ is large, i.e., $\delta > \delta_0 := 0.0005$. This implies that the input size $|A|$ is bounded by $n_0 < 3.016 \times 10^{11}$, and hence, by 1+3 dimension reduction, the problem can be reduced to at most n_0 instances of 3-dimensional congruence testing, taking $O(n \log n)$ time overall.

Remark. Whenever we apply some procedure, e.g., condensing, on A , we will apply the same procedure to the other set B in parallel but we will mostly describe those steps only for the set A . If A and B are congruent, B will undergo exactly the same sequence of steps as A . If a difference manifests itself at any point, we know that A and B are not congruent, and we can terminate.

General Flow. We first give a rough overview of our algorithm, omitting details and some special cases. See Fig. 4.

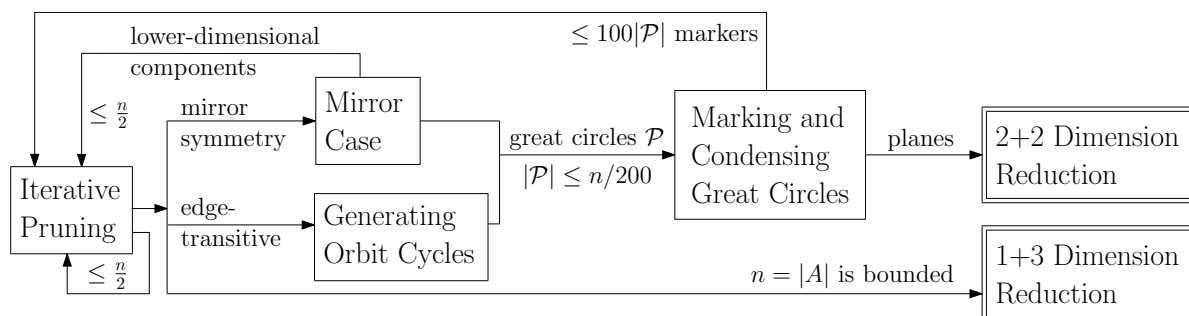


Figure 4: The general flow of the algorithm.

Our goal is to apply the condensing principle repetitively until we can apply either of the two dimension reduction principles to the original input point sets (not to the pruned sets): 1+3 dimension reduction (see Section 14) or 2+2 dimension reduction (see Section 22). The 1+3

dimension reduction makes a line and the orthogonal 3-space invariant, and the 2+2 dimension reduction makes two orthogonal 2-planes invariant.

Whenever the number of pruned points is smaller than some chosen threshold n_0 , we can afford 1+3 dimension reduction as mentioned in the preliminaries. If we find a set \mathcal{P} of equivariant great circles from A and B , such that $1 \leq |\mathcal{P}| \leq C_1$, where $C_1 = 829$ is defined in Section 11, we can trigger 2+2 dimension reduction.

By applying the 2+2 dimension reduction technique described in Section 22, the problem boils down to congruence testing of “labeled” points on a two-dimensional torus under translations. This problem can be solved in $O(n \log n)$ by reducing a point set but still preserving symmetries and by using the Voronoi regions as the neighborhood structure.

The congruence testing algorithm begins with pruning by distance from the origin. We may thus assume without loss of generality that the resulting set A' lies on the unit 3-sphere $\mathbb{S}^3 \subset \mathbb{R}^4$.

In iterative pruning (Section 17), we first check if the minimum distance δ between points of A' is bigger than the threshold δ_0 . If yes, we conclude that $|A'| \leq n_0$ and trigger 1+3 dimension reduction. Otherwise, we construct the closest-pair graph G on A' . Then we prune the edges of G by congruence type of its edge figures. An edge figure consists of two adjacent vertices and all their neighbors. We apply pruning with other criteria to the set A' until the resulting set A_0 cannot be further reduced. Then, all edge neighborhoods in the graph G are congruent. This allows us to find either *orbit cycles* in G , or a subset of A_0 with mirror symmetry. An orbit cycle is a cyclic path $a_1 a_2 \dots a_\ell a_1$ with a rotation R such that $R a_i = a_{(i \bmod \ell) + 1}$; in other words, it is the orbit of the point a_1 under a rotation R .

Mirror symmetry is the symmetry that, for each edge of G , swaps the two endpoints and maps the whole point set onto itself. It implies that the point set A_0 must be related to one of the regular four-dimensional polyhedra, or A_0 (and G) is the Cartesian product of two regular polygons in orthogonal planes. This follows from the classical classification of discrete reflection groups. The former case can be excluded, since $\delta < \delta_0$, and in the latter case, we can proceed to Section 21 for marking and condensing great circles. See Section 19.

Let us look at the case when we have found orbit cycles (Section 18) or a set of pairs of regular polygons in orthogonal planes (Section 19). Geometrically, an orbit cycle lies on a helix around a great circle C and a pair of orthogonal planes intersect a sphere as a pair C, C' of great circles. Thus, we get a collection \mathcal{P} of great circles on \mathbb{S}^3 . We treat these great circles as objects in their own right, and we construct the closest-pair graph on \mathcal{P} . For this, we use the Plücker embedding of the corresponding planes into $\mathbb{S}^5/\mathbb{Z}_2 = \mathbb{RP}^5$, mapping each plane to a pair of antipodal points on the 5-sphere \mathbb{S}^5 . Then, we construct a closest-pair graph of planes with respect to the Plücker distance, defined as a normalized Euclidean distance in the Plücker embedding. Refer to Section 7 for the Plücker embedding and Plücker distance.

For each closest pair (C, D) of great circles in \mathcal{P} , if they are not isoclinic (i.e., the projection of C to the plane containing D is an ellipse but not a circle) the major axis of this ellipse marks two points of C as described in Paragraph “Equivalent Definition by Orthogonal Projections” in Section 6. The set of all markers replaces the set A_0 . This completes a successful condensing step, and we restart and continue pruning as before.

Otherwise, all projected “ellipses” turn out to be circles. In this case, we can find a subfamily of great circles in a special position: they must be part of a *Hopf bundle* of circles. The circles of a Hopf bundle can be mapped one-to-one to points on the 2-sphere \mathbb{S}^2 . See Section 8. We can thus use the condensing procedures for a 2-sphere in three dimensions as in Section 20. This yields a small set \mathcal{P} of at most 12 great circles. Then, we can apply the 2+2 dimension reduction technique.

The details for marking or condensing great circles are actually more complicated, since we might have a phase in which \mathcal{P} is successively pruned, see Section 21.

This concludes the summary of the algorithm. The steps of the algorithm involve several different operations: We need closest-pair graphs in 4 and 6 dimensions. The closest-pair graph

of k points can be calculated in $O(k \log k)$ time in any fixed dimension, by a classical divide-and-conquer approach [6]. We also compare a pair of edge figures in \mathbb{S}^3 but this takes only time $O(1)$ because the maximum degree is at most 12. We also need Voronoi diagrams in two dimensions and convex hulls in three dimensions. Finally, we need to sort lists of numbers lexicographically. In summary, we will be able to reduce the size of the current point set A' by a constant factor less than $\frac{1}{2}$, in $O(|A'| \log |A'|)$ time, until dimension reduction is possible.

Hence, we obtain the following theorem.

Theorem 16. *We can decide if two n -point sets A and B in 4-space are congruent in $O(n \log n)$ time and $O(n)$ space.*

17 Iterative Condensing Based on the Closest-Pair Graph: Algorithm C

After pruning by the distance to the origin, we have a set $A \subset \mathbb{R}^4$ of n points with equal distance from the origin. This set may already be the result of some previous condensing steps. Without loss of generality, we may assume to be given a set A on the unit sphere \mathbb{S}^3 . We construct the closest-pair graph G for a point set A on the 3-sphere and try to condense it.

In this section, we perform an algorithm, which is a sequence of pruning and condensing to obtain a new equivariant point set A' that is in one of the following cases:

- (i) the closest distance δ of A' is greater than δ_0 , or
- (ii) the closest-pair graph of A' has mirror symmetry, or
- (iii) the closest-pair graph of A' is in the edge-transitive case.

If (i) happens, we apply 1+3 dimension reduction. If (ii) or (iii) happens, we proceed to Algorithm M in Section 19 or Algorithm O in Section 18 respectively. Those sections explain how to reduce the cases to 2+2 dimension reduction.

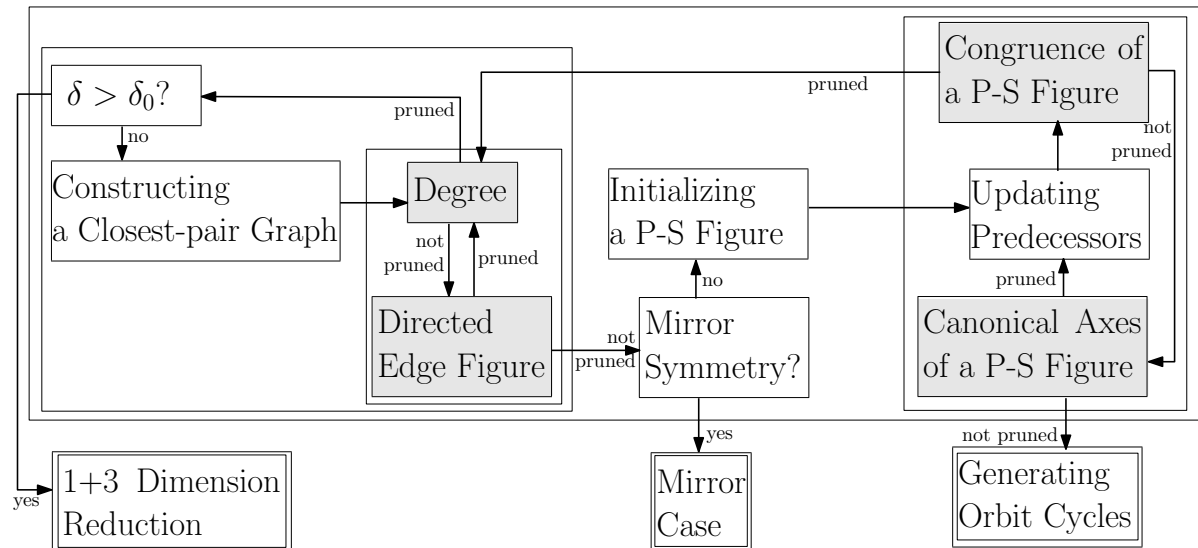


Figure 5: Iterative pruning. The shaded boxes represent pruning/condensing steps with pruning criteria or condensing methods.

We first explain the notion of a *predecessor-successor figure* in G . We have a directed edge uv of G together with a set of *predecessor edges* $p(uv)$ incident to u and a set of *successor edges* $s(uv)$ incident to v , as in Fig. 6a. All edges have the same length, their endpoints lie on the

3-sphere \mathbb{S}^3 , and all predecessor and successor edges form the same angle α with uv . Then the endpoints of these edges lie on two circles. If we reflect the predecessor circle at the bisecting hyperplane of uv , it comes to lie on the successor circle. This results in one circle with a succinct representation of the geometric situation. See Fig. 6a and Fig. 6b. We refer to the endpoints of the predecessor and successor edges as *predecessors* and *successors*.




Figure 6: (a) Predecessors and successors of uv at angle α ; t' is the reflected copy of t on the successor circle. (b) The corresponding predecessor-successor figure, and the torsion angle $\tau(tuvw)$. Predecessors are drawn white and successors black. (c) An edge-transitive predecessor-successor figure with torsion angle τ_0 . (d) Canonical axes. (e) Mirror symmetry.

The following algorithm guarantees that the outcome A' is in one of the promised cases.

Algorithm C (*Prune the closest-pair graph*). We are given a set $A \subset \mathbb{R}^4$ of n points, equidistant from the origin. This set may already be the result of some previous condensing steps. Without loss of generality, we may assume that A lies on the unit sphere \mathbb{S}^3 .

- C1.** [Well-separated points?] Compute the closest distance δ between points of A . If $\delta > \delta_0 = 0.0005$, apply 1+3 dimension reduction. ($|A|$ is bounded by a constant.)
- C2.** [Construct the closest-pair graph.] Construct the closest-pair graph G , and initialize the directed graph D with two opposite arcs for every edge of G . (This takes $O(n \log n)$ time, and the degrees in G are bounded by $K_3 = 12$.)
- C3.** [Prune by degree.] If the indegrees and outdegrees in D are not all the same, prune the vertices by degree, and return to C1, with the smallest class A' taking the role of A . (Otherwise, we now enter a loop in which we try to prune arcs from D .)
- C4.** [Prune by directed edge figure.] The *directed edge figure* of an arc $uv \in D$ consists of uv together with all arcs of D out of v and all arcs of D that lead into u . If the directed edge figures are not all congruent, *prune the arcs of D* by congruence type of directed edge figures, and return to C3. (Here we apply the pruning principle not to points but to arcs and we replace the edge set D by a smaller subset D' . Since the degrees are bounded, we can compare two directed edge figures in constant time.)
- C5.** [Mirror symmetry?] (Now all arcs have the same directed edge figure.) If the directed edge figure of uv is symmetric with respect to the bisecting hyperplane of uv , proceed to Algorithm R (Section 19).
- C6.** [Choose an angle α with no mirror symmetry.] Pick an angle α for which the predecessors and successors are not completely symmetric, that is, the predecessor-successor figure does not look like Fig. 6e.
- C7.** [Initialize Successors.] For every arc $uv \in D$, set $s(uv) := \{vw : vw \in D, \angle uvw = \alpha\}$. (The size of $s(uv)$ is bounded by $K_2 = 5$. We will now enter an inner loop in which we try to prune arcs from the sets $s(uv)$.)

- C8.** [Update predecessors.] Define the predecessor edges by $p(uv) := \{tu : uv \in s(tu)\}$. Build the predecessor-successor figure for each arc, as explained in the text above.
- C9.** [Prune by predecessor-successor figures.] If there are arcs whose predecessor-successor figures are not congruent, prune the arcs of D accordingly, and return to C3.
- C10.** [Check regularity.] (Now all arcs have the same predecessor-successor figure. Each figure must contain the same number k of predecessors and successors, since the total number of predecessors in all figures must balance the total number of successors.) If the predecessor-successor figure consists of two regular k -gons, proceed to Algorithm O for generating orbit cycles, see Section 18. (We call this the *edge-transitive* case. See Figure 6c.)
- C11.** [Prune successors by canonical axes.] Prune $s(uv)$ to a proper nonempty subset by computing *canonical axes* as explained in Section 12, and return to C8.

An example of canonical axes is depicted in Fig. 6d. If canonical axes consist of p axes, we know that $p < k$ because the maximally symmetric case of two regular k -gons (the *edge-transitive* case shown in Fig. 6c) has been excluded in Step C10. The loop from C8–C11 maintains the following loop invariant on entry to Step C10:

$$\textit{There is a position of a successor that is not occupied by a predecessor.} \quad (11)$$

For the reduction in Step C11, we rotate the canonical axes counterclockwise until they hit the first position of type (11). The successors that are intersected by the canonical axes form a nonempty proper subset $s' \subseteq s(uv)$. We thus replace $s(uv)$ for each edge uv by s' and return to Step C8. By construction, we have made sure that (11) still holds. After pruning all successor sets, the predecessor sets are reduced accordingly in Step C8, but this cannot invalidate (11). (The invariant (11) holds on first entry to the loop because of Step C6.)

The algorithm has three nested loops (indicated by indentation) and works its way up to higher and higher orders of regularity. After C3, all vertices have the same degree. After C4, we know that all *pairs* of adjacent vertices look the same. If we exit to Algorithm O in Step C10, we will see that certain chains of *four* points can be found everywhere.

There is the *global loop* that leads back to C1 after each successful pruning of vertices by degree. Since the size of A is reduced to less than a half, we need not count the iterations of this loop. In addition, there is an outer loop that resumes working at C3 after pruning the edges of D , and an inner loop that starts at C8 and is repeated whenever the successor set $s(uv)$ is pruned. In these loops, we maintain that $D \neq \emptyset$ and $s(uv) \neq \emptyset$. In Step C3, if we have removed at least one edge from D , we will either be able to prune by degree, or the degree of all vertices has gone down by at least one. Since the degree is initially bounded by 12, Step C3 can be visited at most 12 times before exiting to C1. Similarly, Step C11 removes at least one element of $s(uv)$, so this loop is repeated at most 5 times before there is an exit to the outer loop in step C9. The most time-consuming step is the construction of the closest-pair graph in Step C2. All other operations take $O(n)$ time, not counting the exits to Algorithms M and O. Thus, the overall time is $O(|A| \log |A|)$.

18 Generating Orbit-Cycles: Algorithm O

We now describe how, in the edge-transitive case, the algorithm O produces a set \mathcal{P} of at most $|A|/200$ great circles. All predecessor-successor figures look like Fig. 6c. The *torsion angle* $\tau(tuvw)$ between a predecessor edge $tu \in p(uv)$ and a successor edge $vw \in s(uv)$ is the oriented angle $\angle(t'ow)$ in the predecessor-successor figure of uv . We define τ_0 as the smallest counterclockwise torsion angle that appears in the predecessor-successor figure. Let $t_0u_0v_0w_0$ be a fixed quadruple with this torsion angle.

- Lemma 17.**
1. For every $a_2a_3 \in s(a_1a_2)$, there is a (unique) edge $a_3a_4 \in s(a_2a_3)$ such that $a_1a_2a_3a_4$ is congruent to $t_0u_0v_0w_0$.
 2. Moreover, there is a unique rotation R_0 that maps $t_0u_0v_0w_0$ to $a_1a_2a_3a_4$.
 3. For every triple $a_1a_2a_3$ with $a_2a_3 \in s(a_1a_2)$, there is a unique cyclic sequence $a_1a_2 \dots a_\ell$ such that $a_i a_{i+1} a_{i+2} a_{i+3}$ is congruent to $t_0u_0v_0w_0$ for all i . (Indices are taken modulo ℓ .)
 4. Moreover, there is a unique rotation matrix R such that $a_{i+1} = Ra_i$. In other words, $a_1a_2 \dots a_\ell$ is the orbit of a_1 under the rotation R .
 5. The points $a_1a_2 \dots a_\ell$ do not lie on a circle.

Proof. We first establish two facts about $t_0u_0v_0w_0$.

$$\text{The four points } t_0, u_0, v_0, w_0 \text{ do not lie on a circle.} \tag{12}$$

$$\text{The three points } t_0, u_0, v_0 \text{ do not lie on a great circle.} \tag{13}$$

If t_0, u_0, v_0, w_0 would lie on a circle (not necessarily through the origin), the predecessor-successor figure of u_0v_0 would have a mirror-symmetric predecessor $t_0u_0 \in p(u_0v_0)$ and successor $v_0w_0 \in s(u_0v_0)$, in contradiction to the choice of α in Step C6 and to the invariant (11). If the points t_0, u_0, v_0 lie on a great circle C , then w_0 must also lie on C , since $\angle t_0u_0v_0 = \angle u_0v_0w_0 = \alpha$, but this would contradict (12).

Now, let a_1, a_2, a_3 be any three points with $a_2a_3 \in s(a_1a_2)$. By the definition of predecessors, $a_1a_2 \in p(a_2a_3)$. We can thus fit $t_0u_0v_0$ to $a_1a_2a_3$ in the predecessor-successor figure of a_2a_3 . Since the points t_0, u_0, v_0 are not on a great circle, they span a three-dimensional subspace, and the rotation R_0 that maps $t_0u_0v_0$ to $a_1a_2a_3$ is uniquely determined. Since all predecessor-successor figures are congruent, this means that the edge v_0w_0 is mapped to some successor edge $a_3a_4 \in s(a_2a_3)$. This establishes Properties 1, 2, and 5.

This process can be continued: Since $a_2a_3 \in p(a_3a_4)$, there is a unique edge $a_4a_5 \in s(a_3a_4)$ such that $a_2a_3a_4a_5$ is congruent to $t_0u_0v_0w_0$, and so on.

By Property 2, there are two unique rotations from $t_0u_0v_0w_0$ to $a_1a_2a_3a_4$ and to $a_2a_3a_4a_5$, and thus there is a unique rotation R' with $R[a_1a_2a_3a_4] = [a_2a_3a_4a_5]$. We have seen that the points $a_1a_2a_3$ are not on a great circle, being congruent to $t_0u_0v_0$ and thus the rotation R is already uniquely specified by the conditions $R[a_1a_2a_3] = [a_2a_3a_4]$. We have therefore established that $a_{i+1} = Ra_i$ for $i = 1, 2, 3$ implies $a_{i+1} = Ra_i$ for $i = 4$. This can be continued by induction, implying that $a_1a_2 \dots$ is the orbit of the point a_1 under the rotation R .

Since the number of points in A is finite, this orbit must return to a_1 and form a directed cyclic path $a_1a_2 \dots a_\ell$ in the closest-pair graph D . This establishes Properties 3 and 4. \square

We call the cyclic paths that are constructed in Lemma 17 *orbit cycles*. The following lemma gives a bound on the number of orbit cycles in terms of $|A|$ when the closest-pair distance is small enough.

Lemma 18. *The number of orbit cycles is at most $|A|/200$ provided that the closest distance δ is smaller than $\delta_0 = 0.0005$.*

Proof. We have $a_{i+1} = R_{\varphi,\psi}a_i$ for all i with an appropriate basis $x_1y_1x_2y_2$ where $R_{\varphi,\psi}$ is defined as (1). We cannot have $\varphi = 0$ or $\psi = 0$, because otherwise the orbit would form a regular polygon $a_1a_2a_3a_4 \dots a_\ell$ in a plane, contradicting Lemma 17.5. Thus, we know that $\varphi, \psi \neq 0$. To get a closed loop, we must have $|\varphi|, |\psi| \geq 2\pi/\ell$. If the projection of a_i to the x_1y_1 -plane has norm r_1 and the projection to the x_2y_2 -plane has norm r_2 with $r_1^2 + r_2^2 = 1$, then the squared distance is

$$\delta^2 = \|a_{i+1} - a_i\|^2 = (2r_1 \sin \frac{|\varphi|}{2})^2 + (2r_2 \sin \frac{|\psi|}{2})^2 \geq 4 \sin^2(\pi/\ell).$$

As a result,

$$\delta \geq 2 \sin \frac{\pi}{\ell} \tag{14}$$

We can get the same bound via Fenchel’s theorem about the length of a closed curve that is not contained in a hemisphere [16]. Thus, if $\delta \leq 2 \sin(\pi/12000) \approx 0.000523$, it is guaranteed that $\ell \geq 12000$, that is, every orbit cycle contains at least 12000 points. On the other hand, each point $u \in A$ belongs to a bounded number of orbit cycles: it has at most $K_3 = 12$ incoming arcs tu , and each arc tu has at most $K_2 = 5$ successor edges $uv \in s(tu)$. The triple tuv specifies a unique orbit cycle, and thus there are at most $12 \times 5 = 60$ orbit cycles through u . The total number of orbit cycles is therefore at most $|A| \cdot 60/12000 = |A|/200$. \square

For each orbit cycle, we can find an appropriate rotation matrix $R_{\varphi,\psi}$. If $\varphi = \pm\psi$, then the orbit of any point under this rotation lies on a great circle, contradicting Lemma 17.5. Thus, we get a unique invariant plane that rotates by the smaller angle in itself. We replace each orbit cycle by the great circle in this invariant plane, yielding the desired set \mathcal{P} of great circles with $|\mathcal{P}| \leq |A|/200$.

Section 19 we need a bound on the number of points in a regular polygon inscribed in a great circle.

Corollary 19. *If the closest distance $\delta < \delta_0$ and there is a regular k -gon inscribed in a great circle of a sphere, then $k > 12000$.*

Proof. The same computation as in Lemma 18 can be used for a great circle. From (14) and $2 \sin(\pi/12000) > \delta_0$, we find that $k > 12000$. \square

19 The Mirror Case: Algorithm R

Algorithm R (*Treat mirror-symmetric closest-pair graphs*). We are given a nonempty directed subgraph D of the closest-pair graph on a point set $A \subset \mathbb{S}^3$ with closest-pair distance $\delta \leq \delta_0$. All directed edges in D have equal edge-figures and exhibit perfect mirror-symmetry, as established in Steps C4 and C5 of Algorithm C. Algorithm R will produce, in an equivariant way, either

- a) a set A' of at most $|A|/2$ points, or
- b) a set \mathcal{P} of at most $|A|/200$ great circles.

- R1.** [Make D undirected.] Construct the undirected version of D and call it G .
- R2.** [Check for eccentric centers of mass.] Compute the center of mass of each connected component of G . If these centers are not in the origin, return the set A' of these centers.
- R3.** [Two-dimensional components?] If each component is a regular polygon with center at the origin, return the set \mathcal{P} of circumcircles of each polygon.
- R4.** [Three-dimensional components?] If each component spans a 3-dimensional hyperplane H , replace the component by two antipodal points perpendicular to H . Return the set A' of these points.
- R5.** [Toroidal grid.] As it was shown in Lemma 14, now each component of D is the product $P \times Q$ of a regular p -gon P and a regular q -gon Q in orthogonal planes, with $p, q \geq 3$. Pick a vertex u from each component. There are four incident edges, and the plane spanned by two incident edges is orthogonal to the plane spanned by the other two edges. Represent the component of G by these two orthogonal planes shifted to the origin. Return the set \mathcal{P} of great circles in these planes (two per component).

The algorithm takes linear time. We still need to show that these cases are exhaustive. After we make D undirected in Step R1, for every edge uv of G , the bisecting hyperplane of uv acts as a mirror that exchanges u and v together with their neighborhoods. Since the reflected mirror hyperplanes act again as mirrors, it follows that every component of the graph is the orbit of some point u_0 under the group generated by reflections perpendicular to the edges incident to u_0 . Thus, the incident edges of a single point determine the geometry of the whole component. The graph may have several components, all of which are congruent.

There are infinitely-many groups of the form $D_2^p \times D_2^q$, where D_2^p is the dihedral group of order $2p$, the symmetry group of the regular p -gon. These groups are treated in Step R5. Except this infinite family, there are only eight other groups, which are excluded by Lemma 14 (see Section 10), because the constant $\delta_1 = 0.07$ from Lemma 14 is (much) larger than δ_0 .

We go through the steps one by one to check if the algorithm achieves the claimed results. Step R2 treats the case when a (lower-dimensional) component does not go through the origin. This includes the cases when G is a matching or a union of “small” regular polygons. Every component contains at least 2 points, and thus $|A'| \leq |A|/2$. Steps R3 and R4 treat the two- and three-dimensional components that are centered at the origin. (The one-dimensional case of a matching cannot be centered at the origin, because then we would have $\delta = 2$.) Suppose we have a two-dimensional component and this component is a regular k -gon inscribed in a great circle (Step R3). From $\delta < \delta_0$ and by Corollary 19, we obtain that $k > 12000$. Thus $|C| \leq |A|/12000 \leq |A|/200$. A three-dimensional component (Step R4) contains at least four points and is reduced to two antipodal points. Again we have $|A'| \leq |A|/2$.

Steps R2–R4 have treated all cases of Coxeter groups which are not full-dimensional, and Lemma 14 excludes all full-dimensional groups which are not of the form $D_2^p \times D_2^q$. Thus, in Step R5, each component of D is the product $P \times Q$ of a regular p -gon P and a regular q -gon Q in orthogonal planes, with $p, q \geq 3$.

Note that a regular p -gon can be generated as the orbit of a point u_0 in D_2^p or in D_2^{2p} , depending on whether we put u_0 in the interior of the fundamental region or on a mirror. But this makes no difference for the resulting point set.

Such a component forms a toroidal $p \times q$ grid. Each vertex has four neighbors. The two polygons P and Q have equal side lengths δ , because otherwise the four neighbors would not be part of the closest-pair graph G . The two incident edges of a vertex u that come from P are orthogonal to the two edges that come from Q , so we can distinguish the two edge classes. (The case when all four edges are perpendicular is the 4-cube with $p = q = 4$ and with reflection group $D_2^4 \times D_2^4$. This case has 16 vertices $(\pm 1/2, \pm 1/2, \pm 1/2, \pm 1/2)$ and minimum distance $\delta = 1$, and it is therefore excluded.) Moreover, the grid cells are geometric squares. Through these squares, the classification of edges into the edges from P and the edges from Q can be transferred consistently to a whole connected component of G . All copies of P in the grid lie in parallel planes, and so do the copies of Q . Accordingly, Step R5 represents each connected component by two orthogonal planes through the origin. We need to show that the component is large. As $P \times Q$ lies on the unit 3-sphere, the circumradii r_P and r_Q of the two polygons satisfy $r_P^2 + r_Q^2 = 1$. Thus, the larger circumradius, let us say r_P , is at least $1/\sqrt{2}$. Since the closest-pair distance is $\delta = 2r_P \sin \frac{\pi}{p} \leq \delta_0 = 0.0005$, we get $\sqrt{2} \sin \frac{\pi}{p} \leq \delta_0$, which implies $p \geq 8886$. Each component contains $pq \geq p$ points and is reduced to two circles. Thus, the algorithm achieves the claimed reduction.

20 Finding Representative Points from a 2-Sphere: Algorithm K

In the following Section 21, we need an auxiliary algorithm to condense a set of great circles that belong to a common Hopf bundle. These circles can be mapped to points on the 2-sphere, and thus, we adapt Atkinson’s algorithm [5] (see Section 4.2) to condense a finite point set in a 2-sphere.

Lemma 20 (Algorithm K). *There is a procedure that reduces a set F of points on the 2-sphere \mathbb{S}^2 to a nonempty set F' of at most $\min\{12, |F|\}$ representative points on \mathbb{S}^2 , in $O(|F| \log |F|)$ time. These points are either the vertices of a regular tetrahedron, a regular octahedron, a regular icosahedron, a single point, or a pair of antipodal points. This mapping from sets F to sets F' is equivariant under rotations and reflections.*

Proof. We repeatedly prune the set F . We start by computing the convex hull of F . If it does not contain the origin, we can immediately output the vector pointing to the center of gravity of the points as a representative. If the hull is one-dimensional or two-dimensional, we get two antipodal points as representatives.

We can thus assume that the hull is a three-dimensional polytope. If the vertex degrees are not all the same, we prune the point set and restart. We can thus assume that the graph of the polytope is regular, and by Euler's formula, the degree d can be 3, 4, or 5. Euler's formula also yields the number f of faces in terms of the number n of vertices: $f = (d-2)/2 \cdot n + 2 \leq \frac{3}{2}n + 2$. We now try to prune by face degrees. If there are different face degrees, we replace F by the centroids of the smallest class of faces, and restart. We have condensed $|F|$ at least by a factor of $3/4$ (the additive term $+2$ is negligible as long as $|F|$ is large). The remaining case is when all face degrees and all vertex degrees are equal. Then the polytope must have the combinatorics of one of the five regular polytopes: the tetrahedron, the octahedron, the icosahedron, the cube, or the dodecahedron. For the cube and the dodecahedron, we take the centroids of the faces and thereby reduce the number of points.

We are left with the case of the tetrahedron, the octahedron, and the icosahedron. If the edges don't have the same lengths, we condense and replace F by the midpoints of the smallest edge class. This will lead to a reduction in the number of points, except for an octahedron and icosahedron with two edge lengths, each occurring at least 6 times or 12 times, respectively. We claim that the triangular faces cannot be all congruent in this case, and we can prune the 8, resp. 20, faces by congruence type, leading to a reduction to at most 4 resp. 10 triangle centroids. If the triangles were all congruent, their sides would all have to be long-long-short or long-short-short. This would mean, for the octahedron, that the 12 edge lengths are divided in the proportion $4 : 8$, and in that case we could have condensed by the edges. The same argument works for the icosahedron, because the 30 edge lengths are divided in the proportion $10 : 20$.

Thus, the only cases where pruning cannot proceed are a tetrahedron, an octahedron, or an icosahedron, with regular triangles as faces. These must be the regular polytopes.

The most expensive part in each pruning step is the computation of the convex hull, which takes $O(n' \log n')$ time, where n' is the size of the current point set. Each step reduces n' by a constant factor of at most $11/12$. Thus the overall running time is $O(|F| \log |F|)$. \square

21 Marking and Condensing of Great Circles: Algorithm M

We have extracted a set \mathcal{P} of at most $|A|/200$ great circles from the point set A , either from the mirror case (Algorithm R in Section 19) or from orbit cycles (Algorithm O in Section 18). Algorithm M obtains a small set A' of *marker points* on each circle so that we can resume Algorithm C, or it exits to Algorithm T for 2+2 dimension reduction (Section 22).

For this purpose, we look at the angles α, β of pairs of circles $C, D \in \mathcal{P}$. If C and D are not Clifford parallel, we can *mark* a pair of points on C and on D , as follows: Let C' be an orthogonal projection of C to the plane through D . Then C' is an ellipse (not a circle). A pair of points on D can be *marked* by the intersection of D and the major axis of C' . Similarly, two points on C are marked.

Doing this for all pairs of circles would lead to a quadratic blowup. Therefore we construct the closest-pair graph on the *set of circles*. We use Plücker coordinates to map great circles of \mathbb{S}^3 to points on \mathbb{S}^5 . We can then compute the closest-pair graph in six dimensions, and every circle has at most K_5 closest neighbors.

The marking approach fails for Clifford-parallel circles. Clifford-parallel pairs manifest a very rigid structure on a sphere. Section 8 describes properties of such pairs. For convenience, we restate these results in the following proposition. They are formulated for right pairs of circles, but they hold equally with left and right exchanged.

Proposition 21. *1. The relation of being a right pair is transitive (as well as reflexive and symmetric) (Corollary 11). An equivalence class is called a right Hopf bundle.*

2. For each right Hopf bundle, there is a right Hopf map h that maps the circles of this bundle to points on \mathbb{S}^2 (Lemma 8).

3. By this map, two Clifford-parallel circles with angle α, α (and with Euclidean distance $\sqrt{2} \sin \alpha$ on the Plücker sphere \mathbb{S}^5) are mapped to points at angular distance 2α on \mathbb{S}^2 (Lemma 12).

4. A circle can have at most $K_2 = 5$ closest neighbors on the Plücker sphere \mathbb{S}^5 that form right pairs.

Proof of Property 4. Property 4 is a direct consequence of Properties 2 and 3. □

Algorithm M (*Mark and condense great circles*). Given a set \mathcal{P} of great circles on \mathbb{S}^3 , this algorithm will produce, in an equivariant way, either a nonempty set A' of at most $100|\mathcal{P}|$ points on \mathbb{S}^3 , or a nonempty set \mathcal{P}' of at most 829 great circles.

The algorithm updates the set \mathcal{P} and maintains an equivalence relation \sim on \mathcal{P} such that all circles in the same equivalence class belong to a common (left or right) Hopf bundle. The common chirality of all bundles is indicated by a variable $chirality \in \{None, Left, Right\}$, where *None* is chosen when the equivalence relation is trivial and all classes are singletons. The size of the equivalence classes is bounded by 12.

M1. [Initialize.] Let every circle $C \in \mathcal{P}$ form a singleton equivalence class.

M2. [Prune by size.] If equivalence classes are of different sizes, choose the size that occurs least frequently. Throw away all equivalence classes that are not of this size, together with the circles they contain.

M3. [Few circles?] If $|\mathcal{P}| \leq 829$, return the set \mathcal{P} .

M4. [Singletons?] If all equivalence classes are singletons, set $chirality := None$.

M5. [Construct closest pairs.] Represent each circle $C \in \mathcal{P}$ by two antipodal points on \mathbb{S}^5 , using Plücker coordinates. Construct the closest-pair graph H on \mathcal{P} with respect to the distances on \mathbb{S}^5 .

M6. [Classify edges.] Partition the edges of H into $E_L \cup E_R \cup E_N$, representing pairs of circles that are left pairs, right pairs, and not Clifford parallel.

M7. [Use non-Clifford-parallel edges.] If $E_N \neq \emptyset$, let $\mathcal{N} := E_N$ and go to Step M12.

M8. [Find non-Clifford-parallel pairs from equivalent circles.] If $E_L \neq \emptyset$ and $chirality = Right$, set

$$\mathcal{N}_C := \{\{C', D\} \mid \{C, D\} \in E_L, C' \text{ is closest to } C \text{ among the circles } C' \parallel_+ C, C' \neq C\}$$

for all $C \in \mathcal{P}$. Set $\mathcal{N} := \bigcup_{C \in \mathcal{P}} \mathcal{N}_C$ and go to Step M12.

M9. (Same as M8, with left and right exchanged.)

- M10.** [Merge classes.] If $E_L \neq \emptyset$ and $chirality \in \{Left, None\}$, merge equivalence classes that contain circles that are connected in E_L . Use Algorithm K to condense each resulting class F to a set F' of at most 12 *representative* circles. Set \mathcal{P} to the set of all representative circles, and put them into the same equivalence class if and only if they come from the same set F' . Set $chirality := Left$, and return to M2.
- M11.** (Same as M10, with left and right exchanged. Since all possibilities are exhausted, the only remaining possibility in this step is to merge classes and return to M2.)
- M12.** [Mark points on circles.] (Each pair in \mathcal{N} is now a non-Clifford-parallel pair of circles.) For each pair of circles $\{C, D\} \in \mathcal{N}$, mark the two points on C that are closest to D , and likewise, mark two points on D . Return the set A' of all marked points.

The crucial properties for the correctness and the running time are formulated in a lemma:

- Lemma 22.** 1. *After the algorithm returns to Step M2 from step M10 or M11, the number of equivalence classes is reduced at least by half.*
2. *Algorithm M terminates in $O(|\mathcal{P}| \log |\mathcal{P}|)$ time.*
3. *In Step M12, \mathcal{N} is a nonempty set of at most $25|\mathcal{P}|$ non-Clifford-parallel pairs.*

Proof. 1. The only possibility for the algorithm to stall is that the edges of H are *within* classes, and no merging takes place in Step M10 or M11. Each class must be one of the five configurations listed in Lemma 20, and the smallest possible angular distance $2\alpha_{\min} \approx 1.10715$ occurs for two adjacent vertices of the icosahedron, and a volume packing argument on \mathbb{S}^5 yields that there can be at most 829 circles with this distance. See Section 11. Then the algorithm exits in Step M3.

2. The most expensive step is the construction of the closest-pair graph in Step M5, which takes $O(|\mathcal{P}| \log |\mathcal{P}|)$ time. The algorithm contains one loop, when returning from M10 or M11 to M2. Since the number of equivalence classes is geometrically decreasing and each class contains at most 12 circles, the overall time is also bounded by $O(|\mathcal{P}| \log |\mathcal{P}|)$.

3. When \mathcal{N} is constructed in Step M7, there can be at most $22|\mathcal{P}|$ such pairs, because the degree in H is bounded by $K_5 \leq 44$. In Step M8, the constructed set \mathcal{N} is nonempty: Every pair $\{C, D\} \in E_L$ produces at least one element of \mathcal{N} , since all equivalence classes contain at least two elements, by M2 and M4. When a pair $\{C', D\}$ in \mathcal{N} is formed, we have a left pair $\{C, D\}$ and a right pair $\{C, C'\}$. If $\{C', D\}$ were also Clifford parallel, we would get a contradiction to transitivity (Proposition 21.1). Each circle $C \in \mathcal{P}$ gives rise to at most $5 \cdot 5 = 25$ pairs, since a circle can have at most 5 Clifford-parallel neighbors in H by Proposition 21.4. \square

We conclude that the algorithm produces at most $100|\mathcal{P}|$ points, four points for each pair in \mathcal{N} .

22 2+2 Dimension Reduction: Algorithm T

If we arrive at Algorithm T, we are in the very last step so we should restore the initial input point sets A and B . This 2+2 dimension reduction is applied when we have already identified at most $C_1 = 829$ pairs of planes P and Q from Algorithm M in Section 21 such that any candidates R of congruence mapping from A to B has to map P to Q .

We begin by choosing a coordinate system (x_1, y_1, x_2, y_2) for A so that P becomes the x_1y_1 -plane, and similarly for B and Q . We then look for rotations R that leave the x_1y_1 -plane invariant. Such rotations have the form $R = \begin{pmatrix} R_1 & 0 \\ 0 & R_2 \end{pmatrix}$ with two orthogonal 2×2 matrices R_1 and R_2 . Since $\det R = \det R_1 \cdot \det R_2 = 1$, we try two cases: (a) R_1 and R_2 are planar rotations; (b) R_1 and R_2 are planar reflections. We can reduce (b) to (a) by applying the

rotation $\begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}$ to A . Thus, it suffices to describe (a), where R has the form $R_{\varphi,\psi}$ in (1), i. e., a combination of a rotation by φ in the x_1y_1 -plane and an independent rotation by ψ in the x_2y_2 -plane. We use polar coordinates in these two planes by setting

$$\begin{pmatrix} x_1 \\ y_1 \\ x_2 \\ y_2 \end{pmatrix} = \begin{pmatrix} r_1 \cos \alpha_1 \\ r_1 \sin \alpha_1 \\ r_2 \cos \alpha_2 \\ r_2 \sin \alpha_2 \end{pmatrix} \quad \text{with } r_1 = \sqrt{x_1^2 + y_1^2} \text{ and } r_2 = \sqrt{x_2^2 + y_2^2}. \quad (15)$$

The rotation $R_{\varphi,\psi}$ changes (α_1, α_2) by adding (φ, ψ) modulo 2π , and leaves r_1 and r_2 unchanged. In other words, $R_{\varphi,\psi}$ acts as a translation on the torus $\mathbb{T}^2 = [0, 2\pi) \times [0, 2\pi)$. We attach the distances (r_1, r_2) to each point $(\alpha_1, \alpha_2) \in \mathbb{T}^2$ as a *label*. The problem becomes therefore a *two-dimensional problem of testing congruence under translation for labeled points on a torus*. We denote the two labeled point sets as \tilde{A} and \tilde{B} , and a point of \tilde{A} can only be mapped to a point of \tilde{B} with the same label.

Points in the x_1y_1 -plane should be considered separately, because α_2 is not unique when $r_2 = 0$, and the same problem occurs for the points in the x_2y_2 -plane. We will defer the treatment of these points to the end of this section, and start with other points first.

We now give an algorithm for the following problem: given two labeled point sets \tilde{A} and \tilde{B} on the torus \mathbb{T}^2 , test if \tilde{A} and \tilde{B} are the same up to translations. We will find a canonical set of \tilde{A} (and \tilde{B}) which is similar to a condensed set. In contrast to a condensed set, we add no new symmetries to a canonical set, by updating labels to preserve complete information. Let $\text{Sym}(A)$ for a set $A \subset \mathbb{T}^2$ denote translational symmetry group of A , i. e., the set of translations that map A to itself and preserve labels, if A is a labeled set.

22.1 A Canonical Set Procedure

Roughly, our goal is to find simplest subsets A' and B' of representative points that still preserve the same symmetries as A and B respectively. In addition, we want to construct A' without arbitrary decisions. If we can find such subsets, it is enough to compare an arbitrary point in A' with that in B' . We construct A' for a given point set A in some space \mathcal{S} and some group Θ of symmetries of \mathcal{S} in two cases: (i) $\mathcal{S} = \mathbb{S}^1$ (the unit circle), and Θ are the rotations of \mathbb{S}^1 . (ii) $\mathcal{S} = \mathbb{S}^1 \times \mathbb{S}^1$ (the flat torus), and Θ are the translations on \mathcal{S} .

Now, we formally define a *canonical set procedure* and explain how to use it. We denote by $\text{Sym}_\Theta(A)$ the symmetry group of A within Θ :

$$\text{Sym}_\Theta(A) = \{ R \in \Theta \mid R(A) = A \text{ and } R \text{ preserves labels} \}.$$

Definition 23. A *canonical set procedure* f_{cano} for a space \mathcal{S} and a subgroup Θ of the symmetries of \mathcal{S} maps every finite set $A \subset \mathcal{S}$ to a set $A' \subset \mathcal{S}$ such that the following properties hold.

1. Symmetries are preserved: $\text{Sym}_\Theta(A') = \text{Sym}_\Theta(A)$
2. $\text{Sym}_\Theta(A')$ acts transitively on A' : For every $p, q \in A'$, there exists $R \in \text{Sym}_\Theta(A')$ that maps p to q .
3. A' is defined in an equivariant way from A : If $RA = B$ for some $R \in \Theta$, and B is mapped to B' , then $B' = RA'$.

We call A' a *canonical set* of A .

It follows from the definition that the canonical set A' is nonempty whenever A is nonempty, provided that Θ is an infinite group. It is easy to see that condition 3 is implied by the \supseteq part of condition 1. Thus, if the procedure is equivariant, we only have to prove that A' does not have more symmetries than A .

A canonical set procedure is used to check congruence of two sets $A, B \subset \mathcal{S}$ under a congruence from the class Θ as follows: We compute their canonical point sets A' and B' . We then

pick two arbitrary points $p \in A'$ and $q \in B'$, we find the unique rotation $R \in \Theta$ that maps p to q , and finally, we simply have to check whether $RA = B$. In the above two cases of \mathcal{S} and Θ , there is always such a unique R .

The correctness of this approach is formulated in the following easy lemma.

Lemma 24. *Suppose that \mathcal{S} is a space in which, for any two points $p, q \in \mathcal{S}$, there is a unique symmetry $R = R_{pq} \in \Theta$ with $Rp = q$. Let A' and B' be canonical sets of A and B for \mathcal{S} and Θ respectively. There exists a congruence in Θ that maps A to B if and only if for any $p \in A', q \in B'$, we have $R_{pq}A = B$.*

Proof. Suppose there is a congruence $R_0 \in \Theta$ such that $R_0(A) = B$ and $p \in A'$ and $q \in B'$ are chosen arbitrarily. Since A' and B' are obtained by a canonical set procedure, $B = R_0(A)$ is mapped to $B' = R_0(A')$. Let p' be the preimage of q under R_0 . Since A' is a canonical point set, there exists a unique $R_1 \in \text{Sym}_\Theta A' \subset \text{Sym}_\Theta A$ satisfying $R_1(p) = p'$ and $R_1(A) = A$. By taking $R = R_0 \cdot R_1$, R is the unique rotation satisfying $R(p) = R_0(R_1(p)) = R_0(p') = q$ and $R(A) = R_0(R_1(A)) = R_0(A) = B$. Also, $R \in \Theta$ since $R_0, R_1 \in \Theta$ and Θ is closed under the multiplication.

The other direction is obvious. □

In the case (i) that $\mathcal{S} = \mathbb{S}^1$, the map from a set of points in \mathbb{S}^1 to a set of *canonical axes* (refer to Section 12) is the obvious canonical set procedure. It yields a set of regularly spaced directions on \mathbb{S}^1 .

A canonical set procedure for case (ii) is presented in Lemma 25 with Algorithm T.

Algorithm T (*A canonical set procedure from a labeled point set on the torus*). The input is a labeled point set \tilde{A} on the torus \mathbb{T}^2 . We assume that two labels can be compared in constant time. The output is an unlabeled *canonical set* \hat{A} as in Definition 23. In addition, \hat{A} should be obtained from \tilde{A} without making any arbitrary decisions.

- T1.** [Prune by labels.] Choose the label that occurs least frequently in \tilde{A} , and let A' be the set of points with this label. (For a while we will now do ordinary pruning, using only the geometry of the point set A' .)
- T2.** [Compute the Voronoi diagram.] Compute the Voronoi diagram of A' on the torus \mathbb{T}^2 . This can be reduced to a Voronoi diagram in the plane by replicating the square region representing the torus together with the set A' 9 times in a 3×3 tiling, see for example [15]. Clipping the result to the central tile yields the Voronoi diagram on the torus in $O(|A'| \log |A'|)$ total time.
- T3.** [Prune by shape.] Translate each point $a \in A'$ to the origin together with its Voronoi cell. If the translated cells are not all equal, replace A' by the subset of points whose cell shape occurs least frequently, and return to T2.
- T4.** [Restore information from the original set \tilde{A} by labeling the points in A' .] (Now all Voronoi cells are translated copies of the same hexagon or rectangle.) Assign each point of \tilde{A} to its Voronoi cell. A point on the boundary is assigned to all incident cells. Now for each point $x \in A'$, collect the points in its cell and translate them so that x lies at the origin. Represent each point as a triple (φ -coordinate, ψ -coordinate, label). Concatenate these triples in lexicographic order into a string of numbers that represents the cell contents, and attach the string as a label to the point x . (This string representation is obtained equivariantly, since two points $x, y \in A'$ get the same string if and only if the two Voronoi cells are exact translated copies of each other, including all points in the cells with their original labels. We have thus preserved complete information about the symmetries of \tilde{A} .)

- T5.** [Compress labels.] Sort the strings and replace each label with its rank in sorted order.
- T6.** [Finalize.] If there are at least two labels, return to T1. Otherwise, return A' as the canonical set \hat{A} .

Lemma 25. *Algorithm T computes a canonical set \hat{A} of a labeled set \tilde{A} on the torus in time $O(|\tilde{A}| \log |\tilde{A}|)$.*

Proof. We first check that \hat{A} has the claimed properties. Property 23.1 consists of two inclusions: We have $\text{Sym}(\tilde{A}) \subseteq \text{Sym}(\hat{A})$, because \hat{A} is obtained in an equivariant way from \tilde{A} . None of the operations which are applied to \tilde{A} to obtain \hat{A} destroys any translational symmetries. The other inclusion $\text{Sym}(\hat{A}) \subseteq \text{Sym}(\tilde{A})$ is ensured by Step T4. (This would work for any set A' .)

To see Property 23.2, note that the Voronoi cell of a point fixes the relative positions of its neighbors. Starting from any point $a \in \hat{A}$ we can reconstruct the whole set \hat{A} if we know the shape of each Voronoi cell. Step T3 ensures that all Voronoi cells are equal, and therefore the reconstructed set $\hat{A} - a$ is the same, no matter from which point a we start. The set \hat{A} forms a lattice on the torus.

Let us analyze the running time. Each iteration of the loop T2–T3 takes $O(|A'| \log |A'|)$ time and reduces the size of A' to half or less. Thus, the total running time of this loop is $O(|\tilde{A}| \log |\tilde{A}|)$, since the initial size of A' is bounded by the size of \tilde{A} .

Steps T4 and T5 involve point location in Voronoi diagrams [28] and sorting of strings of numbers of total length $O(|\tilde{A}|)$. These operations can be carried out in $O(|\tilde{A}| \log |\tilde{A}|)$ time. Thus, each iteration of the whole loop T1–T6 takes $O(|\tilde{A}| \log |\tilde{A}|)$ time. Step T1 reduces the size of \tilde{A} to a half or less after each iteration. Thus, the total running time is $O(|\tilde{A}| \log |\tilde{A}|)$. \square

We are ready for the main result of this section.

Theorem 26. *Given two sets A, B of n points in \mathbb{R}^4 , and two planes P and Q , it can be checked in $O(n \log n)$ time if there is a congruence that maps A to B , and P to Q .*

Proof. We have already seen how to convert input points to points $(\alpha_1, \alpha_2) \in \mathbb{T}^2$ with labels (d_1, d_2) so that congruence testing is reduced to testing congruence by translation on the torus.

We still have to deal with points on coordinate planes. Let $A_1 \subseteq A$ be the points in the x_1y_1 -plane, and let $A_2 \subseteq A$ be the points in the x_2y_2 -plane. If $A_1 \neq \emptyset$, we compute the canonical axes of A_1 , as described in Section 12. They form $k \geq 1$ equally spaced angular directions $\bar{\alpha} + j\frac{2\pi}{k}$ modulo 2π , $j \in \mathbb{Z}$. We know that R must map the canonical axes of A_1 to the canonical axes for the corresponding set B_1 . We incorporate this restriction into an additional label for each point $u \in A \setminus A_1 \setminus A_2$. If u has a polar angle α_1 in the x_1y_1 -plane, we attach to it the difference to the nearest smaller canonical angle:

$$\min\{ \alpha_1 - (\bar{\alpha} + j\frac{2\pi}{k}) \mid j \in \mathbb{Z}, \alpha_1 - (\bar{\alpha} + j\frac{2\pi}{k}) \geq 0 \}$$

We also have to test if A_1 and B_1 are congruent by simply overlaying the canonical axes and testing if they are equal.

If $A_2 \neq \emptyset$, we treat this in the same way and attach an additional label to the points in $A \setminus A_1 \setminus A_2$.

After this, we can just compute the canonical set from Lemma 25 and thereby reduce the congruence test to an equality test between sets \hat{A} and \hat{B} , as described before in Lemma 24. If there are no points outside the xy -plane and the zw -plane, the problem reduces to two independent congruence tests in these planes. The whole algorithm takes $O(n \log n)$ time. \square

There is an alternative algorithm with the same time complexity that does not use Lemma 24.

Alternative Proof of Theorem 26. We can also use Algorithm T for comparing two labeled sets \tilde{A} and \tilde{B} on the torus as follows. We run Algorithm T in parallel for \tilde{A} and \tilde{B} . We must make sure that all steps run identically. In particular, the sorted strings in Step T5 must be the same for \tilde{A} and \tilde{B} . If the algorithm runs to the end without finding a difference between \tilde{A} and \tilde{B} , and if the Voronoi cells of the canonical sets \hat{A} and \hat{B} are equal, we know that \tilde{A} and \tilde{B} are congruent by translation. There are $|\hat{A}| = |\hat{B}|$ translations that map \tilde{A} to \tilde{A} . Points on coordinate planes can be taken care of by the same way as the previous proof, and the overall algorithm also takes $O(n \log n)$ time. \square

23 Concluding Remarks

We can substitute our algorithm for 4 dimensions as the base case in the recursive dimension-reduction algorithms of Akutsu [1] and Brass and Knauer [10]. The deterministic algorithm of Brass and Knauer runs now in $O(n^{\lceil (d-1)/3 \rceil} \log n)$ time, and the randomized Monte Carlo algorithm of Akutsu takes time $O(n^{\lfloor (d-1)/2 \rfloor / 2} \log n)$ for $d \geq 9$ and $O(n^{3/2} \log n)$ for $d \leq 8$.

It is likely that our algorithm can be simplified. The constants in the bounds are certainly larger than the actual bounds.

23.1 Practical Implementability

Instead of using the Real-RAM model, it also makes sense to test congruence with an error tolerance ε , but this problem is known to be NP-hard even in the plane [14, 20] as mentioned earlier in Section 2. However, the problem becomes polynomial if the input points are sufficiently separated compared to ε . We are confident that our algorithm, when it is implemented with standard floating-point arithmetic and with appropriate error margins to shield equality tests from round-off errors, would decide approximate congruence in the following weak sense. Given a tolerance ε that is large compared to the machine precision, and two sets A and B whose points are separated by more than, say 10ε , the algorithm would correctly distinguish the instances that are congruent with a tolerance of ε from the instances that are not even congruent with a tolerance of, say 100ε . Between ε and 100ε , the algorithm is allowed to make errors. Such a result will require a thorough analysis of the numerical stability of the operations in our algorithm.

23.2 Regularity and Related Open Questions

Does Local regularity imply global regularity? The general theme in our four-dimensional algorithm raises the question whether *local* regularity implies *global* regularity. In other words, when the neighborhoods of all the points look the same, does it imply that a symmetry group acts transitively on the point set, if the neighborhoods are sufficiently large? It would be interesting to prove such a statement quantitatively.

Geometric classification of 4-dimensional point groups. A *point group* is a discrete subgroup $O(d)$, or in other words, a finite group of orthogonal $d \times d$ matrices. Hessel's Theorem [27] gives an explicit classification of the point groups for $d = 3$. A somewhat implicit formulation of Hessel's theorem is stated as follows.

Theorem 27. *A three-dimensional point group is one of the following:*

- (a) *the symmetry group of one of the five three-dimensional platonic solids,*
- (b) *the symmetry group of a prism over a regular polygon,*
- (c) *or a subgroup of one of above groups.*

This theorem can be derived without much effort from Lemma 20 (Algorithm K).

The four-dimensional point groups were already enumerated [31, 12]. The question is whether the four-dimensional point group can be enumerated with geometrically interpretable characteristics. The following is a conjecture by Rote [29].

Conjecture 28. *A four-dimensional point group is either*

- (a) *the symmetry group of one of the five four-dimensional regular polytope,*
- (b) *a direct product of lower-dimensional points groups,*
- (c) *or a subgroup of one of the above groups.*

Since the list of four-dimensional point groups is known, it should in principle be easy to settle this conjecture. However, it would be nice to established it directly, without resorting to the previous classification. Felix Klein, in his book on non-Euclidean geometry from 1928, writes that the determination of the four-dimensional point groups should be “easy” [23, pp. 240–241]¹.

Tilings on a sphere. Another interesting question is the classification of regular and semiregular tilings on spheres. It is possible that our techniques can shed light on symmetries on the 3-sphere, or on regular and semiregular tilings of S^3 .

References

- [1] T. Akutsu. On determining the congruence of point sets in d dimensions. *Computational Geometry: Theory and Applications*, 4(9):247–256, 1998.
- [2] H. Alt, K. Mehlhorn, H. Wagnen, and E. Welzl. Congruence, similarity, and symmetries of geometric objects. *Discrete & Computational Geometry*, 3(1):237–256, 1988.
- [3] D. Asimov. The grand tour: a tool for viewing multidimensional data. *SIAM Journal on Scientific and Statistical Computing*, 6(1):128–143, 1985.
- [4] M. J. Atallah. On symmetry detection. *IEEE Transactions on Computers*, 100(7):663–666, 1985.
- [5] M. D. Atkinson. An optimal algorithm for geometrical congruence. *Journal of Algorithms*, 8(2):159–172, 1987.
- [6] J. L. Bentley and M. I. Shamos. Divide-and-conquer in multidimensional space. In *Proceedings of the eighth annual ACM symposium on Theory of computing*, pages 220–230. ACM, 1976.
- [7] M. Berger. *Geometry*, volume I. Springer-Verlag, 1996.
- [8] M. Berger. *Geometry*, volume II. Springer-Verlag, 1996.
- [9] Å. Björck and G. H. Golub. Numerical methods for computing angles between linear subspaces. *Mathematics of Computation*, 27(123):579–594, 1973.
- [10] P. Brass and C. Knauer. Testing the congruence of d -dimensional point sets. *International Journal of Computational Geometry and Applications*, 12(1&2):115–124, 2002.
- [11] J. H. Conway, R. H. Hardin, and N. J. Sloane. Packing lines, planes, etc.: Packings in Grassmannian spaces. *Experimental Mathematics*, 5(2):139–159, 1996.

¹Die Kugeldrehungen der euklidischen Geometrie lassen sich eineindeutig den wesentlich voneinander verschiedenen Quaternionen sowie den Schiebungen jeder der beiden Arten so zuordnen, dass die zugehörigen Gruppen isomorph sind. [...] Die Kenntnis dieser Isomorphie ist sehr nützlich; denn mit ihrer Hilfe ist es z. B. leicht, die endlichen, d. h. die nur aus endliche vielen Bewegungen bestehenden *Bewegungsgruppen* der räumlichen elliptischen Geometrie zu bestimmen. [Emphasis in the original]

- [12] J. H. Conway and D. A. Smith. *On Quaternions and Octonions: their Geometry, Arithmetic, and Symmetry*. AK Peters, Wellesley, Massachusetts, 2003.
- [13] H. S. M. Coxeter. *Regular Polytopes*. Dover Publications, 3rd edition, 1973.
- [14] C. Dieckmann. *Approximate Symmetries of Point Patterns*. PhD thesis, Freie Universität Berlin, 2012.
- [15] N. P. Dolbilin and D. H. Huson. Periodic Delone tilings. *Periodica Mathematica Hungarica*, 34(1-2):57–64, 1997.
- [16] W. Fenchel. Über Krümmung und Windung geschlossener Raumkurven. *Mathematische Annalen*, 101(1):238–252, 1929.
- [17] G. H. Golub and C. F. Van Loan. *Matrix computations*, volume 3. Johns Hopkins University Press, 2012.
- [18] P. T. Highnam. Optimal algorithms for finding the symmetries of a planar point set. *Information Processing Letters*, 22(5):219–222, 1986.
- [19] J. E. Hopcroft and R. E. Tarjan. A $V \log V$ algorithm for isomorphism of triconnected planar graphs. *Journal of Computer and System Sciences*, 7(3):323–331, 1973.
- [20] S. Iwanowski. Testing approximate symmetry in the plane is NP-hard. *Theoretical Computer Science*, 80(2):227–262, 1991.
- [21] H. Kim. *Congruence Testing for Point Sets in 4-Space*. PhD thesis, Freie Universität Berlin, 2016. submitted.
- [22] H. Kim and G. Rote. Congruence testing of point sets in 4-space. In S. Fekete and A. Lubiw, editors, *32st International Symposium on Computational Geometry (SoCG 2016)*, Leibniz International Proceedings in Informatics (LIPIcs). Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2016. to appear.
- [23] F. Klein. *Vorlesungen über nicht-euklidische Geometrie*, volume 28 of *Grundlehren der math. Wissenschaften*. Springer, 1928.
- [24] D. E. Knuth, J. H. Morris, Jr, and V. R. Pratt. Fast pattern matching in strings. *SIAM Journal on Computing*, 6(2):323–350, 1977.
- [25] G. Manacher. An application of pattern matching to a problem in geometrical complexity. *Information Processing Letters*, 5(1):6–7, 1976.
- [26] H. P. Manning. *Geometry of Four Dimensions*. Macmillan, 1914.
- [27] G. E. Martin. *Transformation Geometry: An Introduction to Symmetry*. Springer, 1982.
- [28] F. P. Preparata and M. I. Shamos. *Computational Geometry: an Introduction*. Springer, 1985.
- [29] G. Rote. Congruence testing of point sets in three and four dimensions—results and techniques (invited talk). In I. Kotsireas, S. Rump, and C. Yap, editors, *Sixth International Conference on Mathematical Aspects of Computer and Information Sciences—MACIS 2015*, volume 9582 of *Lecture Notes in Computer Science*, page 10 pages. Springer-Verlag, 2016.
- [30] K. Sugihara. An $n \log n$ algorithm for determining the congruity of polyhedra. *Journal of Computer and System Sciences*, 29(1):36–47, 1984.

- [31] W. Threlfall and H. Seifert. Topologische Untersuchung der Diskontinuitätsbereiche endlicher Bewegungsgruppen des dreidimensionalen sphärischen Raumes. *Mathematische Annalen*, 104(1):1–70, 1931.
- [32] J. D. Ullman, A. V. Aho, and J. E. Hopcroft. *The Design and Analysis of Computer Algorithms*. Addison-Wesley, Reading, 1974.
- [33] Y.-C. Wong. Differential geometry of Grassmann manifolds. *Proceedings of the National Academy of Sciences of the United States of America*, 57(3):589–594, 1967.