Improved orientation sampling for indexing diffraction patterns of polycrystalline materials

Larsen, Peter Mahler; Schmidt, Søren

Published in:
Journal of Applied Crystallography

Link to article, DOI:
[10.1107/S1600576717012882](https://doi.org/10.1107/S1600576717012882)

Publication date:
2017

Document Version
Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

Citation (APA):
Larsen, P. M., & Schmidt, S. (2017). Improved orientation sampling for indexing diffraction patterns of polycrystalline materials. *Journal of Applied Crystallography*, 50(6). <https://doi.org/10.1107/S1600576717012882>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Improved orientation sampling for indexing diffraction patterns of polycrystalline materials

Peter Mahler Larsen* and Søren Schmidt

Department of Physics, Technical University of Denmark, 2800 Kongens Lyngby, Denmark. *Correspondence e-mail: pmla@fysik.dtu.dk

Received 12 July 2017

Accepted 8 September 2017

Edited by Th. Proffen, Oak Ridge National Laboratory, USA

Keywords: orientation sampling; pattern indexing; forward modelling; quaternions; spherical coverings.

Orientation mapping is a widely used technique for revealing the microstructure of a polycrystalline sample. The crystalline orientation at each point in the sample is determined by analysis of the diffraction pattern, a process known as pattern indexing. A recent development in pattern indexing is the use of a brute-force approach, whereby diffraction patterns are simulated for a large number of crystalline orientations and compared against the experimentally observed diffraction pattern in order to determine the most likely orientation. Whilst this method can robustly identify orientations in the presence of noise, it has very high computational requirements. In this article, the computational burden is reduced by developing a method for nearly optimal sampling of orientations. By using the quaternion representation of orientations, it is shown that the optimal sampling problem is equivalent to that of optimally distributing points on a four-dimensional sphere. In doing so, the number of orientation samples needed to achieve a desired indexing accuracy is significantly reduced. Orientation sets at a range of sizes are generated in this way for all Laue groups and are made available online for easy use.

1. Introduction

In many types of diffraction experiments, the aim is to determine the orientation of the diffracted crystallite volume which creates the experimentally observed pattern. For example, when studying a multigrain sample with the three-dimensional X-ray diffraction technique (3DXRD; Poulsen *et al.*, 2001), a 'grain map' is constructed by finding the crystalline orientation at each point in the sample. The process of determining the crystalline orientation from a diffraction pattern on the detector is known as pattern indexing. Throughout this article, we use the term 'experimental pattern' to denote an image of a diffraction pattern as recorded on a detector.

The most widely used pattern indexing methods work 'backwards' from features in the observed data to an orientation. Such methods are typically highly efficient, but can fail in the presence of noise. A well known example is in electron backscatter diffraction (EBSD), where the Hough transform is used to find lines in the backscattered Kikuchi pattern, from which the orientation can be determined (Adams *et al.*, 1993). Under noisy conditions, however, the Kikuchi lines can no longer be reliably identified and the indexing process fails as a consequence.

The desire to analyse diffraction patterns under less-than-ideal conditions has motivated the development of forward-modelling-based pattern indexing, also known as dictionary-based indexing. In a forward model, rather than working backwards from the data, the orientation is found using a brute-force approach. A dictionary is constructed by selecting

a set of orientations and generating simulated patterns for each of them. A requirement for simulating patterns is that the crystal phase is known *a priori*, or, if indexing a multiphase material, that the set of candidate phases is known.

To index an experimental pattern, it is compared against every simulated pattern in the dictionary, and the dictionary pattern with the highest similarity determines the orientation (in a multiphase material this also determines the phase). Here, the similarity is determined by the difference in the pixel intensities in the simulated and experimental patterns. By using the full image information (*i.e.* all pixel intensities), the similarity exhibits a continuous degradation with increasing noise, as opposed to the catastrophic degradation exhibited when looking for specific features in the experimental pattern.

A significant drawback of the forward-modelling approach, however, is the computational effort required: each experimental pattern must be tested against every dictionary pattern. Since, the accuracy of the pattern indexing process depends on the granularity of the set of dictionary orientations, a more accurate indexing requires a larger set. Increasing the number of dictionary orientations, however, increases the time required to index a pattern. Since the objectives of increased accuracy and reduced running time are in opposition to each other, we ask the question, how can we achieve the highest accuracy with the fewest dictionary orientations? In this article, we describe a method for doing so with the use of quaternions.

1.1. Measurement of dictionary orientation sets

Orientations can be conveniently represented using unit quaternions (Altmann, 2005). Briefly, a quaternion is a four-dimensional vector of the form $q = \{w, ix, jy, kz\}$, where w, x, y and z are real numbers and i, j and k are imaginary numbers which generalize the better-known complex numbers. Unit quaternions represent points on a four-dimensional hypersphere, a space formally known as $\mathbb{S}^3$ and which consists of all vectors that satisfy $(w^2 + x^2 + y^2 + z^2)^{1/2} = 1$. This space is a double covering of $SO(3)$, the group of rotations in three-dimensional Euclidean space ($\mathbb{R}^3$). The double covering relationship means that $-q$ and q represent the same orientation, which is evident when considering the quaternion-derived rotation matrix:

$$\mathbf{U}_q = \begin{bmatrix} 1 - 2y^2 - 2z^2 & 2xy - 2wz & 2xz + 2wy \\ 2xy + 2wz & 1 - 2x^2 - 2z^2 & 2yz - 2wx \\ 2xz - 2wy & 2yz + 2wx & 1 - 2x^2 - 2y^2 \end{bmatrix}. \quad (1)$$

It can be seen that in each element of $\mathbf{U}_q$ the sign of the quaternion cancels out. By using the quaternion representation, the problem of selecting an optimal set of dictionary orientations is equivalent to finding an optimal distribution of a set of points on $\mathbb{S}^3$. To do so, we must first decide what constitutes a good distribution.

The misorientation between two orientations in quaternion form, p and q , is given by

$$\alpha(p, q) = 2 \arccos | \langle p, q \rangle |, \quad (2)$$

where $\langle p, q \rangle$ denotes the inner product of p and q . In many previous studies, dictionary orientation sets are quantified by the misorientation between neighbouring orientations: for example, the average value of $\alpha(p, q)$ over all pairs of nearest neighbours p and q . This may be adequate when the orientation set has a known grid-like structure, but it does not constitute a universal measure of quality. To illustrate this with a pathological example, consider an orientation set, $\mathbf{Q}$, where all orientations lie at the same point. The misorientation between all pairs of orientations is zero, that is

$$\alpha(p, q) = 0 \quad \forall p \in \mathbf{Q}, \quad q \in \mathbf{Q}, \quad (3)$$

yet the set constitutes the worst possible dictionary. A good measure of quality should instead consider the misorientation between the dictionary set and any possible experimental orientation. We define the error term as the maximum misorientation between these two, *i.e.* how far can an experimental orientation lie from the dictionary? More specifically, this error term is given by

$$\alpha_{\max} = \max_{x \in SO(3)} \min_{q \in \mathbf{Q}} \alpha(x, q). \quad (4)$$

This quantity can be minimized by solving the spherical covering problem in $\mathbb{S}^3$. Given N hyperspherical caps of equal radius, r , called the covering radius, the spherical covering problem asks how to arrange the caps to cover the surface of $\mathbb{S}^3$ with minimal r . We describe this problem in detail in §2.

By creating orientation sets with a small covering radius, we can either reduce the number of orientations required to achieve a desired error tolerance (thereby reducing the running time of forward-modelling pattern indexing) or simply improve the error distribution for a fixed number of orientations. Creation of such sets is the principal contribution of this work.

1.2. Previous work

Forward modelling has been successfully applied in many types of diffraction-based experiments, including the indexing of three-dimensional X-ray diffraction microscopy data (Li & Suter, 2013; Schmidt, 2014), EBSD data (Chen *et al.*, 2015) and electron channelling patterns (Singh & De Graef, 2017). Any forward-modelling method requires a discretization of $SO(3)$. Whilst many such discretization methods have been developed, here we consider only three which are both successful and commonly used amongst crystallographers.

Yershova *et al.* (2010) have developed an incremental infinite sequence based on the Hopf fibration. The method generates orientations deterministically, with proven maximal dispersion reduction when used as a sequence. Furthermore, the orientation sets are isotitudinal, which permits expansion into spherical harmonics (Dahms & Bunge, 1989), refinable and can be generated on the fly. Whilst the method has many desirable properties, it is developed for the purpose of robot-motion planning and is not easily integrated with crystallographic fundamental zones. To remedy this, Roşca *et al.* (2014) have developed ‘cubochoric’ coordinates, in which an area-preserving Lambert projection is used to map points

Table 1

Summary of properties of different methods of orientation set generation.

Existing methods prioritize fast generation and a grid-like structure. In our work we optimize the covering radius at the expense of all other properties. The optimality gap for a set of N orientations is the percentage difference of its covering radius to that of the simplex bound (*cf.* §2.5)

Method	Fast generation	Refinable	Isolatitudinal	Isochoric	Crystallographic	Optimal gap at $N \simeq 10^5$ (%)
Random sampling	✓	✓	–†	–	✓	127
Hopf fibration	✓	✓	✓	–	–	59.9
Cubochoric	✓	✓	✓	✓	✓	40.8
Octahedral b.c.c.	✓	✓	–†	–	–‡	15.4
Present work	–	–	–†	–	✓	4.64

† Non-isolatitudinal sets do not permit an expansion into spherical harmonics, though any orientation set can be expanded into hyperspherical harmonics (Mason & Schuh, 2008; Mason, 2009). ‡ These orientation sets can be mapped out into seven of 11 Laue group fundamental zones (*cf.* §2.4).

from a cubic grid onto any desired crystallographic fundamental zone in $SO(3)$. A different approach, developed by Karney (2007) for use in molecular modelling, is to generate sets which attempt to solve the spherical covering problem. Inspired by the observation that body-centred cubic (b.c.c.) grids solve the covering problem in $\mathbb{R}^3$, b.c.c. grids are constructed in Rodrigues–Frank (RF) space (Frank, 1988; Morawiec & Field, 1996) in order to generate good coverings in $SO(3)$.

Table 1 summarizes the properties of the different methods of generating orientation sets. Each of the three aforementioned techniques attempts to solve slightly different problems and involves different trade-offs as a consequence, although one feature they have in common is fast generation. We take an alternative approach, sacrificing other properties in pursuit of creating the ‘best’ possible orientation sets. Whilst this approach requires a significant up-front computational effort, this is a good trade-off when the resulting sets will subsequently be used many times. We emphasize that, whilst the orientational error is critical to forward modelling, there are many other sources of error in any modality [see Ram *et al.* (2017) for a comprehensive analysis in an EBSD context].

The rest of this article is organized as follows: in §2 we define the spherical covering problem on $\mathbb{S}^d$, show how this relates to the problem of finding an optimal set of orientations and derive a conjectured lower bound. We describe the generation of orientation sets in §3. Results on the covering radius and error distributions of the resulting orientation are given in §4. Lastly, the advantages and drawbacks of the method presented are discussed in §5.

2. Error quantification of orientation sets

In order to compare different orientation sets we must define a measure of quality. Here, we describe the covering radius of a set, which we argue is the canonical error measure since it determines the maximum possible error. We will first describe the sphere covering problem for Euclidean and spherical

geometries, and then show that the problem of generating optimal orientation sets is a special case of the spherical covering problem.

2.1. Spherical coverings

The sphere covering problem is best known in Euclidean geometries. In $\mathbb{R}^d$, it asks ‘for the most economical way to cover d -dimensional space with equal overlapping spheres’ (Conway & Sloane, 1998). Optimal coverings are known for $d = 1$ and $d = 2$, which are equally spaced points on a line and a hexagonal lattice, respectively, and optimal lattice coverings are known for $1 \leq d \leq 5$.

The presence of curvature in spherical geometries renders the covering problem vastly more challenging. In $\mathbb{S}^d$, the spherical covering problem asks for the most economical way to cover the surface of $\mathbb{S}^d$ with equal overlapping hyperspherical caps. In $\mathbb{S}^1$, the optimal covering is a set of N points with angle $2\pi/N$ between adjacent points. For $d > 1$, however, there is no general formula for determining the optimal spherical covering. Furthermore, unlike in $\mathbb{R}^d$, the configuration of the optimal covering depends on the number of points in the covering. For example, for $d = 2$, the known optimal configurations are the vertices of the tetrahedron, the octahedron and the icosahedron. Hardin *et al.* (2017) have found putatively optimal coverings for $d = 2$ at other values of N , but these have been found using numerical optimization and are not provably optimal.

2.2. Covering radius and covering density

For coverings on $\mathbb{S}^d$, the two (equivalent) measures of quality are the covering radius and the covering density. Given a discrete collection of points $\mathbf{P} = \{p_1, p_2, \dots, p_N\} \in \mathbb{S}^d$, the covering radius, θ , is defined as the largest angular distance between any point in $\mathbb{S}^d$ and $\mathbf{P}$, that is

$$\theta = \max_{x \in \mathbb{S}^d} \min_{p \in \mathbf{P}} \arccos \langle x, p \rangle, \quad (5)$$

where $\langle x, p \rangle$ denotes the inner product of x and p . Then, $\mathbf{P}$ covers the surface of $\mathbb{S}^d$ with $N = |\mathbf{P}|$ equal hyperspherical caps of radius θ . The covering density, $\tau_d(\theta)$, is given by the ratio of the sum of the surface area of the caps to the surface area of a unit d -sphere,

$$\tau_d(\theta) = N \frac{C_d(\theta)}{S_d(1)}, \quad (6)$$

where

$$C_d(\theta) = \int_0^{\tan(\theta)} \frac{S_{d-1}(r)}{(1+r^2)^2} dr, \quad S_{d-1}(\theta) = \frac{d\pi^{d/2}}{\Gamma(d/2+1)} \theta^{d-1}. \quad (7)$$

$S_{d-1}(\theta)$ is the surface area of the d -sphere of radius θ and $C_d(\theta)$ is the surface area of a hyperspherical cap of radius θ (*cf.* Appendix A for derivation). To find the covering radius, we need to determine the Voronoi cell of each point $p_i \in \mathbf{P}$. The Voronoi cell of point p_i , denoted $\text{Vor}(p_i)$, consists of all points of $\mathbb{S}^d$ that are at least as close to p_i as to any other p_j . More specifically

$$\text{Vor}(p_i) = \{x \in \mathbb{S}^d \mid \arccos\langle x, p_i \rangle \leq \arccos\langle x, p_j \rangle \ \forall j\}. \quad (8)$$

Since the vertices of the Voronoi cells are the points which locally maximize the angular distance from $\mathbf{P}$, the covering radius is determined by the Voronoi vertex that lies furthest from $\mathbf{P}$.

The Voronoi cells of a set of points in $\mathbb{S}^d$ are not easy to calculate directly, so instead we calculate the (hyperspherically constrained) Delaunay triangulation. The Delaunay triangulation, $\text{DT}(\mathbf{P})$, is a set of hyperspherical simplices whose vertices are points in $\mathbf{P}$ which satisfy the empty-sphere condition, that is, no points in $\mathbf{P}$ lie inside the circumhypercap of any simplex in $\text{DT}(\mathbf{P})$. Associated with each simplex is a Voronoi vertex, which lies at the centre of the simplex circumhypercap (the circumcentre). The Delaunay triangulation, Voronoi cells, and simplex circumhypercaps and circumcentres are illustrated in $\mathbb{S}^2$ in Fig. 1. We now show how to calculate the circumcentre of a simplex.

Theorem. For a hyperspherical simplex $\mathbf{t} \in \text{DT}(\mathbf{P})$ with vertices $\{p_1, p_2, p_3, \dots, p_{d+1}\} \in \mathbb{S}^d$, the position of the circumcentre, X , is equal to the unit normal vector of the d -dimensional hyperplane on which the vertices of $\mathbf{t}$ lie.

Proof. Let $\mathbf{S} = \{s_1 = p_2 - p_1, s_2 = p_3 - p_1, s_3 = p_4 - p_1, \dots, s_d = p_{d+1} - p_1\}$ and let $X \in \mathbb{S}^d$ be the circumcentre of $\mathbf{t}$. Then, per definition, X must satisfy

$$p_i X = p_1 X \quad \forall i. \quad (9)$$

Subtracting $p_1 X$ from each side gives

$$s_i X = 0 \quad \forall i. \quad (10)$$

The unit length of X follows from requiring $X \in \mathbb{S}^d$.

To find X , we calculate the normalized d -fold vector cross product (Brown & Gray, 1967) of $\mathbf{S}$. Since every hyperplane has two (opposite) unit plane normals, X has two solutions,

Figure 1

Left: a putatively optimal spherical covering for 28 points in $\mathbb{S}^2$ [point set from Hardin *et al.* (2017)]. The solid lines indicate the spherically constrained Delaunay triangulation. The dashed lines indicate the Voronoi cells. Right: the same points; each simplex in the Delaunay triangulation has a circumcap, the centre of which (marked in red) lies at a Voronoi cell vertex. The maximum simplex circumradius determines the covering radius of the point set.

which correspond to the centre of the simplex hypercircumcap and its dual. However, given that $|\mathbf{P}| \geq d + 2$ only one of these solutions fulfils the empty-sphere condition, which is the one which satisfies $\langle X, p_i \rangle > 0 \ \forall i$. This corresponds to the smaller of the two hypercircumcaps.

For a set of points on $\mathbb{S}^d$, the vertices of each simplex $\mathbf{t} \in \text{DT}(\mathbf{P})$ can be found by calculating the convex hull of $\mathbf{P}$, as shown in Fig. 2. If we denote the circumradius of a simplex $\mathbf{t}$ by $\varphi(\mathbf{t})$, equation (5) can be restated as

$$\theta = \max\{\varphi(\mathbf{t}) \mid \mathbf{t} \in \text{DT}(\mathbf{P})\}, \quad (11)$$

which provides a practical solution to equation (5): the covering radius of a point set is simply the maximum simplex circumradius.

2.3. Orientation sets

The problem of finding a good spherical covering is immediately relatable to the problem of finding good sets of orientations. As described previously, rotations can be represented by quaternions, which are points on $\mathbb{S}^3$. The maximum rotational angle between a point $x \in \text{SO}(3)$ and a point set $\mathbf{P}$, also called the maximum misorientation, is given by

$$\begin{aligned} \alpha_{\max} &= 2 \max_{x \in \mathbb{S}^3} \min_{p \in \mathbf{P}} [\arccos\langle -x, p \rangle, \arccos\langle x, p \rangle], \\ &= 2 \max_{x \in \mathbb{S}^3} \min_{p \in \mathbf{Q}} \arccos\langle x, p \rangle, \end{aligned} \quad (12)$$

where $\mathbf{Q} = \mathbf{P} \cup \{-p \mid p \in \mathbf{P}\}$. It can be seen that, for a point set with antipodal symmetry, $\alpha_{\max} = 2\theta$, that is, the maximum misorientation is twice the covering radius. Thus, the problem of finding a set of rotations with the lowest maximum misorientation is equivalent to finding an optimal spherical covering for a point set with antipodal symmetry on $\mathbb{S}^3$.

2.4. Integration with crystallographic symmetries

Equation (12) shows that a set of $2N$ points with antipodal symmetry represents a set of N rotations. A set of orientations generated in this way covers the whole space of $\text{SO}(3)$ and is immediately applicable to pattern indexing of materials with triclinic (C_1) Bravais lattices. For materials with higher-order symmetry, though, a dictionary set which covers all of $\text{SO}(3)$ is

Figure 2

Convex hull (left) and the spherically constrained Delaunay triangulation (right) of 22 points on the sphere. The triangulations exist in $\mathbb{R}^3$ and $\mathbb{S}^2$, respectively, but the vertices of each simplex are the same. Data from Hardin *et al.* (2017).

wasteful, since only the fundamental zone orientations (He & Jonas, 2007) are needed. A naive approach for selecting fundamental zone orientations is to generate a full covering of $SO(3)$ and then simply ‘cut out’ the desired region; this introduces artifacts at the boundaries of the fundamental zone which increase the covering radius significantly. Instead, we apply the symmetry of the desired point group during generation of the orientation sets.

Given a set of basis points $\mathbf{B} = \{b_1, b_2, \dots\}$ and a quaternion group $\mathbf{G} = \{g_1, g_2, \dots\}$, we can create a set of points with the symmetry of $\mathbf{G}$ by

$$\mathbf{P} = \{b \otimes g \mid b \in \mathbf{B}, g \in \mathbf{G}\}, \quad (13)$$

where $\otimes$ denotes quaternion multiplication. If $\mathbf{P}$ is to represent a set of orientations [cf. equation (12)], $\mathbf{G}$ must be a superset of antipodal symmetry (C_1). The finite quaternion groups which meet this requirement are (Conway & Smith, 2003)

$2I_{60}$	The binary icosahedral group
$2O_{24}$	The binary octahedral group
$2T_{12}$	The binary tetrahedral group
$2D_n$	The binary dihedral group
$2C_n$	The binary cyclic group

With the exception of the binary icosahedral group, each of these is used to describe the generators of the 11 Laue groups (Morawiec, 2004), $C_1, C_2, C_3, C_4, C_6, D_2, D_3, D_4, D_6, T$ and O . By the application of a symmetry group, the problem of finding a good spherical covering for a chosen crystallographic fundamental zone is reduced to a problem of finding an optimal configuration of the basis points, which is a much smaller problem.

The Laue groups can be divided into two sets:

$$\{C_2, C_4, D_2, D_4, T, O\} \quad (14)$$

and

$$\{C_3, C_6, D_3, D_6\}, \quad (15)$$

where the elements of each are subsets of O and D_6 , respectively (C_1 is trivially a subset of both). This means that, if we generate sphere coverings with O and D_6 applied according to equation (13), then by an appropriate mapping of the fundamental zone orientations we obtain sphere coverings for all Laue groups, without the aforementioned boundary artifacts. The Laue group subset relationships are shown in Appendix B.

2.5. Derivation of the simplex bound on $\mathbb{S}^3$

In addition to knowing the covering radius and density of a point set, it is useful to know how far from optimality a set is. We can estimate the optimality gap with a lower bound.

The simplex bound is a classic result which gives an upper bound on the density of sphere packings and a lower bound on the density of sphere coverings. It has been proven for packings in $\mathbb{R}^d$ (Rogers, 1958) and $\mathbb{S}^d$ (Böröczky, 1978), and for coverings in $\mathbb{R}^d$ (Coxeter *et al.*, 1959) and $\mathbb{S}^2$ (Tóth, 1964). Böröczky has conjectured that it is a lower bound on $\mathbb{S}^3$ (Böröczky, 2004). Despite lacking a proof, we will use the

simplex bound on $\mathbb{S}^3$ to estimate the optimality of our point sets, as it is ‘intuitively obvious’.

The premise of the simplex covering bound is that the lowest covering density can be achieved with regular simplices; this concept is illustrated in Fig. 3. Regular simplices tessellate in $\mathbb{R}^1$ and $\mathbb{R}^2$. In $\mathbb{R}^d$ for $d \geq 3$, regular simplices do not tessellate, and thus the simplex covering density is an unattainable lower bound. As stated previously, regular simplices tessellate in $\mathbb{S}^2$ for three configurations: the tetrahedron, the octahedron and the icosahedron. Thus, the simplex bound is tight for these configurations only, and is provably unattainable for any other number of vertices. In $\mathbb{S}^3$, regular simplices tessellate only in the 5-cell, the 16-cell and the 600-cell. If Böröczky’s conjecture is correct, the simplex bound is tight only for these configurations. Since no description of the simplex bound covering density on $\mathbb{S}^3$ could be found in the literature, we derive an expression for it here.

Consider a hyperspherical cap on $\mathbb{S}^3$ of radius θ and volume $C_3(\theta)$. We denote the inscribed regular spherical tetrahedron $T(\theta)$. At each of the four vertices of $T(\theta)$ is a hyperspherical cap of radius θ . Each of these caps intersects $T(\theta)$ with solid angle $\Omega(\theta)$, giving a volume of intersection of $C_3(\theta)\Omega(\theta)/(4\pi)$. Now $T(\theta)$ is covered by the four equal volumes of intersection. The covering density, $\tau_{\mathbb{S}^3}$, is the ratio of the sum of the four volumes of intersection to the volume of $T(\theta)$:

$$\tau_{\mathbb{S}^3}(\theta) = 4C_3(\theta) \frac{\Omega(\theta)}{4\pi} \frac{1}{\text{Vol}[T(\theta)]}. \quad (16)$$

Here,

$$C_3(\theta) = \pi[2\theta - \sin(2\theta)], \quad (17)$$

$$\Omega(\theta) = 3\psi(\theta) - \pi, \quad (18)$$

$$\psi(\theta) = \arccos\left[\frac{4\cos^2(\theta) - 1}{8\cos^2(\theta) + 1}\right], \quad (19)$$

$$\text{Vol}[T(\theta)] = \{-\text{Re}(L) + \pi[\arg(-Q) + 3\psi(\theta)] - \frac{3}{2}\pi^2\} \bmod 2\pi^2, \quad (20)$$

Figure 3
Illustration of the simplex bound in $\mathbb{R}^2$, shown here because of the difficulty of visualizing the simplex bound in $\mathbb{S}^3$. Regular simplices in $\mathbb{R}^2$ are equilateral triangles, which tessellate. At the vertices of each triangle (of circumradius r) is a circle of radius r . The area of intersection between a triangle and a circle is a circular sector of angle $\pi/3$. Each triangle is covered by three equal areas of intersection. The covering density is therefore the ratio of the sum of the three areas of intersection to the area of the triangle: $\tau_{\mathbb{R}^2} = 2\pi/3(3)^{1/2}$. In $\mathbb{R}^d$ the covering density is independent of r , which is not the case in $\mathbb{S}^d$ for $d \geq 2$ owing to a lack of tessellation.

$$Q = 3 \exp[-2i\psi(\theta)] + 4 \exp[-3i\psi(\theta)] + \exp[-6i\psi(\theta)], \quad (21)$$

$$L = \frac{1}{2} \left(\text{Li}_2(Z_0) + 3 \text{Li}_2\{Z_0 \exp[-4i\psi(\theta)]\} - 4 \text{Li}_2\{-Z_0 \exp[-3i\psi(\theta)]\} - 3\psi(\theta)^2 \right), \quad (22)$$

$$Z_0 = \frac{-6 \sin^2[\psi(\theta)]}{Q} + \frac{2(\{\cos[\psi(\theta)] + 1\}^3 \{1 - 3 \cos[\psi(\theta)]\})^{1/2}}{Q}. \quad (23)$$

where $\psi(\theta)$ is the dihedral angle of $T(\theta)$. The terms in equations (17)–(19) are derived in Appendix A. Equations (20)–(23) are a simplification of Murakami's formula for the volume of a spherical tetrahedron (Murakami, 2012), for the case where all six dihedral angles are equal (a regular spherical tetrahedron).

The covering density can be used to estimate the optimality gap of a point set. For a set of N points with covering radius θ , the lower bound on the covering radius θ^* can be found by rearranging the density expression in equation (6):

$$N = \frac{2\pi^2 \tau_{\mathbb{S}^3}(\theta^*)}{C_3(\theta^*)}, \quad (24)$$

where $2\pi^2$ is the surface area of $\mathbb{S}^3$. The optimality gap of the point set is then $\theta/\theta^* - 1$. Since $\tau_{\mathbb{S}^3}(\theta^*)$ is a nontrivial expression, we find θ^* numerically.

3. Method of orientation set generation

We now describe the method for generating point sets with small covering radii. The direct problem formulation with the application of symmetry is shown in Table 2. This is essentially just a restatement of equations (11) and (13).

The problem of finding optimal spherical coverings is difficult; in addition to being an NP-hard problem (van Emde Boas, 1981), the objective function is non-differentiable, and the 'fitness landscape' is non-convex and has many local minima. One possible solution approach [used by Hardin *et al.* (2017) to generate coverings in $\mathbb{S}^2$] is to use direct search. This

Table 2

Direct model for minimizing the covering radius of a point set in $\mathbb{S}^3$.

The point set $\mathbf{P}$ is composed of a basis set, $\mathbf{B}$, on which a chosen symmetry group, $\mathbf{G}$, acts. The covering radius, θ , is calculated using the Delaunay triangulation of $\mathbf{P}$.

Variables	$\mathbf{B} = \{b_1 \in \mathbb{S}^3, b_2 \in \mathbb{S}^3, \dots\}$
Parameters	$\mathbf{G} = \{g_1 \in \mathbb{S}^3, g_2 \in \mathbb{S}^3, \dots\}$
Minimize	$\theta = \max\{\varphi(\mathbf{t}) \mid \mathbf{t} \in \text{DT}(\mathbf{P})\}$
Subject to	$\mathbf{P} = \{b \otimes g \mid b \in \mathbf{B}, g \in \mathbf{G}\}$

overcomes the non-differentiability of the objective function, but repeated solution from many different starting configurations is required to find the globally optimal configuration. Furthermore, owing to the poor scaling of direct search methods with increasing problem size, this approach is not practical since we wish to create very large orientation sets.

Since it is unlikely that we will find globally optimal solutions for large point sets with direct search, we will instead attempt to find good solutions with an indirect method. We proceed as follows: an initial set of orientations is created by sampling randomly from a uniform distribution on $SO(3)$ (Shoemaker, 1992). The covering radius is then successively reduced, firstly by using gradient descent to find a configuration which is a local minimizer of the Riesz energy. Secondly, a smoothing procedure is used to improve the characteristics of the Delaunay triangulation. Lastly, a local optimization procedure is used to further refine the solution. We present no theoretical basis for the choice of methods or for the order in which the methods are applied. Rather, empirical experimentation has shown that the method is effective and produces point sets with a small covering radius.

The motivation for choosing these methods is illustrated in Fig. 4. The random point set has a large covering radius. By minimizing the Riesz energy the covering radius is significantly reduced. The covering radius can be further reduced as shown in the optimal covering. The effect of the smoothing procedure is not shown here, as it is visually very similar to the Riesz energy and optimal covering configurations. In the rest of this section we describe each method in detail.

3.1. Riesz energy minimization

For a set of points $\mathbf{P} = \{p_1, p_2, \dots, p_N\} \in \mathbb{S}^d$, the Riesz energy is defined as

$$E_s(\mathbf{P}) = \begin{cases} \sum_{i \neq j}^N \frac{1}{|p_i - p_j|^s} & \text{if } s > 0, \\ \sum_{i \neq j}^N \log \frac{1}{|p_i - p_j|} & \text{if } s = 0. \end{cases} \quad (25)$$

The problem of finding optimal Riesz energy configurations is well studied, most commonly for ($d = 3, s = 1$) (also known as the Thomson problem) (Erber & Hockney, 1991; Altschuler *et*

Figure 4

Delaunay triangulations and Voronoi cells of three different point sets in $\mathbb{S}^2$, for $N = 130$. Left: points sampled uniformly from a random distribution. Centre: the global minimum configuration for the Riesz energy, here for $s = 1$ [point set from Wales & Ulker (2006)]. Right: putatively optimal spherical-covering configuration [point set from Hardin *et al.* (2017)]. Point sets in $\mathbb{S}^2$ are used here for illustrative purposes only, because of the difficulty of visualizing $\mathbb{S}^3$.

al., 1994; Wales & Ulker, 2006), but also for ($d = 4, s = 1$) (Altschuler & Pérez-Garrido, 2007), and in the general case (Hardin & Saff, 2004; Rakhmanov *et al.*, 1995). The sphere packing problem is equivalent to solving for $s = \infty$.

Cohn & Kumar (2007) have shown that there exist configurations for certain values of N which are universally optimal, that is, globally optimal solutions for every value of s . The known universally optimal configurations for $d = 3$ are the tetrahedron, the 16-cell and the 600-cell. The vertices of these polyhedra are conjectured to be global optima for the sphere covering problem, since their Delaunay triangulations consist of regular spherical tetrahedra (*cf.* §2.5). However, for any value of N for which a universally optimal configuration does not exist, there is no value of s which for a configuration minimizing $E_s(\mathbf{P})$ guarantees an optimal spherical covering. As such, we will select a value of s on the following basis: Kuijlaars *et al.* (2007) have shown that the set of points $\mathbf{P}$ which minimizes $E_s(\mathbf{P})$ is well distributed when $d - 1 \leq s < d$. We will select $s = 2$ since longer-range potentials exhibit fewer local minima (Wales & Ulker, 2006). We have used the PR+ conjugate gradient method (Wright & Nocedal, 1999) to find a local minimum of $E_s(\mathbf{P})$. The resulting configuration is a good intermediate solution with a small covering radius.

3.2. Optimal Delaunay triangulation smoothing

Minimizing the Riesz energy of a point set reduces the covering radius whilst considering only the relative positions of the points. We can obtain a further reduction in covering radius by considering the positions of a point set *and* the simplices in its Delaunay triangulation. This is a well studied problem in the computational geometry community known as tetrahedral meshing. Given a set of points sampled from an object (*e.g.* a teapot model) the objective is to move the points in order to create a ‘nice’ Delaunay triangulation (the mesh) whilst preserving the shape of the object. Chen (2004) defines an optimal Delaunay triangulation (ODT) as a set of points which minimizes the energy function

$$E_{\text{ODT}} = \frac{1}{d+1} \sum_{i=1 \dots N} \int_{\Omega_i} \|p - p_i\|^2 dp, \quad (26)$$

where Ω_i is the 1-ring of p_i (the volume bounded by p_i and its simplicial neighbours). Minimization of this energy results in a Delaunay triangulation whose simplices have a low circumradius to inradius ratio. Alliez *et al.* (2005) have shown that, for a given point, the position which minimizes E_{ODT} is

$$p_i^* = \frac{1}{\text{Vol}(\Omega_i)} \sum_{\mathbf{t} \in \Omega_i} \text{Vol}(\mathbf{t}) C(\mathbf{t}), \quad (27)$$

where $\text{Vol}(\mathbf{t})$ and $C(\mathbf{t})$ are, respectively, the volume and circumcentre of simplex $\mathbf{t}$. They have shown that the energy can be minimized with guaranteed convergence by alternately constructing the Delaunay triangulation and moving the vertices to their optimal positions using equation (27).

For our applications the ‘object’ whose shape we must preserve is simply $\mathbb{S}^3$. As such, after the optimal vertex posi-

Table 3

Model for reducing the covering radius of a local neighbourhood of a point set.

The point set $\mathbf{P}$ is composed of a basis set, $\mathbf{B}$, on which a chosen symmetry group, $\mathbf{G}$, acts. The basis set, $\mathbf{B}$, consists of an active set, $\mathbf{A}$, which defines the local neighbourhood to be optimized, and a constant set, $\mathbf{C}$, which contains the remaining points. The covering radius, θ , is again calculated using the Delaunay triangulation, though only of the points which either are active or share a simplicial neighbour with an active point.

Variables	$\mathbf{A} = \{a_1 \in \mathbb{S}^3, a_2 \in \mathbb{S}^3, \dots\}$	(1)
Parameters	$\mathbf{C} = \{c_1 \in \mathbb{S}^3, c_2 \in \mathbb{S}^3, \dots\}$	(2)
	$\mathbf{G} = \{g_1 \in \mathbb{S}^3, g_2 \in \mathbb{S}^3, \dots\}$	(3)
Minimize	$\max\{\varphi(\mathbf{t}) \mid \mathbf{t} \in \text{DT}(\mathbf{P}) \wedge \mathbf{t} \cap \mathbf{A} \neq \emptyset\}$	(4)
Subject to	$\mathbf{B} = \mathbf{A} \cup \mathbf{C}$	(5)
	$\mathbf{P} = \{b \otimes g \mid b \in \mathbf{B}, g \in \mathbf{G}\}$	(6)

tion has been calculated using equation (27), the vertex position is normalized in order to bring it back onto $\mathbb{S}^3$. We also calculate $\text{Vol}(\mathbf{t})$ for a spherical tetrahedron (Murakami, 2012) rather than for a Euclidean tetrahedron. Despite the intended use for Euclidean geometries, we have found that this method works very well in practice in $\mathbb{S}^3$, which is probably because of the small local curvature of $\mathbb{S}^3$ for large point sets.

3.3. Local refinement

As a last step in the process of reducing the covering radius, we use an optimization procedure to iteratively refine a succession of local neighbourhoods. We do so by generalizing the direct problem, by iteratively dividing $\mathbf{B}$ into an active set $\mathbf{A}$ and a constant set $\mathbf{C}$. We then minimize the maximum circumradius of the simplices with a vertex in $\mathbf{A}$. A description of the optimization problem is given in Table 3.

Whilst the smallest active set consists of a single vertex, we find that optimizing the vertices of one whole simplex at a time gives better results. To do so, we alternately construct the Delaunay triangulation and then optimize each simplex in turn. The order in which the simplices are optimized is determined by their circumradius, from largest to smallest. After each update the chosen symmetry group is reapplied to the basis set in order to maintain a consistent point set.

Since the minimization of the maximum value of a set is a non-differentiable objective function, we use the Nelder–Mead method (Nelder & Mead, 1965) to optimize the above function as it is a derivative-free method. In order to avoid dealing with the implicit constraint $|p| = 1 \ \forall p \in \mathbf{P}$, we represent the vertices using RF vectors. Representing the vertices as RF vectors during optimization has the added benefit of reducing the number of variables, which is particularly advantageous when using the Nelder–Mead method. Since an RF vector representation of any 180° rotation has infinite magnitude, we rotate the local neighbourhood under consideration to $\{1, 0, 0, 0\}$ prior to optimization, and back again after optimization.

4. Results

Fig. 5 illustrates how each stage of the optimization process affects the solution quality. The initial random sampling

results in a distribution of simplex circumradii that is approximately Gaussian. Minimization of the Riesz energy significantly reduces the mean and variance of the simplex circumradii, as well as the number of simplices. The distribution resembles a bimodal Gaussian distribution, which suggests an ordered underlying simplex structure.

Application of ODT smoothing reduces the mean and variance of the of simplex circumradii and results, again, in an approximately Gaussian distribution. Lastly, the objective of the local refinement procedure is to minimize the maximum simplex circumradius. It can be seen that this produces a peak around the maximum circumradius with a tail of smaller circumradii below this.

Fig. 6 shows a comparison of our method with the methods discussed in §1, in the range $N = [960, 200\,000]$. We have

applied $2I_{60}$ symmetry, as it requires a small basis set and thus allows us to quickly generate coverings of the full space of S^3 . For each value of N , we have applied our method from 200 random starting configurations and taken the point set with the lowest covering radius. It can be seen that the resulting sets have a lower covering radius than the other methods, at both small and large values of N .

Furthermore, our method displays a smooth decrease in covering radius with increasing N , which is highlighted by the almost constant covering density. We do not claim optimality for any of our point sets; in most cases the covering radius of the best point set was unique amongst the 200 runs. As such we can conclude that lower covering radii could be obtained simply by increasing the number of runs, though this is very time consuming for large point sets.

Figure 5
Histograms showing the change in simplex circumradius at each stage in the optimization for a point set with antipodal symmetry and $N = 20\,000$. The histograms show the simplex circumradius distribution after (a) initial random sampling, (b) Riesz energy minimization, (c) ODT smoothing and (d) local refinement. The maximum circumradius is reduced at every stage.

Figure 6
Comparison of the covering radius (left) and the covering density (right) of random sampling from a uniform distribution with antipodal symmetry, incremental grids based on the Hopf fibration (Yershova *et al.*, 2010), cubochoric grids (Rořca *et al.*, 2014), b.c.c. grids with binary octahedral symmetry (Karney, 2007) and our method. For the random sampling, a mean of 10^5 runs was used. For the incremental grids based on the Hopf fibration, the covering radius was calculated at every value of N in the range shown. For the covering radius optimized point sets (our method), the best result of 200 runs was used.

Table 4

Conjectured optimality gaps for covering-radius-optimized configurations, with $2I_{60}$ symmetry applied.

N is the number of points in each set, θ is the covering radius and θ^* is the covering radius of the simplex bound on S^3 , conjectured to be a lower bound (Böröczky, 2004). The optimality gap percentage is $100(\theta/\theta^* - 1)$.

N	θ (°)	θ^* (°)	Optimality gap (%)
8†	60.00	60.00	0.00
120†	22.24	22.24	0.00
1920	9.05	8.73	3.68
3960	7.20	6.85	5.05
6000	6.27	5.96	5.07
7920	5.71	5.44	4.95
9960	5.27	5.04	4.67
12000	4.96	4.73	4.71
13920	4.72	4.50	4.76
15960	4.50	4.30	4.54
18000	4.33	4.13	4.74
19920	4.18	4.00	4.61
24000	3.93	3.76	4.72
27960	3.72	3.57	4.31
31920	3.56	3.41	4.38
36000	3.43	3.28	4.47
39960	3.31	3.17	4.62
43920	3.21	3.07	4.67
48000	3.11	2.98	4.49
60000	2.89	2.77	4.35
79920	2.63	2.51	4.48
99960	2.44	2.33	4.64
139920	2.19	2.09	4.83
180000	2.01	1.92	4.84

† $N = 8$ and $N = 120$ are the point sets containing the vertices of the 16-cell and 600-cell, respectively, included here to highlight the tightness of the simplex bound for point sets consisting of regular tetrahedral cells.

The optimality gaps of some selected point sets generated using our method are shown in Table 4. The gaps are below 6% at every value of N . In the Euclidean limit ($N \rightarrow \infty$) the curvature of S^3 in a local area is effectively zero. For this reason, the optimal covering in a local area should be a b.c.c. lattice, since this is the best known covering in $\mathbb{R}^3$. Since a b.c.c. lattice has a higher covering density than the simplex bound,

Figure 7

Error histogram for a covering-radius-optimized point set with 6×10^6 orientations and a cubochoric set with 6.3×10^6 orientations. Here, both point sets cover the full space of $SO(3)$, which corresponds to indexing a material with a triclinic crystal lattice. The covering-radius-optimized point set has a lower maximum error (1.00 versus 1.72°) and a better overall distribution of errors.

the optimality gaps presented here leave room for improvement.

4.1. Practical application

The results presented in Fig. 6 demonstrate the evolution of the different methods with increasing size, though all at small sizes. For a practical pattern-indexing application, much larger point sets are needed. Furthermore, whilst the covering radius of a set specifies the maximum error, the distribution of errors is also of practical interest. Fig. 7 compares the error histograms of a covering-radius-optimized set and a cubochoric set, which is used for comparison owing to its use in the widely used *EMsoft* microscopy software (De Graef, 2017). In order to generate the error histogram, 10^8 random orientations were sampled; for each sampled orientation, the misorientation is calculated to the nearest orientation in the dictionary set. A KD-tree (Bentley, 1975) is used to quickly find the closest dictionary orientation. In addition to a smaller maximum error, the covering-radius-optimized set has a better overall error distribution. This is achieved despite the use of a smaller number of orientations.

The maximum error of the covering-radius-optimized set is 72% smaller than that of the cubochoric set. In the Euclidean limit $\theta \propto n^{-1/3}$, which suggests that a cubochoric set would require approximately 5 times as many points to achieve the same maximum error.

Using the symmetry relationships described in §2.4, we have created orientation sets for every Laue group with maximum misorientations (2θ) of <1, 2, 3, 4 and 5°, with optimality gaps of less than 6% for every set. The orientation sets are available online (Larsen & Schmidt, 2017).

5. Summary

We have shown how to construct a near-optimal sampling of orientations. First we demonstrated that the sampling problem is equivalent to the problem of how to distribute points on a hypersphere. We then showed that the best measure of quality for a point set is the covering radius, as this determines the maximum orientational error. With the minimum covering radius as the objective, we created sets of orientations at a range of sizes for use in pattern indexing, and demonstrated that the number of orientation samples needed to achieve a desired indexing accuracy is significantly reduced as a consequence.

In addition to an exact calculation of the covering radius, which measures the quality of a set of orientations (smaller is better), we derived a lower bound on the covering radius, which sets a limit on the best attainable quality. The difference between the achieved covering radius and the theoretical limit allows us to quantify the optimality of orientation sets, which we used to show that the sets we created are within 6% of the optimal covering radius.

In order to use the method for indexing of diffraction patterns, we have shown how symmetry groups can be imposed during orientation sampling, without introducing any

edge-effect artifacts. Using this approach we have demonstrated how to sample from the crystallographic fundamental zone of any of the 11 Laue groups.

Existing methods for sampling orientations have prioritized properties such as a refinable grid-like structure, fast generation and the ability expand into spherical harmonics; we have instead chosen to optimize the maximum error (the covering radius) above all else. This also means that the sampling method has very high computational requirements; the largest point set requires approximately four days of computation time. Nonetheless, we claim that this is a good trade-off, since a point set must only be generated once for each desired error level, and affords a significant performance improvement every time a pattern is subsequently indexed.

APPENDIX A

Simplex bound derivation

A1. Volume of a hyperspherical cap

The volume of a hyperspherical cap in $\mathbb{S}^d$ can be calculated by projection into RF space. Since RF space is radially symmetric about the origin, the projection of a hyperspherical cap with radius θ and centre coordinates $\{1, 0, 0, \dots, 0\} \in \mathbb{S}^d$ is a sphere with radius $r = \tan(\theta)$ centred at the origin. Thus, the volume of the cap is the radial integral of the product of the surface area of a $(d-1)$ -sphere with the RF space density:

$$C_d(\theta) = \int_0^{\tan(\theta)} \frac{S_{d-1}(r)}{(1+r^2)^2} dr \quad \text{where} \quad S_{d-1}(\theta) = \frac{d\pi^{d/2}}{\Gamma(d/2+1)} \theta^{d-1}. \quad (28)$$

For a hyperspherical cap in $\mathbb{S}^3$, this gives

$$C_3(\theta) = \int_0^{\tan(\theta)} \frac{4\pi r^2}{(1+r^2)^2} dr = \pi[2\theta - \sin(2\theta)]. \quad (29)$$

This is the same as the result derived by Morawiec (2004, 2010), but without normalization.

A2. Edge length of a regular spherical tetrahedron in $\mathbb{S}^3$

Owing to the radial symmetry of RF space, the RF projection of a regular spherical tetrahedron with centre coordinates $q_c = \{1, 0, 0, 0\} \in \mathbb{S}^3$ is a tetrahedron with centre coordinates $v_c = \{0, 0, 0\}$ and vertex coordinates

$$\begin{aligned} v_1 &= \{k, k, k\}, & v_2 &= \{k, -k, -k\}, \\ v_3 &= \{-k, k, -k\}, & v_4 &= \{-k, -k, k\}. \end{aligned} \quad (30)$$

From this, we obtain the vertex coordinates in $\mathbb{S}^3$:

$$\begin{aligned} q_1 &= \frac{1}{(1+3k^2)^{1/2}} \{1, k, k, k\}, \\ q_2 &= \frac{1}{(1+3k^2)^{1/2}} \{1, k, -k, -k\}, \\ q_3 &= \frac{1}{(1+3k^2)^{1/2}} \{1, -k, k, -k\}, \\ q_4 &= \frac{1}{(1+3k^2)^{1/2}} \{1, -k, -k, k\}. \end{aligned} \quad (31)$$

The circumradius of the tetrahedron is given by the arc length from the centre to any of the vertices:

$$\theta = \arccos\langle q_c, q_i \rangle = \arccos\left[\frac{1}{(1+3k^2)^{1/2}}\right] \quad \forall i. \quad (32)$$

The edge length of the tetrahedron is the arc length between any two vertices:

$$l = \arccos\langle q_i, q_j \rangle = \arccos\left(\frac{1-k^2}{1+3k^2}\right) \quad \forall i \neq j. \quad (33)$$

Using equations (32) and (33) we can express the edge length in terms of the radius:

$$l = \arccos\left[\frac{4\cos^2(\theta) - 1}{3}\right]. \quad (34)$$

A3. Dihedral angle and solid angle of intersection

Let $\{q_1, q_2, q_3, q_4\}$ be the vertices of a regular hyperspherical simplex in $\mathbb{S}^3$ with the following coordinates:

$$\begin{aligned} q_1 &= \{1, 0, 0, 0\}, & q_2 &= \left\{\cos(l), -a, -\frac{a}{3^{1/2}}, z\right\}, \\ q_3 &= \left\{\cos(l), a, -\frac{a}{3^{1/2}}, z\right\}, & q_4 &= \left\{\cos(l), 0, -\frac{2a}{3^{1/2}}, z\right\}, \end{aligned} \quad (35)$$

where

$$a = \left[\frac{1 - \cos(l)}{2}\right]^{1/2}, \quad z = \left\{\sin^2(l) - \frac{2}{3}[1 - \cos(l)]\right\}^{1/2}. \quad (36)$$

When projected into RF space the tetrahedron has vertices

$$\begin{aligned} v_1 &= \{0, 0, 0\}, & v_2 &= \frac{1}{\cos(l)} \left\{-a, -\frac{a}{3^{1/2}}, z\right\}, \\ v_3 &= \frac{1}{\cos(l)} \left\{a, -\frac{a}{3^{1/2}}, z\right\}, & v_4 &= \frac{1}{\cos(l)} \left\{0, -\frac{2a}{3^{1/2}}, z\right\}. \end{aligned} \quad (37)$$

The dihedral angle of the tetrahedron is then given by

$$\psi(l) = \arccos\left\langle \frac{v_2 \times v_3}{|v_2 \times v_3|}, \frac{v_2 \times v_4}{|v_2 \times v_4|} \right\rangle = \arccos\left[\frac{\cos(l)}{2\cos(l) + 1}\right]. \quad (38)$$

Using equation (34) we can express the dihedral angle in terms of θ :

$$\psi(\theta) = \arccos\left[\frac{4\cos^2(\theta) - 1}{8\cos^2(\theta) + 1}\right]. \quad (39)$$

Table 5

Generators for the seven Laue groups which are subsets of O .

%	O	T	D_4	D_2	C_4	C_2	C_1
{1, 0, 0, 0}	✓	✓	✓	✓	✓	✓	✓
{0, 0, 0, 1}	✓	✓	✓	✓	✓	✓	
{0, 1, 0, 0}	✓	✓	✓	✓			
{0, 0, 1, 0}	✓	✓	✓	✓			
$\{2^{1/2}/2, 0, 0, 2^{1/2}/2\}$	✓		✓		✓		
$\{2^{1/2}/2, 0, 0, -2^{1/2}/2\}$	✓		✓		✓		
$\{0, 2^{1/2}/2, 2^{1/2}/2, 0\}$	✓		✓				
$\{0, -2^{1/2}/2, 2^{1/2}/2, 0\}$	✓		✓				
$\{1/2, 1/2, -1/2, 1/2\}$	✓	✓					
$\{1/2, 1/2, 1/2, -1/2\}$	✓	✓					
$\{1/2, 1/2, -1/2, -1/2\}$	✓	✓					
$\{1/2, -1/2, -1/2, -1/2\}$	✓	✓					
$\{1/2, -1/2, 1/2, 1/2\}$	✓	✓					
$\{1/2, -1/2, 1/2, -1/2\}$	✓	✓					
$\{1/2, -1/2, -1/2, 1/2\}$	✓	✓					
$\{1/2, 1/2, 1/2, 1/2\}$	✓	✓					
$\{2^{1/2}/2, 2^{1/2}/2, 0, 0\}$	✓						
$\{2^{1/2}/2, -2^{1/2}/2, 0, 0\}$	✓						
$\{2^{1/2}/2, 0, 2^{1/2}/2, 0\}$	✓						
$\{2^{1/2}/2, 0, -2^{1/2}/2, 0\}$	✓						
$\{0, 2^{1/2}/2, 0, 2^{1/2}/2\}$	✓						
$\{0, -2^{1/2}/2, 0, 2^{1/2}/2\}$	✓						
$\{0, 0, 2^{1/2}/2, 2^{1/2}/2\}$	✓						
$\{0, 0, -2^{1/2}/2, 2^{1/2}/2\}$	✓						

Table 6

Generators for the five Laue groups which are subsets of D_6 .

%	D_6	D_3	C_6	C_3	C_1
{1, 0, 0, 0}	✓	✓	✓	✓	✓
$\{1/2, 0, 0, 3^{1/2}/2\}$	✓	✓	✓	✓	
$\{1/2, 0, 0, -3^{1/2}/2\}$	✓	✓	✓	✓	
{0, 0, 0, 1}	✓		✓		
$\{3^{1/2}/2, 0, 0, 1/2\}$	✓		✓		
$\{3^{1/2}/2, 0, 0, -1/2\}$	✓		✓		
{0, 1, 0, 0}	✓	✓			
$\{0, -1/2, 3^{1/2}/2, 0\}$	✓	✓			
$\{0, 1/2, 3^{1/2}/2, 0\}$	✓	✓			
$\{0, 3^{1/2}/2, 1/2, 0\}$	✓				
$\{0, -3^{1/2}/2, 1/2, 0\}$	✓				
{0, 0, 1, 0}	✓				

The solid angle is then given by

$$\Omega(\theta) = 3\psi(\theta) - \pi = 3 \arccos \left[\frac{4 \cos^2(\theta) - 1}{8 \cos^2(\theta) + 1} \right] - \pi. \quad (40)$$

Since v_1 lies at the origin, this is also the solid angle of intersection of a regular hyperspherical simplex and a hyperspherical cap placed at one of its vertices. We can verify that in the Euclidean limit (where the curvature is zero) $\lim_{\theta \rightarrow 0} \Omega(\theta) = 3 \arccos(1/3) - \pi = \arccos(23/27)$, which is the solid angle for a regular tetrahedron in $\mathbb{R}^3$, and that $\Omega(\pi/3) = \pi/2$, which is the solid angle of a tetrahedral cell in the 16-cell.

APPENDIX B

Laue group subset relationships

The subset relationships between the 11 Laue groups are shown in Tables 5 and 6.

Acknowledgements

The authors thank Thomas J. Hardin for advice on spherical and hyperspherical harmonics, Farangis Ram for discussions on the cubochoric method, Nanna Wahlberg, Erik B. Knudsen and Hugh Simons for proofreading and helpful suggestions on the manuscript, and an anonymous referee for several suggestions which have improved this work.

References

- Adams, B. L., Wright, S. I. & Kunze, K. (1993). *Metall. Trans. A*, **24**, 819–831.
- Alliez, P., Cohen-Steiner, D., Yvinec, M. & Desbrun, M. (2005). *ACM Trans. Graph.* **24**, 617–625.
- Altmann, S. L. (2005). *Rotations, Quaternions and Double Groups*. Mineola: Dover Publications.
- Altschuler, E. L. & Pérez-Garrido, A. (2007). *Phys. Rev. E*, **76**, 016705.
- Altschuler, E. L., Williams, T. J., Ratner, E. R., Dowla, F. & Wooten, F. (1994). *Phys. Rev. Lett.* **72**, 2671–2674.
- Bentley, J. L. (1975). *Commun. ACM*, **18**, 509–517.
- Böröczky, K. (1978). *Acta Math. Acad. Sci. Hung.* **32**, 243–261.
- Böröczky, K. (2004). *Finite Packing and Covering*, Vol. 154. Cambridge University Press.
- Brown, R. & Gray, A. (1967). *Comment. Math. Helv.* **42**, 222–236.
- Chen, L. (2004). *Mesh Smoothing Schemes Based on Optimal Delaunay Triangulation, Proceedings of the 13th International Meshing Roundtable*, pp. 109–120. <http://imr.sandia.gov/papers/abstracts/Ch317.html>.
- Chen, Y. H., Park, S. U., Wei, D., Newstadt, G., Jackson, M. A., Simmons, J. P., De Graef, M. & Hero, A. O. (2015). *Microsc. Microanal.* **21**, 739–752.
- Cohn, H. & Kumar, A. (2007). *J. Am. Math. Soc.* **20**, 99–149.
- Conway, J. H. & Sloane, N. J. A. (1998). *Sphere Packings, Lattices and Groups*. New York: Springer.
- Conway, J. H. & Smith, D. A. (2003). *On Quaternions and Octonions*. New York: Springer.
- Coxeter, H., Few, L. & Rogers, C. (1959). *Mathematika*, **6**, 147–157.
- Dahms, M. & Bunge, H. J. (1989). *J. Appl. Cryst.* **22**, 439–447.
- De Graef, M. (2017). *EMsoft – Electron Microscopy Diffraction and Image Computations*. <http://github.com/marcdegraef/EMsoft/>.
- Emde Boas, P. van (1981). *Another NP-Complete Partition Problem and the Complexity of Computing Short Vectors in a Lattice*. Report 81-04. Mathematisch Instituut, Universiteit van Amsterdam, The Netherlands.
- Erber, T. & Hockney, G. (1991). *J. Phys. A Math. Gen.* **24**, L1369–L1377.
- Frank, F. (1988). *Metall. Trans. A*, **19**, 403–408.
- Hardin, D. & Saff, E. (2004). *Not. AMS*, **51**, 1186–1194.
- Hardin, R. H., Sloane, N. J. A. & Smith, W. D. (2017). *Spherical Codes*. Book in preparation. <http://neilsloane.com/coverings/>.
- He, Y. & Jonas, J. J. (2007). *J. Appl. Cryst.* **40**, 559–569.
- Karney, C. F. F. (2007). *J. Mol. Graph. Model.* **25**, 595–604.

- Kuijlaars, A. B. J., Saff, E. B. & Sun, X. (2007). *J. Comput. Appl. Math.* **199**, 172–180.
- Larsen, P. M. & Schmidt, S. (2017). *Orientation Sets Created via Solution of the Hyperspherical Covering Problem*, <http://github.com/pmla/hyperspherical-coverings>.
- Li, S. F. & Suter, R. M. (2013). *J. Appl. Cryst.* **46**, 512–524.
- Mason, J. K. (2009). *Acta Cryst.* **A65**, 259–266.
- Mason, J. K. & Schuh, C. A. (2008). *Acta Mater.* **56**, 6141–6155.
- Morawiec, A. (2004). *Orientations and Rotations: Computations in Crystallographic Textures*. Berlin, Heidelberg: Springer-Verlag.
- Morawiec, A. (2010). *Acta Cryst.* **A66**, 717–719.
- Morawiec, A. & Field, D. (1996). *Philos. Mag. A*, **73**, 1113–1130.
- Murakami, J. (2012). *Proc. Am. Math. Soc.* **140**, 3289–3295.
- Nelder, J. A. & Mead, R. (1965). *Comput. J.* **7**, 308–313.
- Poulsen, H. F., Nielsen, S. F., Lauridsen, E. M., Schmidt, S., Suter, R. M., Lienert, U., Margulies, L., Lorentzen, T. & Juul Jensen, D. (2001). *J. Appl. Cryst.* **34**, 751–756.
- Rakhmanov, E. A., Saff, E. B. & Zhou, Y. M. (1995). *Comput. Methods Funct. Theory*, **5**, 293–309.
- Ram, F., Wright, S., Singh, S. & De Graef, M. (2017). *Ultramicroscopy*, **181**, 17–26.
- Rogers, C. A. (1958). *Proc. London Math. Soc.* **s3-8**, 609–620.
- Roşca, D., Morawiec, A. & De Graef, M. (2014). *Modell. Simul. Mater. Sci. Eng.* **22**, 075013.
- Schmidt, S. (2014). *J. Appl. Cryst.* **47**, 276–284.
- Shoemake, K. (1992). *Graphics Gems III*, pp. 124–132. Academic Press.
- Singh, S. & De Graef, M. (2017). *Microsc. Microanal.* **23**, 1–10.
- Tóth, L. F. (1964). *Regular Figures*. New York, Paris: Elsevier.
- Wales, D. J. & Ulker, S. (2006). *Phys. Rev. B*, **74**, 212101.
- Wright, S. & Nocedal, J. (1999). *Numerical Optimization*, Vol. 2. New York: Springer.
- Yershova, A., Jain, S., LaValle, S. M. & Mitchell, J. C. (2010). *Int. J. Robot. Res.* **29**, 801–812.