

The structure of the cubic coincident site lattice rotation group

Bryan W. Reed,* Roger W. Minich, Robert E. Rudd and Mukul Kumar

Received 10 February 2004

Accepted 30 March 2004

Lawrence Livermore National Laboratory, 7000 East Ave, Livermore, CA 94550-9234, USA.

Correspondence e-mail: reed12@llnl.gov

This work is intended to be a mathematical underpinning for the field of grain-boundary engineering and its relatives. The inter-relationships within the set of rotations producing coincident site lattices in cubic crystals are examined in detail. Besides combining previously established but widely scattered results into a unified context, the present work details newly developed representations of the group structure in terms of strings of generators (based on quaternionic number theory, and including uniqueness proofs and rules for algebraic manipulation) as well as an easily visualized topological network model. Important results that were previously obscure or not universally understood (*e.g.* the Σ combination rule governing triple junctions) are clarified in these frameworks. The methods also facilitate several general observations, including the very different natures of twin-limited structures in two and three dimensions, the inadequacy of the Σ combination rule to determine valid quadruple nodes, and a curious link between allowable grain-boundary assignments and the four-color map theorem. This kind of understanding is essential to the generation of realistic statistical models of grain-boundary networks (particularly in twin-dominated systems) and is especially applicable to the field of grain-boundary engineering.

1. Introduction

1.1. Motivation

The field of grain-boundary engineering aims to improve material properties by controlling the types of grain boundaries in a network (*i.e.* to increase the population of boundaries with desirable characteristics) as well as the topological properties of the networks themselves (*e.g.* reducing the likelihood of a large continuous path of weak boundaries). Significant strides have been made in this field in recent years (*e.g.* Gertsman *et al.*, 1994; Randle, 1996; Kumar *et al.*, 2000; Schuh *et al.*, 2002), resulting in better understanding and control of brittle fracture strength and related phenomena in polycrystalline materials. Further progress demands highly statistically realistic models of grain-boundary networks, as it is becoming clear that what may at first glance appear to be subtle and mathematically obscure sources of correlation in fact have large effects on the real properties of materials (Palumbo *et al.*, 1992; Gertsman & Szpunar, 1998; Kumar *et al.*, 2000; Minich *et al.*, 2002). Not just the populations of certain types of grain boundaries are important; distributions of triple-junction types and longer-range correlations induced by crystallographic consistency requirements need to be considered as well.

In this work, we will consider the inter-relationships among the various rotations that produce coincident-site-lattice

(CSL) misorientations (*i.e.* rotations that map a fraction $1/\Sigma$ of the lattice points from one crystal grain onto lattice points of another, usually adjacent, grain) (Mykura, 1979). These misorientations make up a subgroup of the usual rotation group in three dimensions [$SO(3)$], with the constraint that the rotation-matrix elements are all rational numbers. It is known that, particularly in the case of twin-dominated structures in cubic crystals (containing a large number of boundaries with $\Sigma = 3^n$), neglecting the interconnectivity of this group of rotations can lead to statistically skewed unrealistic results (see the references in the preceding paragraph). For example, the often-discussed pure $\Sigma 3/\Sigma 9$ twin-limited structure (consisting entirely of $\Sigma 3$ and $\Sigma 9$ boundaries in a 2:1 ratio) (Palumbo *et al.*, 1992; Miyazawa *et al.*, 1996; Gertsman, 2001*b*) is easy to construct in two dimensions but does not exist for most three-dimensional networks unless either the $\Sigma 3/\Sigma 9$ ratio is significantly reduced or low-angle $\Sigma 1$ boundaries are brought into the network (Gertsman, 2001*b*) (and, in the latter case, the fraction of $\Sigma 3$ boundaries can exceed $2/3$). Further, the frequently simulated two-dimensional hexagonal and three-dimensional tetrakaidecahedral arrays, while admittedly convenient, house subtle biases owing to their unusual colorability properties. For example, a network of all $\Sigma 9$ boundaries is possible in the tetrakaidecahedral array (as well as for any two-dimensional structure, *via* the four-color map theorem), but for the vast majority of three-dimensional networks no

such structure is possible. At the other extreme, a common misunderstanding of the CSL inter-relationships would lead one to believe that a network of all $\Sigma 9$ boundaries is impossible even in two dimensions, which is assuredly not the case (Gertsman & Tangri, 1995). These examples are developed to illustrate the power of the formalism presented in this paper – once the methods are developed, the demonstration of each of the above assertions simply amounts to drawing and making observations on a small number of graphs. Implicit in these graphs is the entire structure of the CSL rotation group and the manner in which this group governs the network of grain boundaries.

Developing these ideas rigorously requires a significant foray into group theory and number theory, but the reward at the end is that, once our models are developed, we no longer have to think in terms of Euler angles, integer quaternions, rational matrices or axis-angle pairs. Everything about the interconnectivity of the group can be calculated in terms of simple algebraic operations on strings and mappings of the nodes of one graph to the nodes of another. In short, we are trading algebra for geometry, which has proven time and again to be a very powerful trade both in modern math and in physics.

1.2. Overview

The notion of a CSL misorientation has been widely used in discussions of grain-boundary character (Randle, 1996). In many materials, certain low Σ numbers correlate with so-called special boundaries, which may have different strengths, chemical resistances or impurity segregation properties to non-special (or random) boundaries (Randle, 1994, 1996, and references therein). The CSL model has its limitations, notably in its neglect of the grain-boundary plane orientation, but has proven useful enough to serve as a primary means of categorizing boundary types in both theoretical and experimental investigations. A large literature has grown up around this concept since its introduction, particularly in cubic crystals (which will occupy our attention in this work), although the CSL concept may be applied to other crystal systems (Bollmann, 1972; Iwasaki, 1976; MacLaren & Aindow, 1997; Gertsman, 2001a). The application of the properties of this group to twin-dominated systems in particular has seen considerable attention in recent years (e.g. Kopecky *et al.*, 1991; Fortier *et al.*, 1995; Gertsman & Tangri, 1995; Gertsman & Szpunar, 1998; Kumar *et al.*, 2000; Gertsman, 2001a,b, 2002; Schuh *et al.*, 2002, 2003). The present work owes much to these endeavours.

In spite of broad interest, there remains much basic confusion as to the mathematical nature of this group of rotations. Even so fundamental a result as the Σ combination rule governing the set of misorientations around a triple junction (Miyazawa *et al.*, 1996; Gertsman, 2001a) was only clearly identified fairly recently and is still frequently misunderstood, despite clear mathematical proofs and frequent reminders in the literature as to the true nature of the rule. The aim of the current work is to develop the structure of the

CSL rotation group in terms of a network topology and a non-commutative algebraic representation that, once understood, make results such as the Σ combination rule more intuitively clear. By the end, we will be able to determine what triple junction (and quadruple node) assignments are consistent with a given grain topology merely by sketching a few graphs. Complex multiple intersections, which may occur with significant frequency in some structures (Kopecky *et al.*, 1991), can be handled within the same formalism. Specific cases such as the (often neglected) existence of $\Sigma 9$ – $\Sigma 9$ – $\Sigma 9$ triple junctions and the non-existence of $\Sigma 9$ – $\Sigma 27a$ – $\Sigma 27a$ triple junctions will be easily derived by checking the relevant graphs. Any CSL orientation may be expressed as a product of elementary prime- Σ operations, then manipulated using algebraic rules. We show how to extend the formalism beyond the triple junctions, to quadruple nodes and general topologies, and demonstrate how fundamental dimensionality effects may be clarified in the present approach. For instance, the notion of a twin-limited structure is significantly different in two and three dimensions in ways we can begin to quantify simply by understanding the structure of the CSL rotation group. Since the structure of this group is fundamental to grain-boundary engineering, this type of understanding has significant practical implications.

In order to reach this point, we will begin with definitions, followed by a brief discussion of the importance of the $\Sigma 1$ cubic symmetry group. This will allow us to develop the algebraic representation and then the network topology of the CSL rotation group, using examples starting with groups of the type Σp^ω [with p an odd prime number and ω an exponent ranging over all the non-negative integers; this is not to be confused with Σp^n with a specific integer n , which represents a finite subset (but not a subgroup) of the group] but eventually generalizing to Σ values with more than one prime factor. After considering some mathematical details, we will show a simple procedure for drawing a map of any relevant part of the CSL rotation group and applying this map to a set of crystalline grains, thereby determining all of the boundary and junction types for a given orientation assignment. We will include specific examples drawn from the highly important $\Sigma 3^\omega$ ‘twin-related’ group and close with discussions of current limitations and possible future developments.

In summary, the logical progression is as follows: The geometry of grain orientation and grain-boundary misorientation is encoded in the CSL group; quaternions convert the group to an algebra; number theory factorizes quaternions into strings; the units of the string are mapped to links on a graph; a simple graph-theory model is used to understand the Σ combination rule and other non-intuitive results of grain-boundary inter-relationships.

2. Definitions and symmetry

2.1. Basic definitions

This discussion will be limited to lattices with cubic symmetry. The material in this section is well covered in the

literature (Grimmer *et al.*, 1974; Grimmer, 1974; Mykura, 1979; Grimmer, 1984), so we will summarize results without proof, however the results and definitions are quite scattered and notations vary widely, so this section is provided in the interests of clarity. A CSL rotation may be represented by a 3×3 matrix of rational numbers with least common denominator Σ . Σ must be an odd integer. Since this matrix is to represent a rotation, it must be orthogonal (its transpose equals its inverse) and of determinant 1 (which removes the possibility of reflection-rotations). We will use the convention that, if we start with the crystal aligned with the reference axes, the columns of the matrix give the unit vectors of the rotated crystal's principal axes as expressed in the reference coordinate system. Then right-multiplication of a rotation matrix by an element of the proper cubic point group (which we will call $\Sigma 1$, as discussed below) does not change the resulting crystal orientation, and we can think of a rotation as acting to the right on a grain orientation. The same is not in general true for left-multiplication by $\Sigma 1$. Both this convention and the reverse (in which all matrices are inverted, their order of multiplication reversed, and the matrices can be considered to act on the coordinate axes rather than the crystal grains) are in common use in the literature (Goldstein, 1950). The distinction is important, as the formulas for calculating misorientation are different in the different conventions, as we will see.

Besides a matrix, a CSL rotation also may be represented as a quaternion of integers (Grimmer, 1974), which, when normalized to unit magnitude, may be expressed in terms of an axis-angle pair $\theta \mathbf{n}$ (\mathbf{n} being a unit vector specifying the axis and θ the rotation angle) as:

$$q/|q| = [\cos(\theta/2) \quad n_1 \sin(\theta/2) \quad n_2 \sin(\theta/2) \quad n_3 \sin(\theta/2)]. \quad (1)$$

When the quaternion is expressed in lowest terms, the sum of the squares of the integers (which we will call the squared norm, although it is often called simply the norm in the mathematical literature) will equal one of Σ , 2Σ or 4Σ . A quaternion with a prime squared norm will be termed prime. The inverse rotation is given by reversing the sign of either the first or the last three components (the overall sign being meaningless – we will implicitly neglect the overall sign of a quaternion throughout this work, and indeed any scalar multiple of a quaternion will yield the same rotation since we normalize it to unit magnitude before converting to a matrix). We will use either the matrix or the quaternion representation according to convenience, and for simplicity will often speak of a rotation and its matrix or quaternion representations as being one and the same. A distinction must sometimes be made, since matrix left-multiplication is equivalent to quaternion right-multiplication. By default, we will use the multiplication order appropriate for matrices, noting exceptions where they arise.

These definitions suffice for us to construct the group $\Sigma 1$, being the set of CSL rotations with $\Sigma = 1$. We simply generate all column permutations (a factor of $3! = 6$) and sign reversals (a factor of $2^3 = 8$) for the three columns of the 3×3 identity

matrix, resulting in 48 matrices, half of which will have a determinant of +1. $\Sigma 1$ thus has 24 elements and is nothing other than the cubic point group with the rotation-reflections removed. This is not surprising, since each $\Sigma 1$ operation by definition maps a fraction $1/\Sigma = 1$ of the original lattice points onto lattice points. Therefore, as far as the cubic Bravais lattice is concerned, any element of $\Sigma 1$ is equivalent to a null operation.

Eventually, we will show how to factorize any CSL rotation quaternion into a product of prime quaternions and a single (physically meaningless) element of $\Sigma 1$. In order to develop this notion, we shall use a hierarchy of equivalence classes based on the $\Sigma 1$ symmetry group. These ideas of equivalence will enable us to introduce a set of arbitrary conventions (without loss of generality in any physically meaningful sense), which will allow us to assert the uniqueness of the factorization.

The structures of the coincident site lattices themselves will not be discussed here, but are well described in the literature (*e.g.* Grimmer, 1974, 1976). Instead we are only discussing the rotations that produce CSL's and their inter-relationships. The interesting and important theory of triple-junction CSL's (Gertsman, 2001*a,b*) is similarly beyond the current scope.

2.2. Hierarchical categorization of CSL rotations

Since $\Sigma 1$ contains only symmetries of the crystal, it follows that, if we start in the natural coordinate system of a cubic crystal, performing a $\Sigma 1$ operation before any other rotation R is equivalent to rotating by R alone. The set of rotations thus naturally breaks into cosets $R\Sigma 1$ of rotations equivalent to a given R . As a matter of notation, right-multiplying a matrix R by a set of matrices such as $\Sigma 1$ simply means forming all matrices of the form RS with S being any element of $\Sigma 1$, and similarly for left-multiplying. A capital S symbol will always represent an arbitrary element of $\Sigma 1$ in this work. Each coset represents a grain orientation with respect to a reference grain, and we are free to choose any of the 24 elements RS of the coset to represent the rotation. It is an elementary result of group theory that the cosets do not overlap, are all of the same size and cover the entire CSL group.

If, after a rotation R , we perform a $\Sigma 1$ rotation in the frame of the reference grain, then this amounts to rotating by the same angle but about a different (but symmetrically equivalent) axis. For example, if R rotates by 180° about $[110]$, there are S_1 and S_2 in $\Sigma 1$ such that S_1RS_2 is a rotation of 180° about $[011]$. The set of matrices $\Sigma 1R\Sigma 1$, which will have as many as 24^2 elements (possibly fewer depending on symmetries in R), therefore all represent the same 'type' of CSL misorientation. We will call the equivalence class $\Sigma 1R\Sigma 1$ a subtype, to distinguish it from the conventional type defined below. Since two grain orientations are represented by cosets $R_1\Sigma 1$ and $R_2\Sigma 1$, the misorientation between the two grains is represented by the rotations $(R_1\Sigma 1)^{-1}(R_2\Sigma 1) = \Sigma 1R_1^{-1}R_2\Sigma 1$ in the frame of grain 1 (the set $\Sigma 1$, being a group, is the same set as its elementwise inverse). The choice of which elements of $\Sigma 1$ are used in the representations of the two grains is entirely

Table 1

Hierarchical categorization of the CSL rotation group, a subgroup of $SO(3)$.

Each row consists of a union of sets from the row above. In addition, the types may be categorized by symmetry class (Table 2), but not in a manner that fits into this hierarchy. Types and subtypes will be identical unless the type has 48×24 elements [putting it into the (a, b, c, d) class]. The product over p is for all prime factors p of n .

Term	Formula	Significance	Number of elements
Element	R	A single rotation	1
Coset	$R\Sigma 1$	A grain orientation	24
Subtype	$\Sigma 1R\Sigma 1$	Misorientation of an ordered pair of grains	$(1, 4, 6, 8, 12, \text{ or } 24) \times 24$
Type	$(\Sigma 1R\Sigma 1) \cup (\Sigma 1R^{-1}\Sigma 1)$	Misorientation of an unordered pair of grains	$(1, 4, 6, 8, 12, 24, \text{ or } 48) \times 24$
Σn	Σn (n must be odd)	CSL rotations with a fixed ratio n of unit-cell sizes	$m(n) = 24n \prod_p (p+1)/p$
Σn^{op}	$\Sigma 1 \cup \Sigma n \cup \Sigma n^2 \cup \Sigma n^3 \cup \dots$	A group consisting of repeated applications of elements of Σn	∞ (if $n > 1$, odd), 24 (if $n = 1$)
CSL group	$\bigcup_{k=0}^{\infty} \Sigma(2k+1)$	All 3×3 rotation matrices with all rational elements	∞

Table 2

Construction of representatives d_i^p in quaternion form for each of the possible classes of types.

a, b, c and d are distinct positive integers, and the sum of the squares of the elements of each quaternion is equal to $\Sigma = p$, an odd prime. The construction ensures that (i) each coset is represented exactly once, and (ii) the inverse of a representative rotation is also a representative rotation. 180° rotations (with first element equal to zero) are chosen where possible, so that each rotation is its own inverse in these cases.

Class	Number of cosets	Representatives
[1000]	1 ($\Sigma 1$)	[1000]
[1110]	4 ($\Sigma 3$)	[0111] [0 $\bar{1}\bar{1}$ 1] [01 $\bar{1}\bar{1}$] [0 $\bar{1}\bar{1}$ 1]
[ab00]	6	[0 $\pm a$ b 0] [0 0 $\pm a$ b] [0 b 0 $\pm a$]
[aaab]	8	[$\pm b$ $\pm a$ $\pm a$ a]
[aab0]	12	[0 $\pm a$ $\pm a$ b] [0 $\pm a$ b $\pm a$] [0 b $\pm a$ $\pm a$]
[abc0]	24	[0 $\pm a$ $\pm b$ c] [0 $\pm b$ $\pm a$ c] [0 $\pm a$ c $\pm b$] [0 $\pm b$ c $\pm a$] [0 c $\pm a$ $\pm b$] [0 c $\pm b$ $\pm a$]
[aabc]	24	[$\pm a$ $\pm a$ $\pm b$ c] [$\pm a$ $\pm b$ c $\pm a$] [$\pm a$ c $\pm a$ $\pm b$]
[abcd]	48	[$\pm a$ $\pm b$ $\pm c$ d] [$\pm a$ $\pm c$ $\pm b$ d] [$\pm a$ $\pm b$ d $\pm c$] [$\pm a$ $\pm c$ d $\pm b$] [$\pm a$ d $\pm b$ $\pm c$] [$\pm a$ d $\pm c$ $\pm b$]

arbitrary and physically meaningless, so that any element of the same subtype as $R_1^{-1}R_2$ would serve equally well as a representation. Thus this notion of ‘subtype’ is the appropriate category for describing a misorientation from one grain to another, in the reference frame of the first grain. The inverse subtype $\Sigma 1R_2^{-1}R_1\Sigma 1$, which may be different but will have the same Σ , represents the misorientation going from grain 2 to grain 1.

We note in passing that the product $R_2R_1^{-1}$ gives the misorientation rotation in the frame of the reference grain, which is in general of a different subtype (often even a different Σ) to the misorientation as expressed in the frame of either grain. Thus, these matrices may be useful for calculation but are irrelevant in determining the ‘specialness’ of a given grain boundary. If we had used the reverse of our rotation convention (see §2.1), then the roles of $R_1^{-1}R_2$ and $R_2R_1^{-1}$ would be reversed. Thus it is important to clarify which convention is being used. Unfortunately, this is not universally recognized and ambiguous applications of these matrices are easy to find in the literature.

It is customary (Mykura, 1979) to categorize CSL rotations according to what we will term ‘types’ of the form $(\Sigma 1R\Sigma 1) \cup (\Sigma 1R^{-1}\Sigma 1)$ (with as many as 48×24 elements)

and to label the types with a lower-case letter, e.g. the set $\Sigma 27a$ represents one such type of the rational rotation matrices with common denominator 27. The letter is omitted when a given Σ only has one type. Letters are assigned in order of minimum misorientation angle, with ties broken according to the sum of the indices of the minimum-angle rotation axis (h, k, l) , given in lowest terms (Mykura, 1979). Unfortunately, this criterion is still incomplete as it fails to provide an ordering for, e.g., the $\Sigma 229$ types given by quaternions [14 5 2 2] and [14 4 4 1]. If we break ties by first sorting in order of $h^2+k^2+l^2$ and then by $h^3+k^3+l^3$, the ambiguity problem is solved without creating the need to alter any existing tables. The types are the appropriate category for describing a misorientation between two grains in their own reference frames, allowing for interchanging of the labels of the two grains. The ‘specialness’ of a boundary will thus depend on its type. In terms of integer quaternions, all elements of a type may be generated by finding a quaternion in lowest terms $[ABCD]$ with squared magnitude Σ , then generating the set $\{[ABCD], [A+B A-B C+D C-D], [A+C A-C B+D B-D], [A+D A-D B+C B-C], [A+B+C+D A+B-C-D A-B+C-D A-B-C+D], [A+B+C-D A+B-C+D A-B+C+D A-B-C-D]\}$, then finding all the permutations and sign reversals for each element of this set, and finally eliminating redundancy (recognizing that two quaternions that are nonzero scalar multiples of each other are equivalent) (Grimmer, 1984).

Let us consider $\Sigma 3$ as an example. All integer quaternions with squared magnitude 3 will have the same ‘shape’ (i.e. will be equivalent but for permutations and sign reversals) as [0111], which happens to represent a 180° rotation about [111]. Since there is only one such shape, $\Sigma 3$ has only one ‘type’. We can generate every single $\Sigma 3$ quaternion using the methods in the preceding paragraph. We find that they are the permutations and sign reversals of the set of shapes $\{[0111], [0112]$ and $[3111]\}$ (with squared magnitudes 3, 2×3 and 4×3), 96 in total once the overall-minus-sign redundancy is eliminated. These can be separated into 4 cosets of 24 elements each. We can arbitrarily choose one element of each coset to represent it; later in this work we will use the 180° rotations about [111], [1 $\bar{1}\bar{1}$], [$\bar{1}\bar{1}$ 1] and [$\bar{1}\bar{1}$ 1]. These four elements, expressed as rotation matrices [using, e.g., the formulas in Grimmer (1974)], can be right-multiplied by each of the 24 elements of $\Sigma 1$ to generate every possible $\Sigma 3$ rotation.

We have categorized the CSL rotations, in hierarchical order, according to their Σ numbers, types $(\Sigma 1R\Sigma 1) \cup (\Sigma 1R^{-1}\Sigma 1)$, subtypes $(\Sigma 1R\Sigma 1)$ (which are often the same as the types), cosets $R\Sigma 1$, and individual rotations R (summarized in Table 1). The process of mapping the cubic CSL rotation group will then consist of identifying the cosets and determining the types of the misorientations for each pair of cosets. The symmetry properties allow us to do this for a single representative of each coset, with the assurance that the choice of representative will not affect the outcome.

3. Algebraic representations

3.1. Context and statement of the problem

We will start by deriving the properties of an algebraic representation of the CSL group as strings of representative rotations a_i^p with i identifying a coset and the superscript p giving the Σ value (the variable p will always represent an odd prime in this work). The choice of which representative to use for a given coset is arbitrary, but some choices will be much more convenient than others. It is always possible to choose representatives a_i^p such that each coset is represented exactly once, and the inverse of each representative is also a representative (possibly itself). The construction of representatives is simplest in terms of integer quaternions and relies on the fact (Grimmer, 1974, 1976) that the relevant quaternions will fall into only a few classifications with known numbers of cosets for a given type. Table 2 explicitly shows one possible construction for each of these classifications. The types themselves may be easily constructed by finding all the fundamentally distinct solutions over the integers of $w^2 + x^2 + y^2 + z^2 = p$, for a given p , which are easily found by trial and error. Since we can construct the a_i^p for any type, we can combine these constructions to obtain our representatives for the entire set Σp . Σp itself is obtained by right-multiplying the set of a_i^p by the 24 elements of $\Sigma 1$. Each element will come up exactly once in this enumeration. The verification that any two representatives as produced in Table 2 are in different cosets is a straightforward but tedious exercise in algebra and will not be detailed here.

Once we have constructed our representatives, we may (as will be proven below) represent any given CSL as a string representing an ordered product of representatives. Once we have derived a few rules for the identification, uniqueness and manipulation of these strings, we will implicitly have an algorithm for mapping any desired part of the CSL group. Once this is done, we will be able to start drawing the graphs.

Besides the definitions and development we have already presented, and the elementary result that the set of all CSL rotations is a group (obvious when one considers the rotations in rational matrix form), we will use two results from the literature to complete our proofs. The first is the Σ combination rule (Miyazawa *et al.*, 1996; Gertsman, 2001a), which states that, when two CSL matrices are multiplied such that $AB = C$, their Σ values are related by the formula

$$\Sigma_A \Sigma_B = k^2 \Sigma_C, \quad (2)$$

where k is a common factor of Σ_A and Σ_B . Note that this is often misunderstood, such that many authors insist that $k = 1$ for one of the permutations of A , B and C . This is overly restrictive and discounts perfectly valid triple junctions such as $\Sigma 9-\Sigma 9-\Sigma 9$ (with $k = 3$) (Gertsman & Tangri, 1995). For example, $\Sigma 3-\Sigma 15-\Sigma 5$ is valid (with $k = 3$, a common factor of 3 and 15), while $\Sigma 9-\Sigma 5-\Sigma 5$ is not (since $k = 3$ is not a factor of 5). The second result is a count of the number $m(\Sigma)$ of CSL rotations of a given Σ . This has been variously represented (Grimmer, 1973, 1976), and may be derived directly from number theory as

$$m(\Sigma) = 24\Sigma \prod_p \frac{p+1}{p}, \quad (3)$$

where p is the set of prime factors of Σ [defined so that the special case $m(1)$ equals 24]. The factor of 24 comes from the 24-fold redundancy in the representations of CSL's in terms of rotations, and may be omitted if it is understood that we are talking about cosets of $\Sigma 1$ rather than individual rotations. Thus we define $m'(\Sigma) = m(\Sigma)/24$. For example, $m(1) = 24$, $m(3) = 96$ and $m(9) = 288$, with $m' = 1, 4$ and 12 cosets, respectively. In the special case where $\Sigma = p^n$ with $n > 0$, we have $m'(\Sigma) = (p+1)p^{n-1}$. With $n = 1$, we have $(p+1)$ cosets of 24 elements each. Thus, the index i in the representatives a_i^p can be taken to range from 1 to $p+1$.

It follows directly from the Σ combination rule that left- (or right-) multiplying a rotation R with $\Sigma = \sigma$ by an element of Σp will result in either an element of $\Sigma(\sigma/p)$ ('lowering' σ) or one of $\Sigma(\sigma p)$ ('raising' σ). All Σ values are integers, so if p does not divide σ , then all the a_i^p must raise σ . On the other hand, if p divides σ , then some of the representatives a_i^p may lower σ , while others raise it.

What we shall prove is that, in the latter case, all the a_i^p will raise σ except for one, which will lower it. We will use integer quaternions for purposes of this proof. To avoid confusion with the a_i^p which up to now have been matrices, we will use the symbol A_i^p for the quaternion form of the representative rotation a_i^p , recalling that the order of multiplication is reversed between the quaternion and matrix representations.

3.2. Proof of the validity of the algebraic representation

The proof rests on a textbook result (Hardy & Wright, 1938; Conway & Smith, 2003) from number theory that states that a quaternion Q (with integer components) with a composite squared norm N may always be factored (although in general not uniquely) into a product of prime quaternions with integer components, one for each element in the prime-factor expansion of N , and with the factors appearing in any desired order. With this fact in hand, all we need do is interpret this old result in terms of our CSL rotations. The technical complications involving what Conway & Smith term 'unit migration' as well as the factors of 2 that are explicitly dealt with by Hardy & Wright are handled in our case by the recognition of the hierarchy of equivalence classes, the establishment of our arbitrary conventions and the transfor-

mations which we will encounter later in equations (8)–(10). An understanding of such details is not necessary for following the discussion, however, and we only note this issue in passing for the interested reader.

We have hypothesized (using the notions in §2.1) that we have a Q with $N = \sigma, 2\sigma$ or 4σ , and σ containing a factor of p , an odd prime. We may therefore produce a product $q_1 q_2 \dots q_k = Q$ of prime quaternions q_i such that the squared norm of $q_k, |q_k|^2$, is equal to p (to within factors of 2, which are unimportant *via* the arguments in Hardy & Wright, once we recognize that the considerations involving associates and factors of 2 are taken care of by our $\Sigma 1$). Let $Q' = q_1 q_2 \dots q_{k-1}$. Since for any quaternions Q_1 and Q_2 we have $|Q_1|^2 |Q_2|^2 = |Q_1 Q_2|^2$ (Conway & Smith, 2003), we know that $|Q'|^2 =$ one of $\sigma/p, 2\sigma/p$ or $4\sigma/p$, *i.e.* Q' is a CSL rotation in $\Sigma(\sigma/p)$. Our purpose in splitting off this quaternion prime factor q_k is to determine how many of the A_i^p will lower σ , knowing that all the others will raise it.

Now, since q_k is in Σp , there exists exactly one A_i^p in the same coset as q_k , *i.e.*

$$q_k = SA_i^p \quad \text{or} \quad q_k A_{-i}^p = S \quad (4)$$

(to within meaningless scalar factors), where S is a quaternion representing an element of $\Sigma 1$ and A_{-i}^p is the quaternion representative of the inverse rotation of A_i^p , which by construction is guaranteed to be a representative as well. (This is just a notation – we know by construction that $A_{-i}^p = A_j^p$ for some positive j , possibly with $j = i$. There are still only $p + 1$ representatives for a given p .) Thus we have constructed a representative quaternion A_{-i}^p such that $Q A_{-i}^p$ is the same rotation as $Q'S \in \Sigma(\sigma/p)$. That is, A_i^p lowers Q , leaving (again to within meaningless scalar factors) the quaternion Q' . This procedure may be iterated to produce a prime-factor decomposition of Q , as we will see below.

The fact that there is in general more than one expansion $q_1 q_2 \dots q_k = Q$ with $|q_k|^2 = p$ does not in any way restrict the validity of the proof, for the Σ number of a rotation is entirely independent of the choice of representation. Thus the fact that A_i^p lowers Q for one of the $p + 1$ choices of i is an invariant with respect to the choice of expansion of Q . The only arbitrariness is in our *a priori* choices of the A_i^p among all the possible representatives of the cosets of Σp .

Having completed the existence part of the proof, we return to the matrix representations (which will be more convenient for proving uniqueness), recalling that the order of multiplication is reversed. If we represent σ as a product of odd primes in non-decreasing order:

$$\sigma = p_1^{n_1} p_2^{n_2} \dots p_k^{n_k}, \quad (5)$$

where the n 's are exponents. By repeated application of the above algorithm, we may generate a string of representative prime factors of the coset of a rotation matrix R :

$$RS_j = a_{i_1, p_1}^{p_1} a_{i_2, p_1}^{p_1} \dots a_{i_{n_1}, p_1}^{p_1} a_{i_1, p_2}^{p_2} \dots a_{i_{n_k}, p_k}^{p_k}, \quad (6)$$

where the i 's are the various indices identifying the representatives in the string and $S_j \in \Sigma 1$. For example, we might find that a certain $\Sigma 45$ rotation matrix R may be lowered by left-

multiplying by a_{-1}^3 , resulting in a new matrix R' , which itself may be lowered by left-multiplying by a_{-3}^3 , leaving a matrix R'' , which may finally be lowered by left-multiplying by a_{-2}^5 , leaving the $\Sigma 1$ matrix S_8 . Thus, our original matrix R may be expressed as the product $a_1^3 a_3^3 a_2^5 S_8$, with $3 \times 3 \times 5 = 45$ and the subscripts chosen arbitrarily for purposes of the example. This string is sorted in the sense that the prime factors are in non-decreasing order (we could just as easily have produced a reverse-sorted string, a concept that we will return to as it implies some interesting things about loops in the CSL group). The string is also of minimum length for representing R , since by the Σ combination rule we need at least as many prime factors a_i^p as there are elements in the product-of-primes expansion of σ . It is clear that no two adjacent elements of the string are inverses of each other, for then the string would be immediately shortenable to an equivalent string, which we know not to exist. This property of being not immediately shortenable we will abbreviate as NIS. Thus, we have an algorithm for producing a minimal-length sorted NIS string representing the coset of any given CSL rotation.

To prove uniqueness at this point requires merely a bit of combinatorics, by which we find that the number of sorted NIS strings with a prime expansion as given in (5) is exactly

$$M'(\sigma) = \sigma \prod_{j=1}^k \frac{p_j + 1}{p_j}, \quad (7)$$

which is equal to the number $m'(\sigma)$ of cosets in the set $\Sigma \sigma$. This is because the first representative for a given p may be chosen to be any of the $p + 1$ options, but after that we are not allowed to choose a representative that is the inverse of the one we have chosen most recently. By the algorithm given above, any coset will produce a minimal-length sorted NIS string, and clearly a single sorted NIS string (which describes one specific rotation) cannot represent two different cosets. So by applying our algorithm to one chosen element of each coset of a given set $\Sigma \sigma$, we can generate all $M'(\sigma) = m'(\sigma)$ of the sorted NIS strings of the form given in (6), with no chance for redundancy or ambiguity. When this is applied to the whole CSL group, no sorted NIS strings of any size will be left unassigned.

Therefore, we come to the core result, that any CSL rotation is represented by exactly one sorted NIS string. As a corollary, we find that, if a sorted string is not immediately shortenable, then it is not shortenable at all, *i.e.* no shorter string represents the same coset and the prime-factor expansion of the Σ number of the rotation may be read directly off the sequence of representatives.

We should emphasize that this uniqueness comes at the cost of some arbitrary conventions, namely the specific selection of the representatives (or generators) a_i^p and the sorting of the primes in non-decreasing order. These derive from the $\Sigma 1$ symmetry group and the non-commutativity of rotations. If the primes are sorted in a different order (*e.g.* non-increasing), we will obtain a different NIS string (*e.g.* a reverse-sorted one) and there is no guarantee that the representatives in different sortings will have any particular relationship. They may be

identical, they may partially overlap or they may not overlap at all and even be of different types (Barber, 2003). The requirement that the inverse of a representative be a representative was convenient for the sake of the proof but it may be dropped at the expense of having to manipulate more elements of $\Sigma 1$ in the algebra.

3.3. Algebraic rules

Now that we can generate our strings, we need to know how to manipulate them. Taking the inverse of a string is simple; we merely reverse the order of the elements and replace each element with its inverse. Multiplying two strings is done by simple concatenation. Two adjacent inverse elements may be removed or inserted into any string at will. If we restrict ourselves to a set Σp^ω in which all the Σ values are powers of a single prime p , then this completes all the rules for algebraic manipulation. This is because the issue of sorting does not arise, and there is exactly one NIS string for any coset. In other words, if we think of the set Σp^ω as being a collection of cosets linked by the representatives a_i^p , then the set is a so-called 'free group', that is, a group with no relationships among its generators (in this case, the representatives) other than the existence of an inverse. The structure of a free group is quite simple and may be represented as a graph with no loops and the same number of neighbors (in this case $p + 1$) for every cell. This is the basis of the topological model discussed in the next section.

The additional algebraic rules needed to handle cases of Σ having more than one prime factor are significantly more cumbersome but can be compactly expressed in the following forms:

$$a_i^{p_1} a_i^{p_2} = a_k^{p_2} a_l^{p_1} S_m \quad (8)$$

$$S_m a_i^p = a_j^p S_n \quad (9)$$

and
$$S_m S_n = S_k, \quad (10)$$

where the relationships among i, j, k, l and m must be determined in each particular case. The result will depend on the choice of representatives for the two primes p_1 and p_2 . In words, we may swap prime representatives provided we modify their subscripts and account for a possible $\Sigma 1$ discrepancy, we may swap an element of $\Sigma 1$ with a representative provided the subscripts are adjusted accordingly, and we may perform operations based on the group structure of $\Sigma 1$. Equations (10) simply constitute the character table of $\Sigma 1$. As for transformations of the form given in (9), it is possible to work out every possible case by simply finding all the S_m and S_n that satisfy the equation for any given pair of representatives of the same subtype. By definition (recalling §2), representatives of different subtypes will not interact *via* equations of this sort. Once this problem is solved for all the classes given in Table 2, there are no more such relationships to be found. We have done this but the results are space-consuming and will only appeal to a small audience and so are omitted. Note that except in the lowest-symmetry cases there will be more than one m and n solving the equation for a given p, i and j .

The remaining algebraic relationships, as given in (8), have to be generated systematically for each specific case (*i.e.* pair of prime factors) of interest. There is no guarantee that the two pairs of generators even have types in common and so the job must be done for each pair of prime factors involved in a calculation rather than just for the classes given in Table 2. It is a simple matter to program a computer to do this automatically. We have done this for all primes less than 50 (which surely includes all physically meaningful cases) but again the resulting tables would be too large to publish here.

Some examples should clarify the notation. For compactness, we will give the examples in quaternion rather than matrix form (recalling, again, that the order of multiplication is reversed). Equation (8) has the following example from $p_1 = 3, p_2 = 5$:

$$\begin{aligned} & [0 \ \bar{1} \ \bar{1} \ 1][0 \ 0 \ 2 \ 1] \\ &= [0 \ 1 \ 0 \ 0][0 \ \bar{2} \ 1 \ 0][0 \ \bar{1} \ 1 \ \bar{1}] \\ &= [1 \ \bar{3} \ 1 \ \bar{2}], \end{aligned} \quad (11)$$

where the Σ values of the quaternions appearing in this equation are merely their squared magnitudes, in order 3, 5, 1, 5, 3 and 15. This shows two factorizations of the $\Sigma 15$ quaternion $[1\bar{3}1\bar{2}]$, one sorted and the other reverse-sorted, with the $\Sigma 1$ element compensating so as to obtain the same representation of the coset in both cases. An example of (9) is:

$$\begin{aligned} & [0 \ \bar{1} \ \bar{1} \ 1][1 \ 1 \ 0 \ 0] \\ &= -[1 \ 1 \ 0 \ 0][0 \ 1 \ \bar{1} \ \bar{1}] \\ &= [1 \ \bar{1} \ 0 \ 2], \end{aligned} \quad (12)$$

where in this case the Σ values are, in order, 3, 1, 1, 3 and 3 (since the squared norms are 3, 2, 2, 3 and 6). The overall minus sign on the middle expression is meaningless (and would not appear in the matrix representation). Thus, the two rotations of type $\Sigma 3$ have been interconverted by multiplying by appropriate elements of $\Sigma 1$. Finally, (10) has as an example:

$$[0 \ 1 \ 0 \ 0][1 \ 1 \ 1 \ 1] = [\bar{1} \ 1 \ \bar{1} \ 1], \quad (13)$$

in which all quaternions are $\Sigma 1$.

Once all the algebraic transformations [inversion, pair creation and annihilation, and equations (8)–(10)] are worked out for the prime factors of interest, then all calculations may be done entirely in terms of representatives and $\Sigma 1$ elements; the structure of the group is entirely implicit in these rules, and there is no further need to consider matrix or quaternion representations except to link to real-world calculations. In general, these algebraic manipulations would be quite tedious, but the transformation rules in the Σp^ω case are quite simple and allow us to make a number of observations *via* examples. We proceed to do this in the important case of $\Sigma 3^\omega$, known as the twin-related CSL group since all the rotations may be produced by combinations of coherent $\Sigma 3$ twinning operations. This group captures all possible results of interacting $\Sigma 3$ twins, which are very important in a large class of materials; also, it has been noted that any high-angle boundary may be well approximated as a $\Sigma 3^n$ misorientation, with n not very large (Kopecky *et al.*, 1991).

3.4. Examples drawn from $\Sigma 3^\omega$

It should be clear that the type of rotation is encoded in the string, since the string contains all the information as to which coset the rotation belongs to (which, if we recall the hierarchy in Table 1, also identifies the type). In the case of $\Sigma 3^\omega$, the encoding is particularly simple owing to the very high symmetry of the set of generators. Not only are all the generators of the same type (180° rotations about $\langle 111 \rangle$ axes), but each pair of generators has the same relationship (since any pair of 180° $\langle 111 \rangle$ rotations has the same geometrical relationship as any other pair, within symmetry operations of the crystal). What this means algebraically is that $a_i^3 a_j^3$ is the same type of rotation (namely a $\Sigma 9$) for any distinct i and j , and is the identity when $i = j$. This rule for multiplying representatives of $\Sigma 3$ is the algebraic representation of $\Sigma 3$ – $\Sigma 3$ – $\Sigma 9$ and $\Sigma 1$ – $\Sigma 3$ – $\Sigma 3$ triple junctions (to be clarified below). We also see that there are two types of $\Sigma 27$ rotations $a_i^3 a_j^3 a_k^3$ with $i \neq j$ and $j \neq k$, namely the ones with $i = k$ (which turn out to be the $\Sigma 27a$ rotations, as is easily checked by direct calculation) and the ones with $i \neq k$ (the $\Sigma 27b$'s). Simple combinatorics produces the expected result, that there are 24 cosets of $\Sigma 27b$ but only 12 of $\Sigma 27a$, in accord with standard tables (Mykura, 1979). $\Sigma 81a$ rotations are represented as $(a_i^3 a_j^3 a_k^3 a_l^3)^{\pm 1}$, $\Sigma 81b$ as $a_i^3 a_j^3 a_k^3 a_l^3$ (i, j, k and l are all distinct) etc. and again the multiplicities come out as they should, with 48 cosets of the former and 24 of the latter. Gottstein (1984) produced a table relating strings to types on $\Sigma 3^\omega$, but using notations and conventions that differ substantially from ours, and with no proofs of the general validity of the approach.

If we were to consider $p = 5$ or more, the high symmetry among the set of generators in the $p = 3$ case is lost, which accounts for there being, e.g., more than one type of $\Sigma 25$ rotation (so the 'type' of $a_i^p a_j^p$ with $i \neq j$ can vary depending on the choices of i and j if $p = 5$, unlike the case of $p = 3$). In general, the identities and multiplicities of each type of rotation can be traced to the symmetries among the generators and in the representative strings. These symmetries ultimately link back to the CSL-type categorizations of Grimmer (1973, 1974, 1976). Note that, if we had chosen a less symmetric set of representatives of the $\Sigma 3$ rotations, the symmetry properties and type classifications of the higher Σp^ω rotations would not be as obvious – the result that the distinction between a $\Sigma 27a$ and a $\Sigma 27b$ lies in the identity of the first and last elements is peculiar to our choice of representation.

Next, consider the algebraic expression of the Σ combination rule (see Fig. 1). We have three grains with orientations I , A and B , as expressed in the reference frame of grain 1, and the misorientations, in clockwise order and expressed in the frame of the first grain of each pair, are $X_I = A$, $Y_A = A^{-1}B$ and $Z_B = B^{-1}$, with the product $X_I Y_A Z_B = I$ expressing the conservation of orientation (the subscript on a rotation giving the frame in which the matrix is expressed). If all three misorientations are elements of Σp^ω , then they will each be expressible uniquely as a string of representatives a_i^p . When we form the product XY (which must equal Z^{-1} , and thus be equivalent to a string of the same length), we will find that

some number k (possibly zero) of representatives on the right-hand side of string X will cancel with an equal number of representatives from the left-hand side of string Y . The result is that the string Z is shorter than the sum of the lengths of the other two by an even number, i.e.

$$n_Z = n_X + n_Y - 2k. \quad (14)$$

Since $\Sigma = p^n$ for each string, this translates to

$$\Sigma_Z = \Sigma_X \Sigma_Y / p^{2k}. \quad (15)$$

Since p^k is a common factor of Σ_X and Σ_Y (as their strings are of length at least k), this is nothing other than the Σ combination rule in the special case of Σp^ω . Thus the algebraic representation helps provide an intuitive basis for the Σ combination rule, but also takes us one step further, as we can now make statements about the relative likelihood of different values of k in (2). This is because we know that only one of the $p + 1$ possible representatives of a given p will be able to lower the Σ value, with the other p representatives raising it. A similar thing happens when we generalize from Σp^ω to the entire CSL group, but we may have to rearrange the order of factors *via* equations (8)–(10) to produce the maximum number of cancellations.

As a specific example, consider the fact that a $\Sigma 3$ – $\Sigma 9$ – $\Sigma 27b$ triple junction is allowed. An element of $\Sigma 27b$ might be represented as the product $a_1^3 a_2^3 a_4^3$, which may be shortened by left-multiplying by a_1^3 (which is its own inverse), leaving $a_2^3 a_4^3$, an element of $\Sigma 9$. This is the algebraic representation of the 'reaction' $\Sigma 27b = \Sigma 9 + \Sigma 3$, possible at a triple junction. In this case, the conservation of misorientation equation would be expressed as:

$$(a_1^3 a_3^3 a_4^3)(a_3^3 a_4^3)^{-1}(a_1^3) = I. \quad (16)$$

Note the inversion on the $\Sigma 9$ rotation, which is necessitated by the need to go in a consistent direction around the loop. For another example, a $\Sigma 9$ – $\Sigma 9$ – $\Sigma 9$ triple junction might be represented as:

$$(a_1^3 a_2^3)(a_2^3 a_3^3)(a_3^3 a_1^3) = I. \quad (17)$$

It is certainly possible to construct such a triple junction by multiplying pairs of $\Sigma 9$ matrices until the result is a $\Sigma 9$ matrix (Gertsman & Tangri, 1995; Miyazawa *et al.*, 1996). But the algebraic representation makes it clear exactly what form any solution must have and how many solutions are possible – even in cases in which no solution exists. For example, if we try to construct a $\Sigma 9$ – $\Sigma 27a$ – $\Sigma 27a$ triple junction, we find it to be impossible, even though it satisfies the Σ combination rule. Since in our choice of representation a $\Sigma 27a$ will be of the form $a_i^3 a_j^3 a_i^3$ for some i and j , it is clear that there is no pair of $\Sigma 27a$ strings, the product of which will be a $\Sigma 9$ string. Thus we can avoid trying to find such a triple junction by trial and error.

This demonstrates an important result that has been somewhat overlooked in the literature, that the satisfaction of the Σ combination rule is not sufficient to guarantee crystallographic compatibility. It is necessary to keep track not only of the Σ values of the involved rotations, but also their types and even the specific identities of the representatives that

produced the rotations. The restriction in the triple-junction case is relatively weak, in that some combinations of types of rotations with the given Σ values will not be compatible, but as we consider more complicated topologies the added restrictions start to interact with one another. For quadruple nodes, there are Σ labellings of the boundaries that satisfy the Σ combination rule at every triple junction (and thus are locally crystallographically compatible) but are still impossible (see the next section for an example). For large-scale realistic grain-boundary networks, the additional restrictions are complex and quite difficult to work out in general. In a later section, we outline a procedure that is guaranteed to generate crystallographically compatible type labellings of all the junctions in a network of any complexity. We do not know of a computationally efficient solution for the inverse problem, *i.e.* given a grain topology and a set of Σ and type labellings, to determine whether there are any compatible grain-orientation assignments. Using the techniques discussed in this paper, though, we have expressed the problem in terms of graph theory and the algebraic manipulation of strings, so that there is no longer any need to work with explicit representations of rotations. The problem is to find an assignment of one string to each node in the graph of the grain topology such that the type of each link in the graph (representing a grain boundary) matches the labelling we have chosen.

To summarize, we have developed an algebraic representation of any CSL rotation coset (that is, a grain orientation) as a prime-factor decomposition that is unique on Σp^ω and unique subject to a sorting convention on the entire group. The fundamental algebraic rules for manipulating the strings are simple, namely that two adjacent elements with the same p may be simplified if and only if they are inverses, and two elements with different p 's may be swapped if their indices are adjusted according to case-specific but easily discovered rules [equation (8)], possibly with the introduction of $\Sigma 1$ elements in the string, which may be manipulated with additional rules [equations (9)–(10)]. The rules suffice to convert any string whatsoever into an equivalent sorted NIS string, which we know to be unique. To see why, consider that we can always swap prime elements until they are sorted, moving any generated elements of $\Sigma 1$ to the right end of the string, where they do not interfere with the coset definition of equivalence. Since the rules are reversible, this means the rules suffice to generate any string equivalent to a given string. The Σ combination rule, differences among types of rotations, and




Figure 1 Definitions and governing equations for conservation of orientation (also called crystallographic compatibility) around a triple junction.

multiplicities of various types of rotations come directly out of the algebraic rules and the symmetries of the generators a_i^p . Once the generators are constructed and their symmetries identified, we may simply do algebra on the generators and be guaranteed that, for example, any triple junction we produce will be crystallographically correct.

With this, we are able to move on to drawing and interpreting graphs.

4. The topological model

4.1. Topological mappings, with examples from $\Sigma 3^\omega$

We may now take our results and re-describe them in terms of graphs. This requires merely an interpretation of the algebraic representation; no new results need be derived. In fact, we can delve right in, starting with Fig. 2, a map of a portion of $\Sigma 3^\omega$. Each vertex (marked with a square) is a coset, representable by exactly one non-shortenable string of the a_i^3 . The

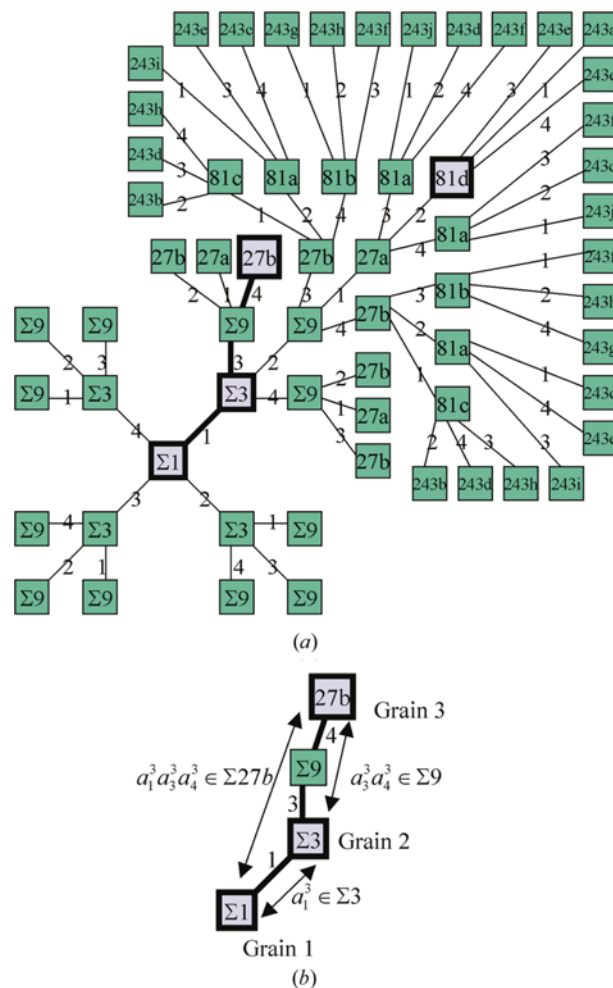


Figure 2 (a) Map of a portion of the $\Sigma 3^\omega$ twin-related group. Each box is one coset of 24 elements. Links are labelled with the index i for the representative a_i^3 linking the two cosets. Examples discussed in the text are highlighted. (b) Highlight of the $\Sigma 3$ – $\Sigma 9$ – $\Sigma 27b$ triple junction example, showing how to generate the triple-junction types from the labels of the links connecting each pair of grain orientations.

index i of the generator that links two adjacent cosets is indicated as a number from 1 to 4 for each link. Since the generators of $\Sigma 3^\omega$ were chosen to be their own inverses, there is no need to make this a directed graph, although this is not always possible for larger primes. The graph is that of a free group with four neighbors per vertex, and labelled such that the Σ value of a vertex is equal to 3^n , with n the number of links between it and the central $\Sigma 1$ vertex, which is the same as the number of elements in its shortest string representation. We have only drawn a fraction of the elements beyond $\Sigma 9$, but the remainder up to $\Sigma 243$ may easily be constructed owing to the symmetry of the structure. Implicit in this graph is the set of distinctions among the different types in terms of symmetries in the strings, as well as the multiplicity of each type.

The string representing a given coset may simply be read off the graph, starting from the $\Sigma 1$ origin. For example, the only $\Sigma 81d$ rotation shown may be reached by a sequence labelled (1,2,1,2), and thus is represented by $a_1^3 a_2^3 a_1^3 a_2^3$. If we were to instead choose a more circuitous route, say along the sequence labelled (1,2,1,3,3,2) (going out on a branch and then returning *via* the same path), the resulting string would be $a_1^3 a_2^3 a_1^3 a_3^3 a_3^3 a_2^3$, which is immediately shortenable *via* the algebraic rule $a_3^3 a_3^3 = I$ to the previous representation. This is the graphical representation of the double elimination rule, that going out and returning along the same path has no effect. Since Σp^ω is a free group in our representation, there is no other way to generate an ambiguity in the generation of a string.

Allowed triple junctions may be immediately generated simply by choosing three of the vertices (which represent the three grain orientations) and tracing the paths that link them. For example, if we choose the $\Sigma 1$ origin, the $\Sigma 27b$ vertex linked to it by the path (1,3,4), and the $\Sigma 3$ vertex linked to $\Sigma 1$ by the line labelled 1 (the relevant cosets and links being highlighted in the figure), we find that the three paths linking the three grains are the rotations $a_1^3 a_3^3 a_4^3$, $a_3^3 a_4^3$ and a_1^3 , and this is exactly the same $\Sigma 3$ – $\Sigma 9$ – $\Sigma 27b$ triple junction we used as an example in a previous section. There is no need for one of the grains to be at the origin; the grains a_2^3 , a_1^3 and $a_1^3 a_3^3$ would also form a $\Sigma 3$ – $\Sigma 9$ – $\Sigma 27b$ triple junction, with all the rotations of the same type as in the first example, and thus entirely equivalent to the first triple junction to within an arbitrary choice of coordinate systems. In other words, two triple junctions that look equivalent on this topological graph really are equivalent in physical terms. This is unsurprising, as a group is in some sense homogeneous, so that the network must look topologically the same from any node. We can also immediately see that there are no $\Sigma 9$ – $\Sigma 27a$ – $\Sigma 27a$ triple junctions (there being no pair of $\Sigma 27a$ vertices separated by a distance of exactly two), while it is easy to construct $\Sigma 9$ – $\Sigma 9$ – $\Sigma 9$ triple junctions (simply choose three of the $\Sigma 3$ vertices, for example).

The fact that double elimination is the only simplification possible on strings in Σp^ω has a practical consequence for the important case of twin-dominated structures, which tend to have clusters of grain boundaries with nearly ideal $\Sigma 3^\omega$ misorientations (Gertsman *et al.*, 1994; Randle, 1996). The

restrictions on the ways in which such clusters may be constructed are central to discussions in grain-boundary engineering (Gertsman *et al.*, 1994; Randle, 1996; Minich *et al.*, 2002). Consider such a cluster, and take some starting grain within it to have the reference orientation (*i.e.* start at the $\Sigma 1$ vertex in the graph). Now if we draw a contiguous path from one grain to the next, we will follow a sequence of vertices in the $\Sigma 3^\omega$ graph, with the misorientation between adjacent grains represented by the path linking their vertices in the graph. If the path through the set of grains is closed, returning to its starting point, then the path through the $\Sigma 3^\omega$ tree must similarly return to the $\Sigma 1$ node, and the only way for this to happen is for the steps to be retraced exactly. So in our $\Sigma 3$ – $\Sigma 9$ – $\Sigma 27b$ example, going from grain 1 to grain 2 we have the string a_1^3 , then from grain 2 to grain 3 we have $a_3^3 a_4^3$ for a total of $a_1^3 a_3^3 a_4^3$, and to get back to grain 1 we must follow a path $a_3^3 a_4^3 a_1^3$, retracing the paths in the tree exactly. This procedure, concatenating the strings in a closed loop of grain boundaries, must retrace the steps regardless of the size of the loop, so long as only ideal $\Sigma 3^\omega$ boundaries are involved. In a highly twin dominated structure, these loops can be very large, involving many grains, and will overlap significantly. The resulting constraint on the structure is quite strong – much stronger than would be implied by the Σ combination rule alone. There will be many imaginable structures that satisfy the Σ combination rule at every triple junction yet fail to work on a larger scale (an example is described below). A large cluster of pure $\Sigma 3^\omega$ misorientations has much fewer (discrete) degrees of freedom in the grain-boundary assignments than would be expected at first glance. All of these constraints are implicit in the structure of Fig. 2(a); now that we have drawn this graph there is no need to consider the explicit matrix or quaternion forms of the rotations so far as crystallographic compatibility is concerned.

This procedure can be generalized to any grain-boundary topology (with each node a grain and each line a boundary) and the entire CSL group. The algorithm is simple: We assign one node in the CSL graph to each node in the grain-boundary topology graph. Then the type of each boundary may be read as the sequence of labels on the shortest path in the CSL graph linking the two CSL nodes corresponding to two adjacent grains. In short, we simply have to specify a relation between the two topologies (the CSL graph and the grain-boundary network graph), and everything about the orientations and misorientations in the entire structure is determined, in a way guaranteed to be crystallographically compatible. This procedure is very straightforward in the case of $\Sigma 3^\omega$, but gets more complex as we add more prime factors.

Some examples should clarify. For a triple junction, the grain topology graph is simply a triangle. For a quadruple node, it is a tetrahedron, with each vertex representing a grain, each edge a grain boundary, and each side a triple junction. Fig. 3 shows examples of possible and impossible constructions on such a tetrahedron, still staying within the $\Sigma 3^\omega$ group. The mapping between the $\Sigma 3^\omega$ graph and the tetrahedral grain topology graph is represented by the capital letters on the $\Sigma 3^\omega$ graph. In the first example (Fig. 3b), each pair of capital letters

is separated by exactly two $\Sigma 3$ links, meaning that all the boundaries are $\Sigma 9$, so that this is a $6\Sigma 9$ quadruple node with four $\Sigma 9$ - $\Sigma 9$ - $\Sigma 9$ triple junctions. The other examples work in the same way. Several of these examples (and some of the associated observations) have already been shown in the literature (Fortier *et al.*, 1995; Miyazawa *et al.*, 1996; Gertsman, 2001a,b), but generally using the relatively cumbersome

matrix representations of the rotations, and probably with some trial and error involved in the generation of some of the examples. We hope the reader will agree that, using the representations developed in this work, these examples and observations may be constructed with more ease, clarity, generality and (in some cases) mathematical rigor. Specifically, besides being able to construct examples of structures that are

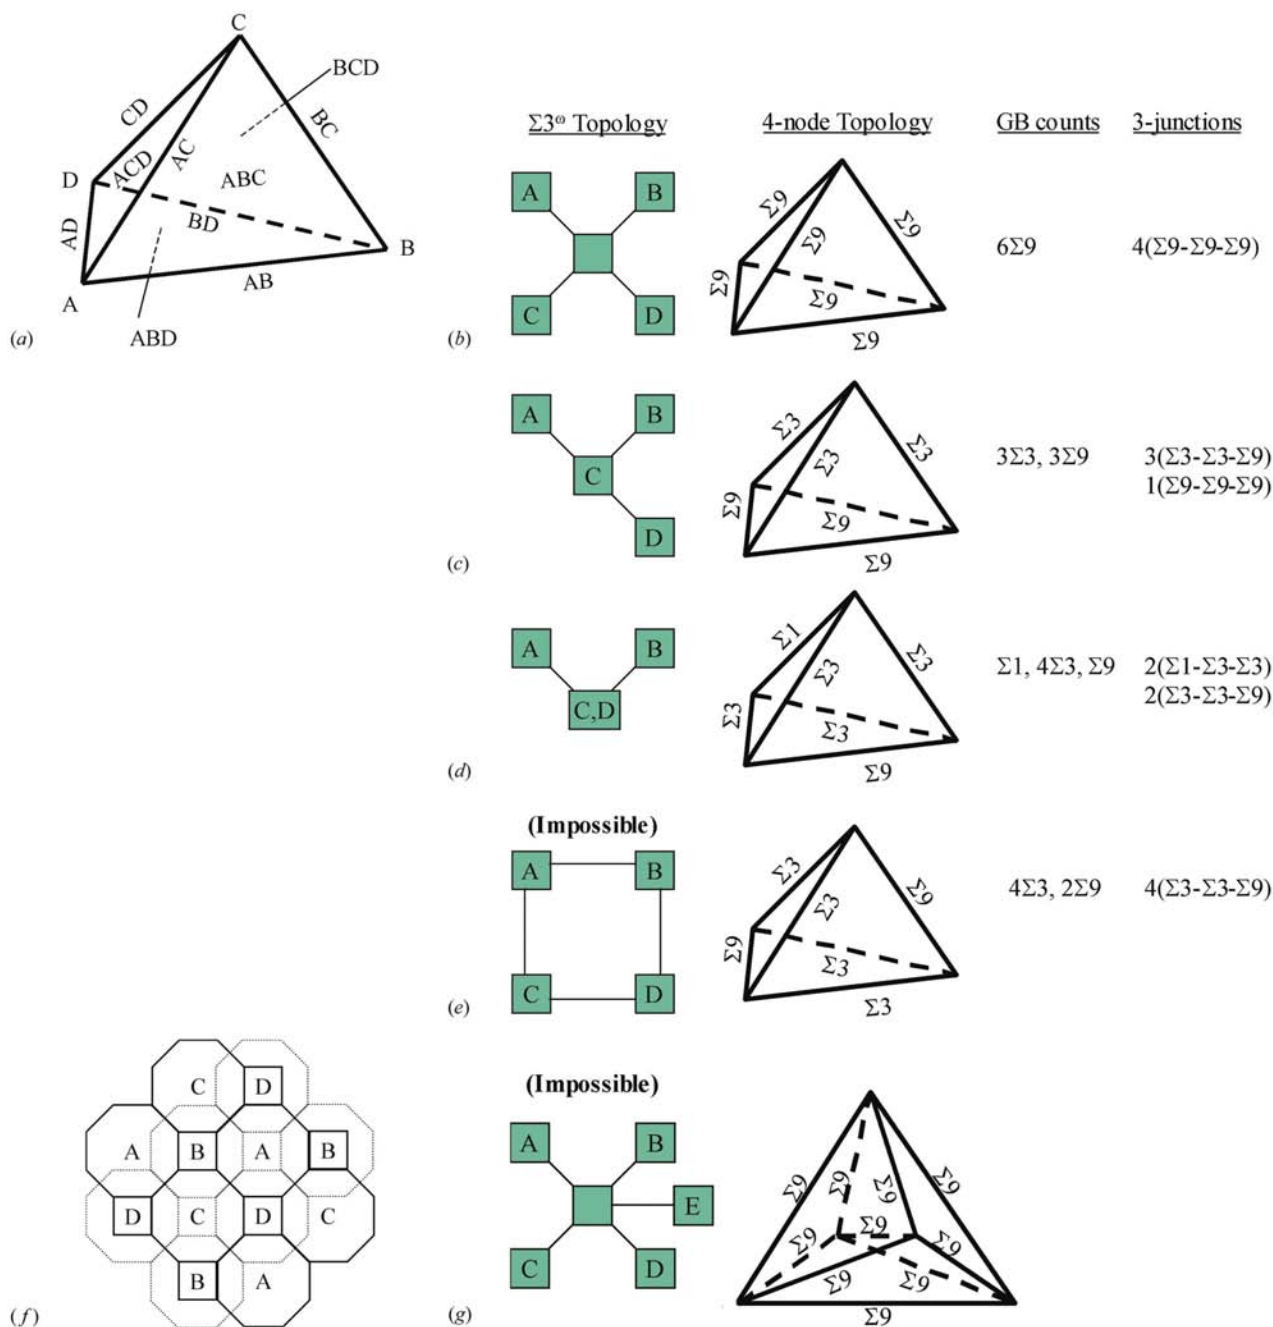


Figure 3 Tetrahedral representation of a quadruple node. (a) Notation, with single letters denoting grains, pairs denoting grain boundaries and triples denoting triple junction lines. (b)–(d) Several examples of crystallographically compatible quadruple junctions. (e) One example that is incompatible yet satisfies the Σ combination rule along all triple lines. (f) Assignments of orientations from (d) to the tetrakaidecahedral lattice, producing a $\Sigma 3$ fraction of $5/7$. Assignments are shown for one plane (solid lines) and the next plane above it (dotted lines). Each succeeding pair of planes will swap the labels in pairs (A, C) and (B, D). This example also illustrates the four-colorability of the network and how the aggregate statistics may differ from the statistics of each quadruple node in three-dimensional networks. If instead we assign the orientations from (c), we obtain a structure with a $\Sigma 3$ fraction of $4/7$ and with only $\Sigma 3$ and $\Sigma 9$ junctions present. (g) A case in which all quadruple nodes are allowable yet the global structure is not.

possible, we can now easily prove that certain large general classes of structures are not possible. With some practice in applying the graphs, one can develop an intuition for the sets of constructions that are and are not possible.

Using the construction in Fig. 3(b), we can see that it is possible to make an entire grain-boundary network with only $\Sigma 9$ boundaries, if and only if the grain topology graph is four-colorable, *i.e.* we can label each vertex with one of only four labels, with no two adjacent vertices having the same label. [Note that the 'if' part of the theorem is implied in the work of Miyazawa *et al.* (1996), but that the 'only if' part is somewhat harder to come by *via* their methods.] We choose the labels as shown in Fig. 3(b), and the construction is complete. Since there is no fifth orientation that has a $\Sigma 9$ relationship to more than one of these, the construction fails for graphs that require more than four colors. Since all planar graphs are four-colorable (Appel & Haken, 1989), the construction is always possible in a plane but almost never possible in three dimensions, where four-colorability is a rare exception (in fact there is no limit to the number of colors potentially required in three dimensions). In practical terms, this means that a thin film with a columnar grain structure has the potential for a higher proportion of $\Sigma 9$ boundaries (potentially 100%) than does a random three-dimensional grain structure. We note that the tetrakaidecahedral lattice often used in grain-boundary models is one of these highly atypical four-colorable three-dimensional networks, which will potentially skew simulation results that depend on the network connectivity, while the three-colorability of the commonly used hexagonal tiling of the plane is similarly atypical for a two-dimensional network.

The second example, Fig. 3(c), shows the only quadruple node that contains $\Sigma 3$ and $\Sigma 9$ junctions and nothing else. Fig. 3(e) shows an attempt at generating another such example and, even though every single triple junction in this construction satisfies the Σ combination rule, the quadruple node itself is impossible. The $\Sigma 3^\omega$ graph for this structure shows why — it would require a loop in the graph, while $\Sigma 3^\omega$ (taken as a set of cosets) is a free group with no loops. Removing any element in this graph with a loop gives us a valid graph; this is the graphical representation of the fact that every triple junction in the attempted construction satisfies the Σ combination rule. This is an example of how higher levels of structure in a system can introduce frustrations that are not evident at the more local levels. It is easy to produce further examples in which *e.g.* the quadruple junctions are all valid yet the complete structure is impossible. Fig. 3(g) is a simple example, showing the simplest possible non-four-colorable network with every boundary labelled as $\Sigma 9$. We speculate that this process continues indefinitely — that to enumerate all such constraints would imply potential correlations at all distance scales, and moreover that these correlations would be very different in graphs of different typical colorabilities (particularly in two *versus* three dimensions). Since even the Σ combination rule by itself has been shown to have a controlling influence on such ensemble behavior as percolation thresholds in two-dimensional twin-dominated systems

(Minich *et al.*, 2002), we propose that perhaps the detailed structure of this group has far deeper physical consequences than has been generally recognized [mirroring observations made by Gertsman and collaborators (multiple references) for example]. Obviously, the identity and processing history of any material will also have controlling influences, but our results suggest that these influences must be considered in the context of the mathematical structures governing the topology. Until now, a detailed, rigorously derived, intuitively visualizable map of the twin-related rotation group has not been available. We hope that our contribution will help to shed some light on this issue.

Let us now consider the notion of a twin-limited structure. This concept has been used by several researchers (Palumbo *et al.*, 1992; Miyazawa *et al.*, 1996; Gertsman, 2001b) but with important differences in definition. Generally, it refers to a structure with the largest possible number fraction of $\Sigma 3$ boundaries. This fraction is often taken to be $2/3$, but this result depends strongly on the context, as has been noted (Gertsman & Szpunar, 1998; Gertsman, 2001b). The dimensionality is certainly important, as is the manner in which $\Sigma 1$ boundaries are handled.

Note that the $\Sigma 3/\Sigma 9$ ratio in the allowed pure $\Sigma 3/\Sigma 9$ quadruple node is 1:1 (Fig. 3c), which differs from the result in two dimensions, with a pure $\Sigma 3/\Sigma 9$ twin-limited structure containing nothing but $\Sigma 3$ and $\Sigma 9$ junctions in a 2:1 ratio. So again we see a significant dimensionality effect, so that the twin-limited structures are notably different in two and three dimensions. One should keep in mind, however, that the ratio of boundary types in a three-dimensional structure can differ from the ratio of boundary types in the average quadruple node — even if there is only one type of quadruple node [here is where our interpretation departs from that of Gertsman (2001b)]. A specific construction demonstrating this will be presented below. So it is possible to get a $\Sigma 3/\Sigma 9$ ratio greater than 1:1 in a three-dimensional structure consisting of only $\Sigma 3$ and $\Sigma 9$ boundaries (see the caption of Fig. 3), but a ratio approaching the 2:1 possible in two dimensions would put enormous constraints on the structure and is unlikely to occur in practice (it might be impossible — we have not proven it one way or the other).

Yet it is possible to have $2/3$ of the junctions in a quadruple node be of the $\Sigma 3$ type, as shown in Fig. 3(d), if we include a low-angle $\Sigma 1$ grain boundary. A proper consideration of this case would require development of the concept of small deviations from ideal CSL orientations, which is beyond the scope of this paper but which has been discussed in the literature (*e.g.* Frary & Schuh, 2003). A simple approach is to introduce small-angle rotation matrices and to develop approximation techniques depending on the commutativity to first order of small rotations, which is essentially what is done by Frary & Schuh (2003), although their statistical results are puzzling [*e.g.* the uniform distribution over their limiting tetrahedra, which should not occur with an unbiased selection of deviations over $SO(3)$ in which small rotation angles θ would appear with a probability proportional to θ^2]. $\Sigma 1$ boundaries are sometimes neglected in discussions of grain-

boundary networks, in part because experimentally it is difficult to distinguish a nearly perfect $\Sigma 1$ boundary from localized lattice strain. This example should serve as a reminder that the low-angle boundaries can play an essential role in the network topology.

Allowing the low-angle $\Sigma 1$ boundaries also allows us to construct three-dimensional networks with $\Sigma 3$ fractions above the $2/3$ limit that applies in two dimensions. Fig. 3(f) shows how to assign orientations to a tetrakaidecahedral array to reach a $\Sigma 3$ fraction of $5/7 = 71.4\%$, with the remaining boundaries equally split between $\Sigma 1$ and $\Sigma 9$. The upper limit to the $\Sigma 3$ fraction in a pure $\Sigma 1/\Sigma 3/\Sigma 9$ system is likely to be somewhat more than this.

4.2. Extension beyond $\Sigma 3^{\omega}$

Fig. 4 shows examples of how to draw graphs for more than one prime factor. We start with Fig. 4(a), a graph of $\Sigma 5^{\omega}$ up to $\Sigma 125$. The representative quaternions in this case are chosen to be $a_{\pm 1}^5 = [\pm 2 1 0 0]$, $a_{\pm 2}^5 = [\pm 2 0 1 0]$ and $a_{\pm 3}^5 = [\pm 2 0 0 1]$. This differs from the choice in Table 2 but makes the symmetry somewhat more apparent and also lets us show how to draw graphs in cases in which the a_i^p are not self-inverses. The graph is a directed graph, with each label showing the path followed if we multiply by a_i^5 (in the direction of the arrow) or the inverse a_{-i}^5 (against the arrow), where i now ranges from 1 to 3. The symmetry of the set of representatives of $\Sigma 5$ is more akin to the faces of a cube, compared to the tetrahedral symmetry of the $p = 3$ case where each of the four generators has an equivalent relationship to each of the other generators. The symmetry of the graph for $\Sigma 5^{\omega}$ is noticeably weaker than that for $\Sigma 3^{\omega}$, for instance (as already noted) there are two different types of $\Sigma 5^2$ rotations but only one of $\Sigma 3^2$ rotations. Also, since each vertex has six nodes instead of four (three each incoming and outgoing labelled 1, 2 and 3), the tree grows quite quickly and we only include a small representative sample of the network. The sample is sufficient to reconstruct the entire tree up to $\Sigma 125$ by use of symmetry.

Fig. 4(b) then shows a portion of the graph of $\Sigma 3^{\omega 5^{\omega}}$, that is, all CSL matrices with all prime factors of Σ in the set $\{3, 5\}$. This tree expands even faster than in the $\Sigma 5^{\omega}$ case, and only enough is drawn to show the basic idea. Labels are omitted not only to reduce clutter but to avoid certain complications which we will consider shortly. Essentially, at every node in the $\Sigma 5^{\omega}$ tree (shown in heavy lines), we may place the origin of a $\Sigma 3^{\omega}$ tree. We know from the unique sorted factorization properties that each coset of $\Sigma 3^{\omega 5^{\omega}}$ will appear exactly once in this tree. But we could also have drawn a $\Sigma 3^{\omega}$ tree and put a $\Sigma 5^{\omega}$ tree at every node, and this too would have reached each coset exactly once. This is the graphical representation of the difference between sorting and reverse sorting. What this means is that, from each of the four $\Sigma 3$ nodes, there are six $\Sigma 5$ lines linking it to six of the 24 $\Sigma 15$ nodes. We did not draw these lines in, for two reasons. First, the diagram would have been hopelessly cluttered. Second, the detailed pattern of connectivity of the lines is rather meaningless (although some aspects are invariants), resulting as it does from the arbitrary




Figure 4 Additional examples incorporating $\Sigma 5^{\omega}$. (a) Map of a portion of $\Sigma 5^{\omega}$. Links are labelled with the alternative a_i^5 described in the text, in the direction given by the arrows for $i = 1, 2, 3$ and in the opposite direction for $i = -1, -2, -3$. (b) A small portion of $\Sigma 3^{\omega 5^{\omega}}$. The a_i^5 links are highlighted. (c) Illustrating the ambiguity that makes labelling graphs with more than one prime factor difficult.

choice of representatives a_i^3 and a_i^5 . The pattern would have to be determined in detail for any specific choice of representatives. This would be the graphical representation of equation (8) in the previous section. Equations (9)–(10) have no representation in these kinds of graphs, since they deal with the properties of the symmetry group $\Sigma 1$, which is simplified in the graphs by letting each node represent a coset with respect to $\Sigma 1$.

The impact of the $\Sigma 1$ element appearing in (8) should be considered more closely. Consider Fig. 4(c), a small portion of the $\Sigma 3^{\omega} \Sigma 5^{\omega}$ graph. We show a single $\Sigma 15$ node that may be reached either by taking first a $\Sigma 3$ path and then a $\Sigma 5$ path from the origin, or by taking a $\Sigma 5$ path and then a $\Sigma 3$ path. The two products of pairs of generators will produce the same coset, but will not in general be equal, differing by a factor of some element of $\Sigma 1$ as in (8). This appears to be an inevitable consequence of the non-commutativity of rotations about different axes and the loss of symmetry involved in combining two sets of representatives with different p 's. Essentially, we end up in two different coordinate systems that are equivalent under cubic symmetry as expressed in the target grain. Performing an additional rotation (say, multiplying by a_i^3) will yield a different result depending on which coordinate system is used. Even the Σ number of the result is in question, for one of the a_i^3 paths from our $\Sigma 15$ node might lead to a $\Sigma 45$, the other to a $\Sigma 5$. Algebraically, this ambiguity is expressed as:

$$a_i^3 a_j^5 a_i^3 = a_k^5 a_l^3 S_m a_i^3 \quad (18)$$

and we note that, while $S_m a_i^3$ is guaranteed to be the same type as a_i^3 , it will not always be in the same coset, and so is likely to be represented by some a_i^3 other than that with $i = 1$. So it is in general not possible to label the links on the graph unless we choose a convention, which can be the same convention as before, such that the link between two cosets is the one that would appear in a sorted NIS string representation. The types of grain boundaries between any two nodes will then have to be calculated using the rules in equations (8)–(10).

In short, the graphs for Σp^{ω} are much easier to work with, since the way we are drawing the graphs suppresses the information about which element is used to represent each coset, but this information is important for cases involving more than one prime factor since the same coset may be reached by more than one shortest path in these cases.

We have now seen the graphical analogs to all the essential aspects of the algebraic representation, namely unique sorted and reverse-sorted representations of the cosets, determination of boundary types from grain orientations, the rules for manipulation of strings, and the determination of what boundary type assignments are consistent with crystallographic constraints such as the Σ combination rule (and generalizations of the Σ combination rule to more complex topologies). With the explicit construction of the representatives in Table 2, it is possible to derive all of the algebraic transformations and draw all of the graphs for any desired part of the CSL group. Fortunately, only the elementary transformations such as those in (8) need be calculated for each case in

terms of quaternions or matrices. Once this is complete (as we have done for all physically meaningful cases), and the types of rotations are identified in terms of their conventional notation (e.g. $\Sigma 27a$ versus $\Sigma 27b$ is identified as a particular symmetry in the representative string), then everything may be done in terms of the prime representatives a_i^p and the symmetry elements S_i and there will no longer be any need to deal with matrix multiplications or quaternion algebra. The advantages are quite significant in the important case of $\Sigma 3^{\omega}$, where we only need to consider strings from a four-letter alphabet to do every possible calculation.

5. Summary and conclusions

As the CSL model is generally taken to be just one of the first steps in describing grain boundaries, so the formalism developed in this work can be a starting point for more in-depth investigations. Deviations from the ideal CSL orientations, for instance, can be introduced by including small random rotations, possibly in an approximation scheme taking advantage of the first-order commutativity of small rotations and their random nature. This has the potential to produce reasonable statistical calculations for large topologies without the need to calculate all the rotations in detail. Small intragrain strains (King, 1999) may be dealt with in a similar manner. Generalization of our approach to the O -lattice concept (Bollmann, 1972; MacLaren & Aindow, 1997) may also be of interest, and consideration of non-cubic lattices is also a natural extension of the approach. Unfortunately, without cubic symmetry the natural mapping between integer quaternions and CSL rotations is lost. The related proofs are likely to be correspondingly less accessible.

Instead of generalizing, we can also proceed in the opposite direction, to simpler models which may be more clear in our formalism than they would be otherwise. Significant results have been obtained in the literature from simple considerations of the crystallographic constraints in a small portion of the $\Sigma 3^{\omega}$ group (Fortier *et al.*, 1995; Kumar *et al.*, 2000; Schuh *et al.*, 2003), including the effects of the Σ combination rule and quadruple-node constraints on the statistical behavior of grain-boundary topologies. Since we can now easily represent the structure of the entire $\Sigma 3^{\omega}$ group with extremely simple manipulations of strings with a four-letter alphabet, the generalization of such results should be greatly facilitated. It is easy to generate an exhaustive list of all the quadruple nodes with all- $\Sigma 3^{\omega}$ boundaries up to a fixed limit, for instance. Thus we will avoid both neglecting real possibilities (the classic example being the $\Sigma 9$ – $\Sigma 9$ – $\Sigma 9$ triple junction, which is neglected almost as a matter of course in some segments of the literature) and wasting time on apparent possibilities that are in fact crystallographically inconsistent [such as the $4\Sigma 3/2\Sigma 9$ quadruple node in Fig. 3(e)].

The abstract concept of the constraining effect of the group connectivity on the grain-boundary network is an interesting and important notion that can be very confusing at first glance. Considered in terms of mapping one graph to another such as in Fig. 4, the idea becomes much more intuitive and accessible,

and some results that were puzzling or poorly understood in the literature become more obvious in light of these simple pictures. In addition to clarifying old results, we have made observations regarding connections among the CSL group structure, the dimensionality and colorability of graphs, and the nature of twin-limited structures in different topologies. We have every hope that our approach will facilitate rapid progress in the statistical understanding of grain-boundary networks.

Thanks to David Barber for useful discussions on the factorization properties of quaternions. This work was performed under the auspices of the US Department of Energy by University of California, Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

References

- Appel, K. & Haken, W. (1989). *Every Planar Map is Four-Colorable*. Providence, RI: American Mathematical Society.
- Barber, D. (2003). Personal communication.
- Bollmann, W. (1972). *Surf. Sci.* **31**, 1–11.
- Conway, J. H. & Smith, D. A. (2003). *On Quaternions and Octonions: their Geometry, Arithmetic, and Symmetry*, pp. 23–24, 55–63. Natick, MA: A. K. Peters Ltd.
- Fortier, P., Aust, K. T. & Miller, W. A. (1995). *Acta Metall. Mater.* **43**, 339–345.
- Frary, M. & Schuh, C. A. (2003). *Acta Mater.* **51**, 3731–3743.
- Gertsman, V. Y. (2001a). *Acta Cryst.* **A57**, 369–377.
- Gertsman, V. Y. (2001b). *Acta Cryst.* **A57**, 649–655.
- Gertsman, V. Y. (2002). *Acta Cryst.* **A58**, 155–161.
- Gertsman, V. Y. & Szpunar, J. A. (1998). *Scr. Mater.* **38**, 1399–1404.
- Gertsman, V. Y. & Tangri, K. (1995). *Scr. Metall. Mater.* **32**, 1649–1652.
- Gertsman, V. Y., Tangri, K., & Valiev, R. Z. (1994). *Acta Metall. Mater.* **42**, 1785–1804.
- Goldstein, H. (1950). *Classical Mechanics*, pp. 100–101. New York: Addison-Wesley.
- Gottstein, G. (1984). *Acta Metall.* **32**, 1117–1138.
- Grimmer, H. (1973). *Scr. Met.* **7**, 1295–1300.
- Grimmer, H. (1974). *Acta Cryst.* **A30**, 685–688.
- Grimmer, H. (1976). *Scr. Met.* **10**, 387–391.
- Grimmer, H. (1984). *Acta Cryst.* **A40**, 108–112.
- Grimmer, H., Bollmann, W. & Warrington, D. H. (1974). *Acta Cryst.* **A30**, 197–207.
- Hardy, G. H. & Wright, E. M. (1938). *An Introduction to the Theory of Numbers*, pp. 300–316. Oxford: Clarendon Press.
- Iwasaki, Y. (1976). *Acta Cryst.* **A32**, 59–65.
- King, A. (1999). *Interface Sci.* **7**, 251–271.
- Kopezky, C. V., Andreeva, A. V. & Sukhomlin, G. D. (1991). *Acta Metall. Mater.* **39**, 1603–1615.
- Kumar, M., King, W. E. & Schwartz, A. J. (2000). *Acta Mater.* **48**, 2081–2091.
- MacLaren, I. & Aindow, M. (1997). *Philos. Mag. Lett.* **76**, 25–32.
- Minich, R. W., Schuh, C. A. & Kumar, M. (2002). *Phys. Rev. B*, **66**, 052101-1–4.
- Miyazawa, K., Iwasaki, Y., Ito, K. & Ishida, Y. (1996). *Acta Cryst.* **A52**, 787–796.
- Mykura, H. (1979). *Grain Boundary Structure and Kinetics*, pp. 435–456. Metals Park, OH: American Society for Metals.
- Palumbo, G., Aust, K. T., Erb, E., King, P. J., Brennenstuhl, A. M. & Lichtenberger, P. C. (1992). *Phys. Status Solidi A*, **131**, 425.
- Randle, V. (1994). *Acta Metall. Mater.* **42**, 1769–1784.
- Randle, V. (1996). Editor. *The Role of the Coincidence Site Lattice in Grain Boundary Engineering*. London: Institute of Materials.
- Schuh, C. A., Kumar, M. & King, W. E. (2002). *Acta Mater.* **51**, 687–700.
- Schuh, C. A., Minich, R. W. & Kumar, M. (2003). *Philos. Mag.* **83**, 711–726.