Three-dimensional Euclidean nets from two-dimensional hyperbolic tilings: kaleidoscopic examples

S. J. Ramsden, V. Robins and S. T. Hyde*

Department of Applied Mathematics, Research School of Physics and Engineering, Australian National University, Canberra, ACT 0200, Australia. Correspondence e-mail: stephen.hyde@anu.edu.au

We present a method for geometric construction of periodic three-dimensional Euclidean nets by projecting two-dimensional hyperbolic tilings onto a family of triply periodic minimal surfaces (TPMSs). Our techniques extend the combinatorial tiling theory of Dress, Huson & Delgado-Friedrichs to enumerate simple reticulations of these TPMSs. We include a taxonomy of all networks arising from kaleidoscopic hyperbolic tilings with up to two distinct tile types (and their duals, with two distinct vertices), mapped to three related TPMSs, namely Schwarz's primitive (P) and diamond (D) surfaces, and Schoen's gyroid (G).

1. Introduction

The science of networks has grown beyond the inherent mathematical interest of graph theory, thanks to the intimate connection between complex systems and networks. Those connections are now a prime focus of statistical physics, where the study of random networks has been catalysed by interesting results concerning 'small-world' and 'scale-free' systems (Watts, 2003; Barabasi, 2002). Complementary to these endeavours is the study of 'crystal nets': triply periodic embeddings of graphs in three-dimensional Euclidean space. Interest in crystal nets stems from their fundamental relevance to condensed materials (Hyde *et al.*, 2008) from atomic crystals (Wells, 1977; O'Keeffe & Hyde, 1996; Klee, 2004) to novel framework materials (Ockwig *et al.*, 2005; Öhrström & Larsson, 2005), including carbon polymorphs (Strong *et al.*, 2004), zeolites (Treacy *et al.*, 1997; International Zeolite Association, 2008), related oxide (Zou *et al.*, 2008) and alumino-phosphate (AlPO) materials (Li *et al.*, 2008), imidazolate (ZIF) frameworks (Banerjee *et al.*, 2008) and metal-coordination polymeric materials (Blatov *et al.*, 2004; Batten, 2001). The microdomains of some liquid-crystalline phases of soft materials, including amphiphile and copolymer assemblies and derivative mesoporous solids, are also characterized by crystal nets (Hyde & Schroeder, 2003).

This paper is concerned with the enumeration of crystal nets (or just 'nets'), a topic that has gained momentum in recent years, helped by advances in tiling theory that we also exploit here. The technique we use is complementary to a more conventional approach that employs three-dimensional tiling theory (Delgado-Friedrichs *et al.*, 1999; Blatov *et al.*, 2007), in that we build the nets largely within non-Euclidean (hyperbolic) two-dimensional space, then project to three-dimensional (Euclidean) space. Any systematic enumeration of crystal nets is constrained by the issue of combinatorial explosion; distinct approaches therefore explore different regions of the universe of crystal nets. Indeed, our results are largely unduplicated by three-dimensional techniques (Hyde *et al.*, 2006).

The underlying methodology of our approach is to decorate a surface with all allowed symmetric tilings of a given complexity, then use the topology of the tile boundaries and their embedding in the surface to define a net. To guide the reader through these various stages of net building, we adopt the convention that all tilings, composed of vertices, edges and faces, are denoted by upper-case names, while nets, composed of vertices and edges only, have lower-case names.

Our enumeration is constrained by considering a particular class of surfaces to reticulate. For various reasons that will be described in the paper, we choose triply periodic minimal surfaces (TPMSs), which give a natural bridge between two-dimensional hyperbolic tilings and three-dimensional Euclidean nets. A particular choice of TPMSs provides a filter with which to control the enumeration both in two-dimensional hyperbolic and three-dimensional Euclidean space. Here we use the simplest (cubic genus-3) TPMS, namely Schwarz's primitive (P) and diamond (D) surfaces, and Schoen's gyroid (G).

Tilings of the TPMSs, called *E-tilings*, may be lifted to their universal covering space, the two-dimensional hyperbolic plane $\mathbb{H}^2$. Tilings there, that we call *U-tilings*, give rise to two-dimensional hyperbolic nets that we call *h-nets*. Alternatively we can form a compact quotient space of the TPMS, where primitive lattice translations are made equivalent to the identity operation, and the resulting surface is the closed genus-3 *tritorus*. Tilings on this finite surface define *O-tilings*;

Table 1
Definitions of the classes of tilings and nets considered here.

Abbreviation	Term	Embedding space	Definition
U-tiling	Universal cover tiling	$\mathbb{H}^2$	Two-dimensional hyperbolic tilings within specific surface-compatible symmetry subgroups and translational unit cell. These tilings are represented by unfolded D symbols augmented with 'cuts'.
h-net	Hyperbolic net	$\mathbb{H}^2$	Two-dimensional hyperbolic nets, distinguished by topology only, represented by the unique maximal-symmetry D symbol with that topology.
O-tiling	Tritorus surface reticulation	Tritorus	Tiling obtained by projecting a U-tiling onto the tritorus.
o-net	Finite net	$\mathbb{E}^3$	Net formed by the edges of an O-tiling embedded in three-dimensional space. Such nets are finite 'molecular' nets and equivalent up to ambient isotopy class.
E-tiling	TPMS surface reticulation	TPMS	Tiling obtained by projecting a U-tiling onto a periodic minimal surface. Such tilings are identified by pairing a U-tiling with a covering map.
e-net	'Epinet'	$\mathbb{E}^3$	Net formed by vertices and edges of the E-tiling, equivalent up to ambient isotopy class.
s-net	'Systre net'	$\mathbb{E}^3$	Three-dimensional periodic net, defined by topology only, represented by the systre key.

subsequent embedding of the tritorus in three-dimensional space leads to a specific threading of edges, forming the *o-net*. Finally, the tilings of a TPMS can give rise to at least two different net types, depending on how much embedding information is used. A net embedded in the TPMS (following the edges of the E-tiling) will share its three-dimensional symmetry, and we label these surface reticulations *e-nets*. Ignoring the surface, and looking only at network topology, we can relax the net in three-dimensional space ($\mathbb{E}^3$) to give canonical, maximally symmetric forms called *s-nets*.

Here we explain in detail the process of generating the various tilings and nets. For ease of reference, we summarize these in Table 1 and display key steps of the enumeration in the flowchart of Fig. 1. A specific example is illustrated in Fig. 2.

In order to illustrate hyperbolic tilings, we require a model of the hyperbolic plane, that – like the surface of a sphere – cannot be readily imaged on the page without distortions. The hyperbolic plane is illustrated in this paper using the Poincaré disc model (Stillwell, 1989). The vast area of hyperbolic space

Figure 1
A flowchart depicting the enumeration process and the relationships between different classes of tilings and nets considered here. 1: Unfold the D symbol within a specific subgroup chart (many-to-many) (§4). 2: Find the minimal-image (maximal-symmetry) D symbols (many-to-one) (§6.1). 3: Map the hyperbolic tiling onto a tritorus (one-to-one) (§5.3). 4: Embed the tritorus, remove faces, leaving edges in three-dimensional space (§6.4). 5: Map the hyperbolic tiling onto a triply periodic minimal surface (one-to-one) (§5.1). 6: Remove faces, leaving embedded edges and vertices (many-to-one) (§6.2). 7: Find the canonical, maximally symmetric form for the e-net topology (many-to-one) (§6.3).

is compressed into a unit disc at the expense of significant foreshortening of distances, that are shrunk more and more towards the perimeter of the disc. This model is attractive in that it is conformal, with no distortion of angles induced by the map from hyperbolic space to the disc; however, hyperbolic lines are imaged as circular arcs that intersect the disc boundary at 90°. The remainder of this introductory section provides an outline of the paper.

The observation that underlies our enumeration scheme is that the intrinsic geometry and symmetry of TPMSs are related to discrete groups of two-dimensional hyperbolic isometries. It is therefore possible to wrap the hyperbolic plane onto a TPMS in almost the same way as the Euclidean plane wraps onto a cylinder. This wrapping is formally defined by a *covering map*, described in detail for the primitive (*P*), diamond (*D*) and gyroid (*G*) surfaces in §2. The determination of suitable covering maps is an essential component of the enumeration process. Further, among the infinite variety of two-dimensional hyperbolic tilings, we choose those examples whose symmetries in $\mathbb{H}^2$ are commensurate with both the local and translational symmetries of the TPMS. Those allowed symmetries were presented in an earlier paper (Robins *et al.*, 2004a). In the current paper we focus on a subset of these: the *kaleidoscopic* groups which are generated by reflections only.

The enumeration of suitable hyperbolic tilings involves a straightforward application of combinatorial tiling theory as developed by Dress, Huson & Delgado-Friedrichs (Dress, 1987; Dress & Huson, 1987; Delgado-Friedrichs, 1994). We give an overview of their techniques in §3; technical definitions are summarized in Appendix B.

In §4 we describe how to embed and unfold these hyperbolic tilings in $\mathbb{H}^2$ so they are compatible with the symmetries of the covering maps. This process involves a number of steps that extend Delaney–Dress theory to account for equivalent tilings on topologically complex TPMSs. This is the technical core of the current paper and involves both combinatorial and group-theoretic algorithms. We then discuss, in §5.1, how these hyperbolic tilings project onto a TPMS, and present the space-group symmetries of the resulting surface reticulations.

In §6 we discard the tiling information and examine the topological structure of the nets defined by edges and vertices.
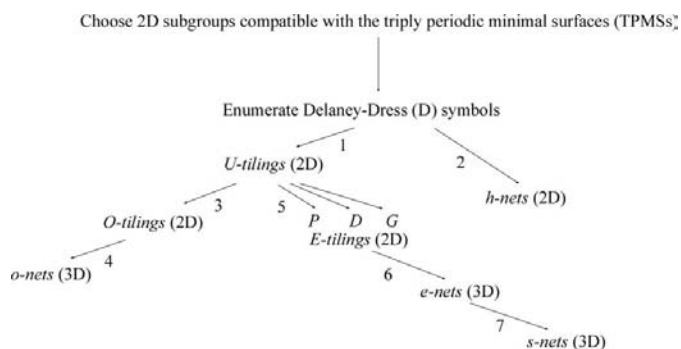
To further clarify our techniques, we illustrate the procedure from start to finish by a fully worked example in §7.

The output of the enumeration process for kaleidoscopic tilings is summarized in §§6 and 9, with emphasis on the variety of e- and s-nets that result. These results are far too extensive to capture in any detail in a publication. Accordingly, we have built the online *Epinet* database (Ramsden *et al.*, 2005), which has detailed searchable catalogues of data and images linking tilings and nets.

In the future, this project will grow to explore non-kaleidoscopic tilings, in directions outlined in §8. Accordingly, we have written this paper to provide a detailed foundation on which an evolving suite of network data will be built. Although foundations invariably make dull reading, they are needed to explicate a richer and more interesting superstructure. The reader is therefore urged to read this paper in conjunction with the *Epinet* database, accessible at http://epinet.anu.edu.au.
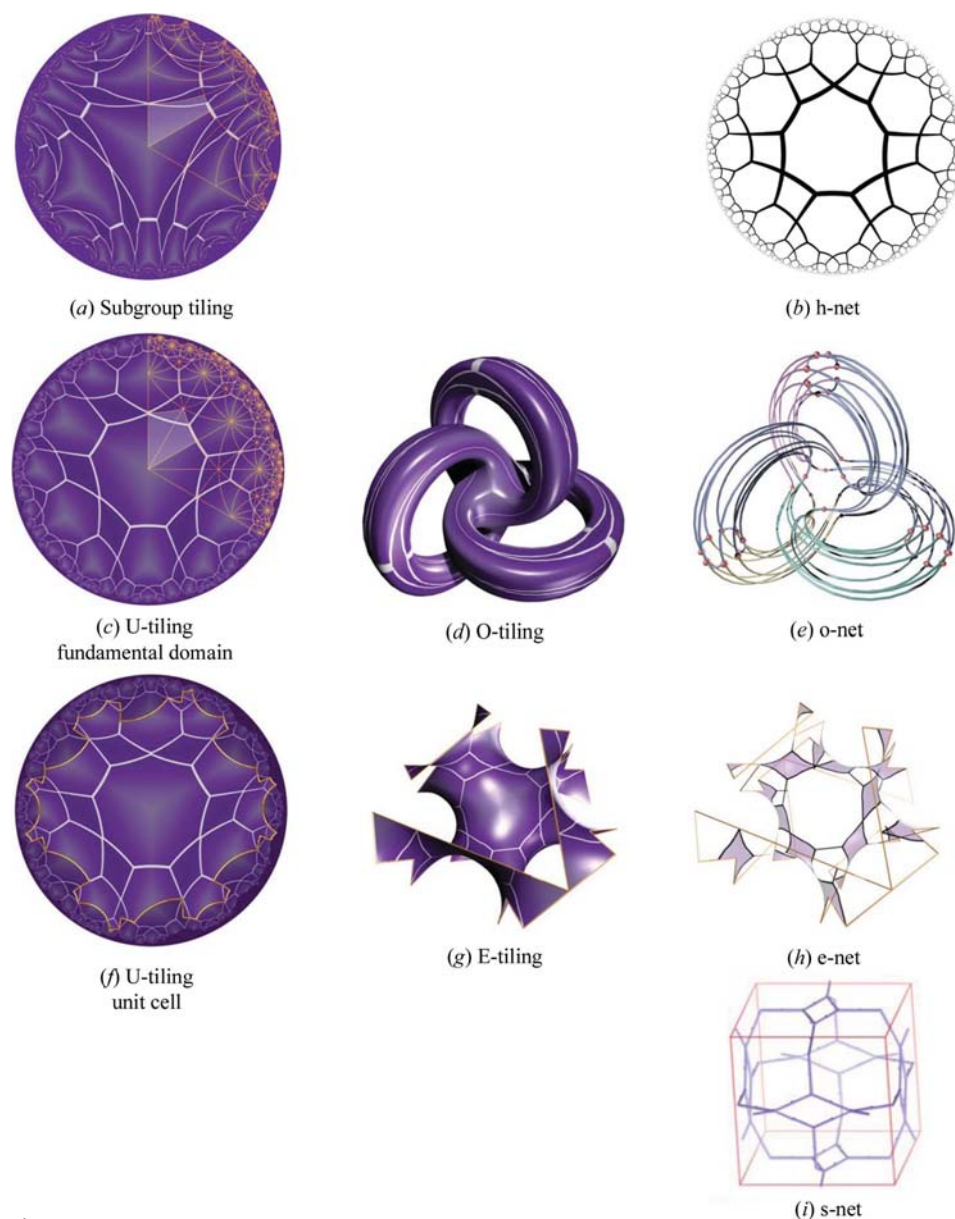
2. Hyperbolic symmetries and the covering maps

Note. Throughout this paper we refer to two-dimensional symmetry groups by their orbifold symbol, a notation introduced by Conway (1992) and described in Appendix A.

The *P*, *D* and *G* surfaces (illustrated in Figs. 3 and 4) each have intrinsic surface symmetry related to the ⋆246 group of hyperbolic reflections. This group is generated by three reflections, R_1, R_2 and R_3, whose mirror lines bound a triangle in $\mathbb{H}^2$ with corner angles of $\pi/4$, $\pi/6$ and $\pi/2$. This geometry induces a set of relations for the group,

$$(R_1R_2)^2 = (R_2R_3)^4 = (R_1R_3)^6 = I,$$

and because the operations are reflections we also have $R_1^2 = R_2^2 = R_3^2 = I$ (the identity). Sadoc & Charvolin (1989) identified translational unit cells of the oriented *P*, *D* and *G* surfaces that pull back to the same dodecagon in the hyperbolic plane. These primitive rhombohedral unit cells and the corresponding $\mathbb{H}^2$ dodecagon are shown in Figs. 3 and 5. Euclidean translation vectors for each surface pull back to the hyperbolic group generated by the six translations that pair opposite edges of the dodecagon. These translations were originally given in Sadoc & Charvolin (1989) and are rewritten

Figure 2
This table of images presents the various stages in our map from the hyperbolic plane (top) to Euclidean space (bottom). (*a*) Delaney–Dress symbol enumeration with ⋆2223 symmetry gives a subgroup tiling, *QC*164, containing five distinct flags. ⋆2223 domains are bounded by orange edges; a single domain is shaded. (*b*) The h-net, *hqc*167, is a maximally symmetric version of the net defined by the edges and vertices of the tiling *QC*164. (*c*) We geometrize the subgroup tiling to form a *U-tiling*, *UQC*183, that respects the local and translational symmetries of the TPMS (in this paper the *P*, *D* and *G* surfaces). The single ⋆2223 (shaded) domain from (*a*) adopts the geometry of a pair of ⋆246 triangles. (*d*) The U-tiling is projected onto a genus-3 torus (tritorus), giving an *O-tiling*. (*e*) The edges of the embedded surface reticulation define an *o-net*. (*f*) A hyperbolic translational cell built from the ⋆2223 domains of (*c*) is highlighted by the orange polygon. (*g*) This polygon projects onto the *D* surface to give an E-tiling, *EDC*183, whose rhombohedral unit cell is shown. (*h*) The projected tile boundaries on the *D* surface form the triply periodic *e-net*, *edc*183, with space group *P*4₂32. A single unit cell is shown, with quadrilateral cycles highlighted in purple for clarity. (*i*) The *s-net*, *sqc*7388, is the highest-symmetry embedding of the e-net topology, with space group *Pm*3̄*n*.

here in terms of the $\star 246$ reflections:

$$t_1 = (R_3R_1R_3R_1R_3R_2)^2,$$
$$t_2 = R_1R_3R_2(R_3R_1R_3R_1R_3R_2)R_3R_1R_3,$$
$$t_3 = R_3R_2(R_3R_1R_3R_1R_3R_2)R_3R_1R_3R_1,$$
$$\tau_1 = (R_3R_1R_2R_3R_1R_2R_3R_1)^2,$$
$$\tau_2 = R_1R_3R_1(R_3R_1R_2R_3R_1R_2R_3R_1)R_3R_1R_2R_3R_2,$$
$$\tau_3 = R_1R_2R_3R_1R_2R_3R_1(R_3R_1R_2R_3R_1R_2R_3R_1)R_3.$$

They satisfy the following identity:

$$\tau_1 t_2 \tau_3^{-1} t_1^{-1} \tau_2 t_3 \tau_1^{-1} t_2^{-1} \tau_3 t_1 \tau_2^{-1} t_3^{-1} = I. \tag{1}$$

Figure 3
Translational unit cells for the P, D and G minimal surfaces. In the left column we show primitive rhombohedral unit cells that are the image of the hyperbolic dodecagon built from 96 triangles (shown in Fig. 5). These primitive unit cells are for the space groups of the oriented surface in each case, *i.e.*, the surface normal direction (and the colouring of both sides) is preserved by the space-group symmetries. In the right column we show $2 \times 2 \times 2$ conventional cubic unit cells for each surface, with the rhombohedral cell inset. The conventional cubic unit cells are those of the space groups of the non-oriented unit cells, which include symmetries that swap sides of the surface.

The subgroup T generated by the t_i and τ_i translations is isomorphic to the fundamental group of a genus-3 torus and therefore has orbifold symbol ∘∘∘.

A covering map, Φ, from $\mathbb{H}^2$ onto a surface in $\mathbb{E}^3$, is a continuous function such that any small disc on the surface pulls back to a countably infinite number of isomorphic copies that form a periodic pattern in the hyperbolic plane. The covering maps we construct have the additional properties of being conformal and compatible with the symmetries of the surface. This means that given a Euclidean symmetry of the surface $S : \mathbb{E}^3 \to \mathbb{E}^3$, there is a corresponding symmetry of the hyperbolic plane, $s : \mathbb{H}^2 \to \mathbb{H}^2$, such that $\Phi \cdot s = S \cdot \Phi$. Thus, the covering map defines a relationship (a group homomorphism) between the hyperbolic symmetry group $\star 246$ and the Euclidean space group of the surface.

In the case of the P surface, the hyperbolic reflections R_1 and R_2 map to Euclidean reflections in (110) and (100) mirror planes, respectively, whilst R_3 maps to a twofold axis of rotation in the [110] directions, lying in the surface (and therefore swaps the sides of the surface). In the D surface R_1 and R_2 map to twofold rotation axes in the [110] and [100] directions, in the surface, and R_3 maps to a reflection in the (110) plane. The Euclidean symmetries of the G surface are related to the rotational subgroup 246 of $\star 246$. While the hyperbolic reflections are intrinsic local symmetries of the G surface, they do not correspond to Euclidean isometries of the entire surface. Instead, the rotation R_1R_2 maps to a twofold axis in the [110] direction, perpendicular to the G surface, R_1R_3 maps to a $\bar{3}$ inversion centre at Wyckoff sites 16a of the space group $Ia\bar{3}d$ in the [111] direction, and R_2R_3 maps to a $\bar{4}$ inversion at Wyckoff sites 24d in the [001] direction. These symmetries are indicated in the surfaces in Fig. 4. Further details can be found elsewhere (Robins *et al.*, 2005; Molnar, 2002).

The $\mathbb{H}^2$ translation subgroup T, introduced above, maps onto the Euclidean translation lattice of each oriented surface space group. We can summarize the action of a covering map by describing how the six generators of T map to three
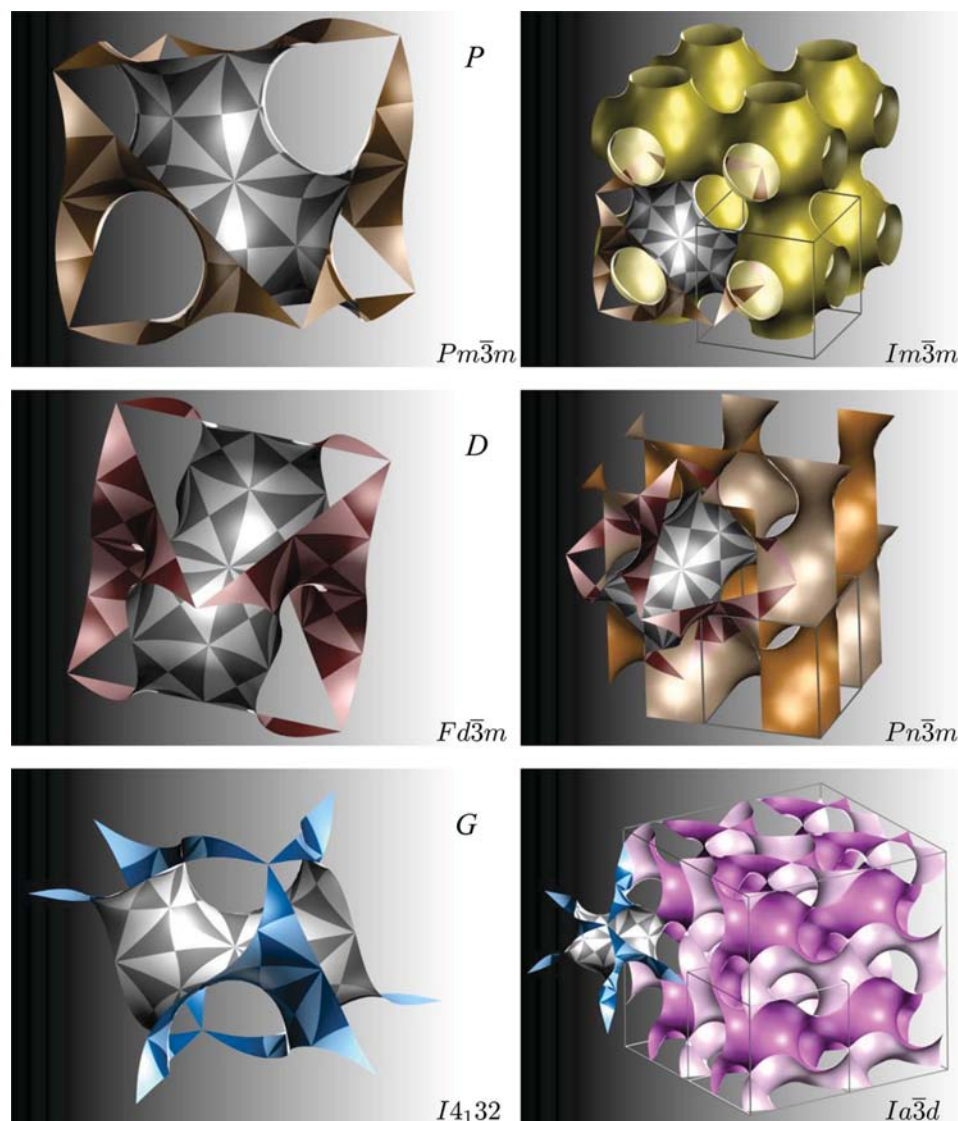
independent Euclidean translations **a**, **b**, **c**. For the P, D and G surfaces we choose the mappings given in Table 2 (other choices are possible and give identical results).

A tiling of the hyperbolic plane will map to a discrete reticulation on a surface only when the symmetries of the tiling are commensurate with the chosen covering map. This means that the tiling should have at least the symmetry of the *fundamental group* of the surface [*i.e.*, the first-order homotopy group $\pi_1(\mathcal{S})$]. Since the P, D and G surfaces have unbounded genus, their fundamental groups have infinitely many generators. These groups are isomorphic, *via* the covering maps, to distinct subgroups of T. A tiling with only
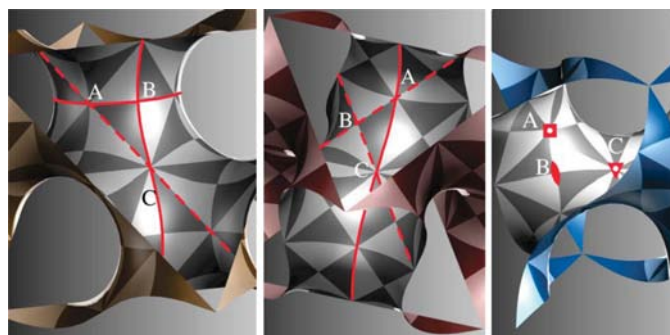

Figure 4
Magnified views of fragments of the P, D and G surfaces (see Fig. 3), with three-dimensional Euclidean symmetry operations marked on the surfaces. The arcs AB, BC and CA correspond to the two-dimensional in-surface (non-Euclidean) reflections R_2, R_1 and R_3, respectively, for all three surfaces (see main text). Left: The P surface, with $Im\bar{3}m$ symmetry (including isometries that swap sides of the surface; these are coloured differently here for clarity): A, B and C are located at the 12d (site symmetry $\bar{4}m2$), 24h ($mm2$) and 8c ($\bar{3}m$) sites, respectively; AB and BC lie in orthogonal mirror planes, CA is a twofold axis. Middle: The D surface ($Pn\bar{3}m$): A, B and C are located at the 6d ($\bar{4}2m$), 12f (222) and 4c ($\bar{3}m$) sites, respectively; AB and BC are twofold axes, CA lies in a mirror plane. Right: The G surface ($Ia\bar{3}d$): A and C are located at the 24d and 16a sites, with point-group symmetries $\bar{4}$ and $\bar{3}$, respectively; B lies on a twofold axis that is normal to the surface.
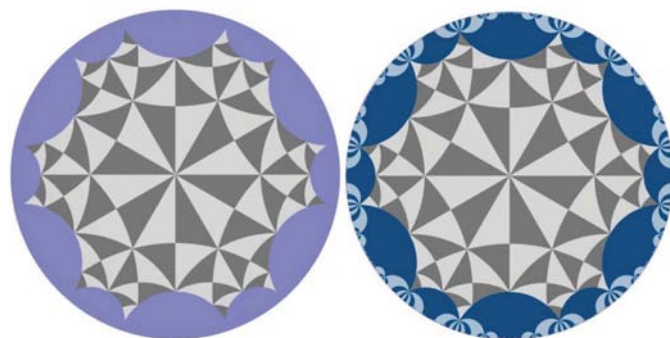

Figure 5
The dodecagonal region of the hyperbolic plane that forms the common preimage for the primitive rhombohedral unit cells of the P, D and G surfaces. The dark and light triangles are $\star$246 domains – there are 96 such triangles in the dodecagonal translation unit. The translations that pair opposite sides of the dodecagon generate an irregular (12, 12) tiling of the hyperbolic plane, illustrated to the right. Eleven additional dodecagons (alternating pale and bright blue) meet at each vertex of the central triangulated dodecagon.

the symmetry of $\pi_1(P)$, say, will map to an aperiodic decoration of the P surface, and will be incompatible with the D and G surfaces. In this paper we study hyperbolic tilings with the full T symmetry, because these tilings are compatible with all three surfaces, and generate triply periodic patterns in $\mathbb{E}^3$. In principle, we can use the T subgroup as the starting point for tiling enumeration, but for convenience we enumerate tilings within symmetry groups that lie between $\star$246 and T; we established earlier that 131 distinct subgroups are suitable (including $\star$246 and T) (Robins *et al.*, 2004a).

Among the 131 subgroups of $\star$246, we have enumerated so far those tilings whose symmetries are generated by reflections only, namely the *kaleidoscopic* subgroups. There are 14 such subgroups, specified by the generators given in Table 3. The subgroup relationships are represented in Fig. 6. The kaleidoscopic subgroups are examples of a much broader class of groups known to mathematicians as *Coxeter groups*. To conform with a forthcoming classification scheme for families of two-dimensional orbifolds, we refer to the orbifolds of kaleidoscopic groups as *Coxeter orbifolds*. Among the 14 kaleidoscopic subgroups there are only 11 different Coxeter orbifolds. The results presented in this paper are derived from these kaleidoscopic subgroups, but the concepts and techniques apply (with a little modification) to subgroups with other types of symmetry.

In the following §3 we give an overview of the relevant algorithms from combinatorial tiling theory. We then discuss in §4 how to take these abstract representations and generate tilings of $\mathbb{H}^2$ that are compatible with the surface-covering maps.

3. Enumerating Delaney–Dress symbols of hyperbolic tilings

The study of tilings has a long history (Grünbaum & Shephard, 1987), but it was only in the 1980s that a finite symbol was found to uniquely encode both the topology and symmetry of an infinite periodic tiling. This result is due to Andreas Dress (Dress, 1987) using earlier work of Matthew Delaney, so we refer to this description of a tiling as a *Delaney–Dress symbol*. Two students of Dress, Daniel Huson and Olaf Delgado-Friedrichs, have extended his mathematical formalism and developed efficient algorithms for enumerating tilings of the sphere, Euclidean and hyperbolic planes, and three-dimensional Euclidean space (Huson, 1993; Delgado-Friedrichs, 2003; Delgado-Friedrichs & Huson, 1999). This body of work is usually referred to as *combinatorial tiling theory*.

Both the symmetry and topology of a two-dimensional tiling are efficiently encoded by a two-dimensional Delaney–Dress symbol, described in detail in Appendix B. In effect, these symbols represent a triangulation of an orbifold. Thus, the enumeration of two-dimensional tilings amounts to the systematic enumeration of all possible triangulations (with certain properties) of a given orbifold.

Delaney–Dress (D) symbols allow exhaustive enumeration of tilings within a symmetry group up to a desired level of

Table 2
The mappings from $\mathbb{H}^2$ translations to Euclidean lattice vectors for the oriented P, D and G surfaces.

For the P surface the vectors form an orthonormal set. For the D surface $\{\mathbf{a}, \mathbf{b}, \mathbf{c}\}$ are face-centred-cubic (f.c.c.) rhombohedral lattice vectors that meet at $60°$ and have length $1/(2^{1/2})$. The translational lattice for the G surface is body-centred-cubic (b.c.c.), and the rhombohedral lattice vectors meet at angles of $109.47°$ and have length $(3^{1/2})/2$ .The given lengths are with respect to the standard (nonprimitive) cubic unit cell in each case.

Surface	t_1	t_2	t_3	τ_1	τ_2	τ_3
P (cubic)	$\mathbf{a}$	$\mathbf{b}$	$\mathbf{c}$	$\mathbf{c} - \mathbf{b}$	$\mathbf{a} - \mathbf{c}$	$\mathbf{b} - \mathbf{a}$
D (f.c.c.)	$\mathbf{a} - \mathbf{c}$	$\mathbf{b} - \mathbf{a}$	$\mathbf{c} - \mathbf{b}$	$-\mathbf{b}$	$-\mathbf{c}$	$-\mathbf{a}$
G (b.c.c.)	$-\mathbf{b}$	$-\mathbf{c}$	$-\mathbf{a}$	$\mathbf{b} + \mathbf{c}$	$\mathbf{a} + \mathbf{c}$	$\mathbf{a} + \mathbf{b}$

complexity. The complexity is initially quantified by the number of distinct symmetry classes of tile, also called the transitivity class. Thus, a tiling in which every tile is related to every other tile by a symmetry of the entire tiling is *tile-1-transitive*. A tiling with k symmetrically distinct tiles is *tile-k-transitive*. The transitivity class of a tiling is incremented by *splitting* all symmetric copies of a chosen tile. The split operations we allow involve adding a single edge to the interior of a tile – either between a vertex and a non-incident edge, between two different edges or between two non-incident vertices.

The second distinction in complexity concerns the internal symmetry of a tile. A tile is *fundamental* if no symmetry of the entire tiling maps it onto itself. If the tile has internal symmetries it is by definition *nonfundamental*. A nonfundamental tile is obtained by *gluing* copies of a fundamental tile. A gluing is possible across an edge of a fundamental tile if the
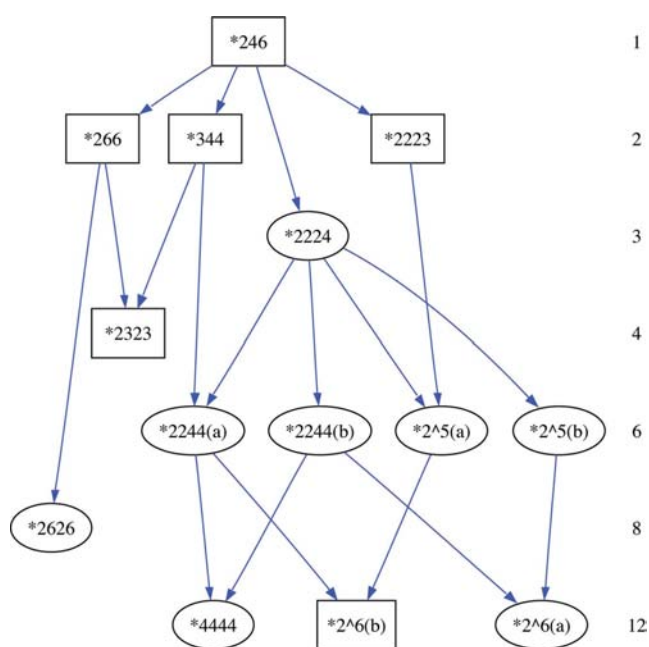

Figure 6
The maximal subgroup lattice for the 14 kaleidoscopic subgroups of ⋆246. The rectangular nodes denote a normal subgroup and ellipsoidal nodes denote a conjugacy class of subgroups. The numbers in the right-hand column are the subgroup index in ⋆246.

edge coincides with a mirror line, or has a centre of twofold rotational symmetry at its midpoint. Alternatively, copies of the fundamental tile may be glued around a vertex if it is the fixed point of a rotational symmetry (*i.e.* at a cone or corner point of the orbifold).

Following Huson (1993), we enumerate Delaney–Dress symbols in the following order. Our enumeration differs slightly from that in Huson (1993) and in Balke & Huson (1996), because we exclude tiles with less than three edges (for reasons given below).

(1) Start with the tile-1-transitive fundamental tilings with the given orbifold. These tilings are called F-tilings. Coxeter orbifolds (those of kaleidoscopic groups) have only one such F-tiling; all other orbifold families have more than one.

(2) Determine all additional tile-1-transitive tilings by applying glue operations to the F-tilings to obtain the FG-tilings. Some example FG tilings are shown in Fig. 7.

(3) Determine all tile-2-transitive fundamental tilings by applying split operations to each F-tiling. The resulting cases are FS-tilings.

(4) Determine all possible tile-2-transitive nonfundamental tilings by applying glue operations to the FS-tilings. The results of a glue operation applied to a single tile class are called FSG-tilings. When glue operations are applied to both tile classes, we get an FSGG-tiling. See Fig. 8 for an illustration of these split and glue operations.

(5) Continue with tile-3-transitive fundamental tilings generated by two split operations, followed by up to three glue operations, and so on.

The above procedure forms tilings ordered by tile-transitivity class. In Table 4, we give the number of D symbols for tile-1- and tile-2-transitive tilings within each of our 11 Coxeter orbifolds.

A simple way to generate additional tilings without further split or glue operations is to introduce the *dual operation* defined by replacing each tile by a vertex and joining these vertices by an edge if the original tiles are adjacent. A number of examples are given in Fig. 9. The dual tiling shares the symmetry of the original tiling. In addition, it has reversed two-dimensional topological characteristics: the degree of a dual vertex is equal to the number of edges in the original polygonal tile and a dual tile is a polygon of order equal to the degree of the original vertex. This is why we require our tiles to have at least three edges: so that the dual vertices are at least degree three.

We do not quite double the number of D symbols by appending the vertex-k-transitive duals to the tile-k-transitive symbols. This is because some tilings are *self dual*, while other symbols may be paired as *mutually dual*; examples are given in Fig. 9. From our lists of tile-1-transitive tilings in the 11 Coxeter orbifolds we find three self-dual D symbols: a regular $\{6, 6\}$ tiling (in ⋆266), a semi-regular $(6, 6)$ tiling (⋆2323), and a regular $\{12, 12\}$ tiling (⋆2626).[1] Among the tile-2-transitive tilings we find 17 self-dual symbols. A pair of D symbols from

---

[1] The Schläfli symbol (p, q) denotes a tiling by p-gons meeting at vertices of degree q; the symbols $\{p, q\}$ denote tilings by regular polygons.

Table 3
The 14 kaleidoscopic symmetry groups compatible with the P, D and G surfaces.

No. in Robins et al. (2004a)	Subgroup name	Index in $\star 246$	Subgroup generators
131	$\star 246$	1	R_1, R_2, R_3
127	$\star 266$	2	$R_1, R_3, R_2R_3R_2$
125	$\star 344$	2	$R_2, R_3, R_1R_3R_1$
124	$\star 2223$	2	$R_1, R_2, R_3R_1R_3, R_3R_2R_3$
123	$\star 2224$	3	$R_1, R_2, R_3R_2R_3, R_3R_1R_3R_1R_3$
119	$\star 2323$	4	$R_3, R_1R_3R_1, R_2R_3R_2, R_1R_2R_3R_2R_1$
107	$\star 2244(a)$	6	$R_2, R_3R_2R_3, R_3R_1R_3R_1R_3, R_1R_3R_2R_3R_1$
103	$\star 2244(b)$	6	$R_1, R_3R_2R_3, R_3R_1R_3R_1R_3, R_2R_3R_1R_3R_1R_3R_2$
102	$\star 2^5(a)$	6	$R_1, R_2, R_3R_2R_3, R_3R_1R_3R_2R_3R_1R_3, R_3R_1R_3R_1R_3R_2R_3R_1R_3R_1R_3$
101	$\star 2^5(b)$	6	$R_1, R_2, R_3R_1R_3R_1R_3, R_3R_2R_3R_1R_3R_2R_3, R_3R_2R_1R_3R_1R_2R_3$
96	$\star 2626$	8	$R_1, R_3, R_2R_3R_1R_3R_1R_3R_2, R_2R_3R_1R_2R_3R_1R_3R_2R_1R_3R_2$
64	$\star 4444$	12	$R_3R_2R_3, R_3R_1R_3R_1R_3, R_1R_3R_2R_3R_1, R_2R_3R_1R_3R_1R_3R_2$
83	$\star 2^6(a)$	12	$R_1, R_3R_1R_3R_1R_3, R_3R_2R_3R_1R_3R_2R_3, R_3R_2R_1R_3R_1R_2R_3, R_2R_3R_1R_3R_1R_3R_2, R_2R_3R_2R_1R_3R_1R_2R_3R_2$
55	$\star 2^6(b)$	12	$R_2, R_3R_2R_3, R_1R_3R_2R_3R_1, R_3R_1R_3R_2R_3R_1R_3, R_1R_3R_1R_3R_2R_3R_1R_3R_1, R_3R_1R_3R_1R_3R_2R_3R_1R_3R_1R_3$

the original tile-transitive enumeration are called 'mutual duals' when they are the dual of one another – *i.e.* the dual symbols are also obtained *via* a sequence of splits and glues from the fundamental tile-transitive tiling. Such a pair generates just two distinct D symbols, rather than four. We find three pairs that are tile-1- and vertex-1-transitive, 15 pairings of a tile-1-, vertex-2-transitive tiling with a tile-2-, vertex-1-transitive one, and 69 tile-2-, vertex-2-transitive pairs. Thus, from our enumeration of 1450 tile-1- and -2-transitive D symbols generated *via* split and glue operations, we obtain a total of 2706 distinct D symbols after considering their duals, corresponding to tilings whose symmetries are those of one of the 131 subgroups. Among those distinct subgroup tilings, 95 examples are either vertex-1- or tile-1-transitive.

For ease of reference, we give each of the 2706 D symbols a distinct name of the form QCn, where $n = 1, 2, 3, \ldots$, characteristic of a subgroup tiling. We capitalize letters to indicate a tiling (as opposed to a net, whose name will contain lower-case letters); 'Q' refers to the P, D and G family of TPMSs with cubic symmetry; the letter 'C' indicates that the subgroup tiling comes from the family of Coxeter orbifolds of the (genus-3) cubic TPMS. The final part of the name is an index, n, that is assigned according to a ranking of the 2706 distinct D symbols by the following sort key:

$$(\min\{v, t\}, v, t, \text{D symbol}).$$

Here v is the vertex transitivity, t is the tile transitivity and the D symbol has a natural ordering discussed by Delgado-Friedrichs (2003). Thus, the subgroup tilings $QC1–QC95$ are a complete enumeration of (vertex or tile)-1-transitive tilings with kaleidoscopic symmetry compatible with the P, D and G surfaces, and $QC96–QC2706$ are all the 2-transitive tilings. Should we choose to continue the split–glue sequence and enumerate tile-3-transitive tilings, the *new* symbols generated can be appended to the end of this list, since any duals that happen to have lower vertex transitivity already appear in the first 2706 subgroup tilings.

4. Embedding and unfolding Delaney–Dress symbols

In the previous section we described how to enumerate Delaney–Dress symbols that encode the symmetry and topology of tilings. Our next step is to take these abstract representations of tilings and determine explicit realizations in the hyperbolic plane that are compatible with the P, D and G surface covering maps. We call this process *embedding*.

The embedding operation is a many-to-many mapping from tilings in simply connected two-dimensional hyperbolic space to tilings on the multiply connected triply periodic minimal surfaces. First, a single D symbol can have multiple realizations that project to distinct surface reticulations. Second, two embedded D symbols (from different subgroups, say) can generate the same surface reticulation. Clarification of the embedding process involves subtleties that require us to go beyond Delaney–Dress tiling theory, since the standard theory is applicable to simply connected spaces only. In order to enumerate distinct tilings of a multiply connected TPMS, we

(a) $\star 246$ F   (b) $\star 246$ FG   (c) $\star 246$ FG
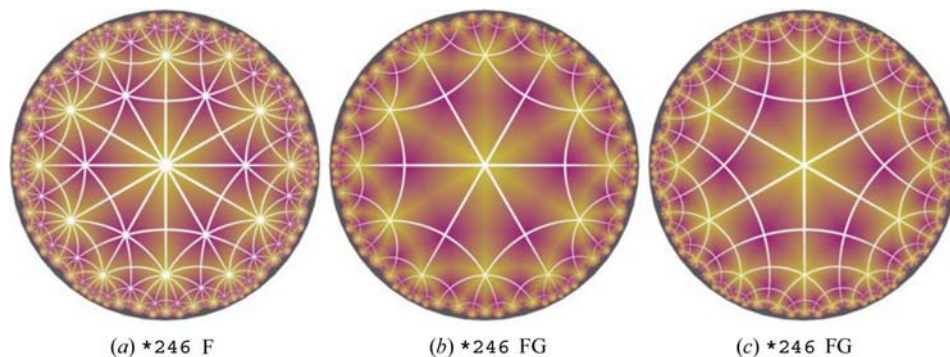
Figure 7
Three tile-1-transitive tilings of the hyperbolic plane with orbifold symmetry $\star 246$: the fundamental tiling (a), and two of the six possible glued tilings (b, c). (b) is obtained from a vertex gluing and (c) from an edge gluing. The background colouring respects $\star 246$ symmetry.

Table 4
A summary of our D-symbol enumeration.

Tilings are enumerated within each symmetry orbifold and grouped according to the glue and/or split operations required to generate the tiling symbol. As explained in the text, each tiling symbol effectively represents two tilings: the tile-transitive version and its vertex-transitive dual. The total number of distinct tiling symbols found within each orbifold is also listed.

	1-Transitive		2-Transitive			
Orbifold symbol	F	FG	FS	FSG	FSGG	No. of tilings
⋆246	1	6	6	42	56	204
⋆266	1	4	4	22	30	112
⋆344	1	4	4	22	38	128
⋆2223	1	5	10	71	115	366
⋆2224	1	5	10	71	115	366
⋆2323	1	3	7	37	67	216
⋆2244	1	5	9	69	123	394
⋆2^5	1	2	5	34	68	201
⋆2626	1	3	7	37	67	216
⋆2^6	1	2	7	53	130	375
⋆4444	1	2	4	19	39	128
Total	52		1398			2706

require an explicit representation of the tiling in the ○○○ translational domain. Finding such a representation is called *unfolding*. We call these embedded and unfolded tilings *U-tilings*.

Before proceeding further, we must clarify what we mean by 'the same surface reticulation'. Since our periodic minimal surfaces have a high degree of symmetry, we say that two surface tilings are equivalent if their Delaney–Dress chamber systems are related *via* a symmetry of the surface, otherwise they are distinct. Symmetries of the surface pull back (through the covering map) to ⋆246 symmetries in the hyperbolic plane, so two U-tilings project to equivalent surface reticulations if their D chambers are related *via* a symmetry from ⋆246.[2]

The process of generating an unfolded D symbol involves a number of steps, whose degree of complexity depends on the symmetry of the original D symbol. To clarify the description, we first discuss the construction for tilings with ⋆246 symmetry (§4.1); tilings with other Coxeter orbifolds are discussed later (§4.2).

4.1. Tilings in ⋆246

First, we introduce a triangulation of the tritorus, called the ⋆246/T chart. This chart forms a bridge between geometry in $\mathbb{H}^2$ and the group structure of ⋆246 modulo the ○○○ translational domain. We then describe how to embed a D symbol in a ⋆246 fundamental domain. Finally, we use the ⋆246/T chart as a scaffold to unfold the ⋆246 D symbol into one with ○○○ symmetry.

4.1.1. The ⋆246/T chart. The ⋆246/T chart is a triangulation of a tritorus with triangles labelled by the 96 distinct ⋆246

words that represent elements in the ⋆246/T quotient group. This chart forms the basis for unfolding a D symbol with ⋆246 symmetry into a D symbol with ○○○ symmetry. We make an explicit connection between the geometry of the ⋆246 triangulation and the reflection group, using the process described below (Fig. 10 will assist visualization).

(1) First, we list the 96 distinct elements of the quotient group ⋆246/T. We do this using the *GAP* package *kbmag* (an acronym for Knuth–Bendix in Monoids and Groups), although any coset enumeration algorithm will suffice (The GAP Group, 2002; Holt, 1998). The quotient-group cosets are represented by minimal words in the reflections R_1, R_2, R_3. The *kbmag* enumeration uses a short-lex ordering derived from a given order for the generators. The 96 elements are listed in full in Appendix C, Table 15.

(2) The ⋆246 fundamental domain is a triangle bounded by mirror lines that meet at angles $\pi/2$, $\pi/4$ and $\pi/6$. A fundamental tiling of $\mathbb{H}^2$ consists of these domains with 4, 8 and 12 triangles meeting at ⋆2, ⋆4 and ⋆6 corner points, respectively. We pick an initial triangle and give it the label of the identity element, I.
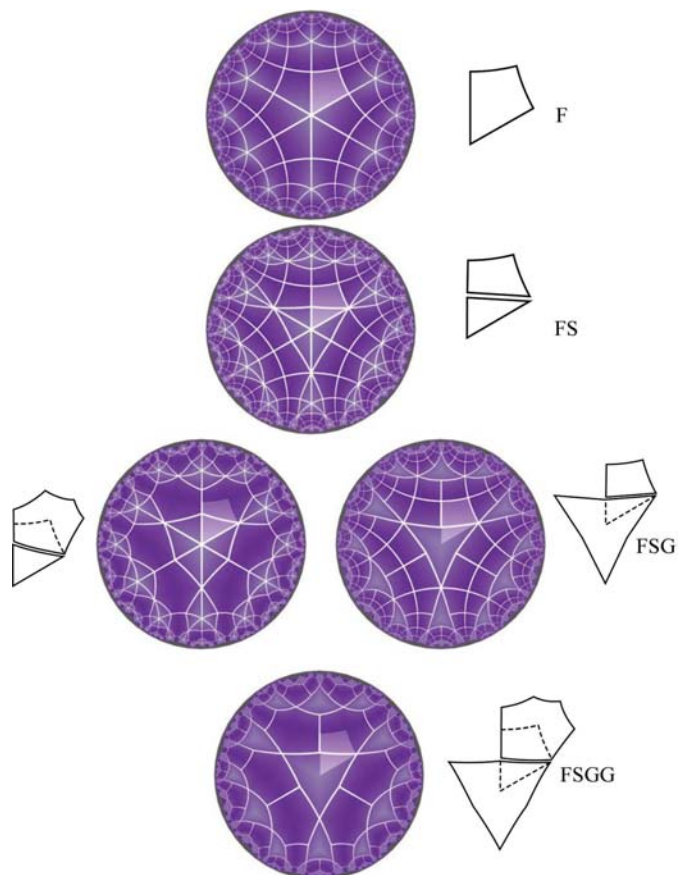

Figure 8
A sequence of tilings illustrating how tile-*k*-transitive tilings are created from fundamental (F) tilings by split (S) and glue (G) operations. The ⋆2223 fundamental tile-1-transitive F tiling is split into a tile-2-transitive tiling of type FS. Two complementary gluing operations produce two new tile-2-transitive tilings of type FSG. Finally, by applying both of these glue operations we obtain a tiling of type FSGG.

---

[2] There is an important exception to this when considering chiral tilings projected onto the *G* surface; see Robins *et al.* (2005) for details. Chirality (in $\mathbb{H}^2$) is not an issue here, since we are working with kaleidoscopic symmetry groups.
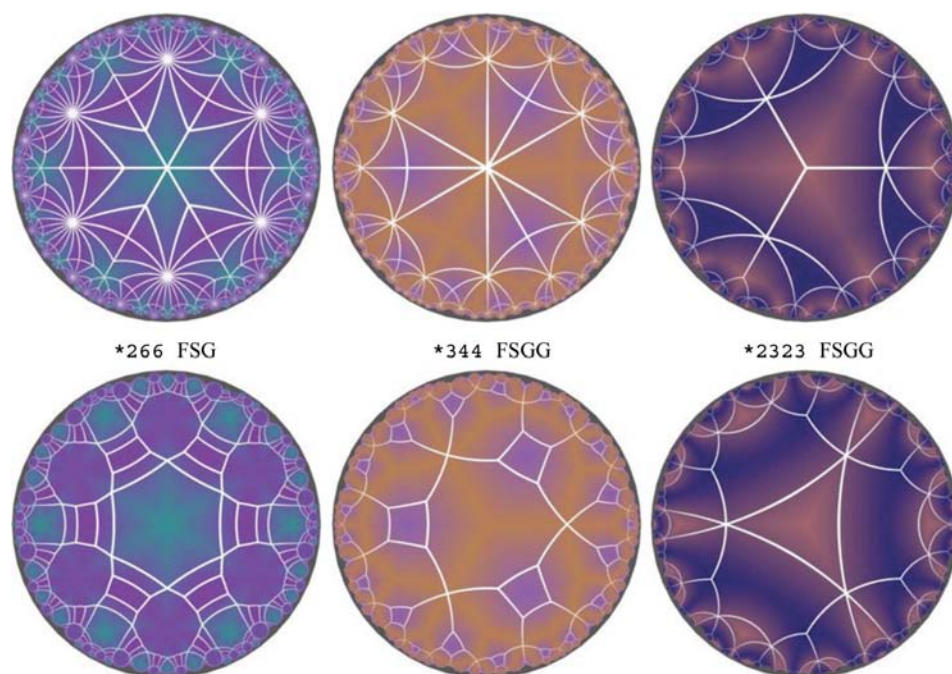
Figure 9
Illustrations of duality. Each column shows a tile-2-transitive tiling above its dual vertex-2-transitive tiling, together with orbifold and tiling class labels (where F, S and G refer to the fundamental tiling, split and glue operations, respectively). The tilings with symmetry ⋆266 illustrate the general case in which the dual D symbol is not present in the initial tile-transitive enumeration. In contrast, the tilings with symmetry ⋆344 are both tile-2- and vertex-2-transitive. Each is generated *via* a different FSGG sequence applied to the ⋆344 fundamental tiling, and this pair of tilings is called 'mutually dual'. The third column of tilings with symmetry ⋆2323 shows a self-dual example.

(3) Since the reflections R_1, R_2 and R_3 map the identity triangle onto its three neighbours, these triangles are labelled accordingly. The R_1 triangle is the neighbour opposite the ⋆4 corner point, R_2 lies across the shortest edge (opposite the ⋆6 corner), and R_3 is the neighbour sharing the hypotenuse of the identity ⋆246 triangle, opposite the ⋆2 point.

(4) We continue labelling each triangle by the minimal R-word that maps the identity triangle onto it. By continuity, the neighbours opposite the ⋆4, ⋆6 and ⋆2 corners of a given triangle, V, will be VR_1, VR_2 and VR_3. For example, the neighbours of R_1 are $R_1R_1 = I$, R_1R_2 and R_1R_3, and the neighbours of R_2 are $R_2R_1 = R_1R_2$, $R_2R_2 = I$ and R_2R_3. The process of converting each neighbour-word VR_i to its minimal version is called *word reduction*, and is performed using the *kbmag* package.

(5) The translational periodicity is encoded by performing word reduction within the ⋆246/T quotient group, rather than the full ⋆246 group. The diagram in Fig. 10 shows, for example, the ⋆2-neighbour of triangle V, with label VR_3, which lies outside the 96 elements of the dodecagon. In fact triangle VR_3 is the image of triangle W under the translation t_2. In the quotient-group word reduction, we find $VR_3 \rightarrow t_2W$. Thus, in the compact tritorus, triangles V and W are neighbours along the edge opposite their ⋆2 corners. We attach the t_2 translation to this adjacency information to mark a boundary or cut line that enables us to convert from the compact surface triangulation to the universal cover in $\mathbb{H}^2$. We also use the cut lines to

transform tilings and networks from $\mathbb{H}^2$ to surface reticulations in $\mathbb{E}^3$. This is described in §5.1.

There is a nice correspondence, hinted at above, between the ⋆246 isometry that maps the identity triangle onto some image triangle and a path through the triangulation starting at I. For example, the isometry that maps the identity triangle onto the W triangle is $R_3R_1R_3R_1R_2R_3R_1$, see Fig. 10. By convention, the isometry operation acts from the left, so the order of successive reflections is read from right to left. In contrast, the path through the triangulation associated with this ⋆246-word is read from left to right. It starts at I, visits its 'R_3-neighbour' (the triangle opposite its ⋆2-corner), then this triangle's 'R_1-neighbour' (the triangle opposite a ⋆4 corner) and so on, ending at the triangle labelled W. Similarly, triangle V is reached by a path starting at I then (reading from left to right) visiting neighbours according to the word $R_1R_3R_1R_2R_3R_1$, a twofold rotation about the circled vertex in Fig. 10.

Finally, the translation that maps W onto VR_3 may be encoded by a path in the triangulation from W to VR_3 passing through I,

$$W^{-1} \cdot V \cdot R_3 = R_1R_3R_2\,R_1R_3R_1R_3 \cdot R_1R_3R_1\,R_2R_3R_1 \cdot R_3 = t_2.$$

Word reduction to the expression for t_2 given in equation (1) uses the identity $(R_1R_3)^6 = I$ which implies that the central part of the word reduces as

$$R_1R_3R_1R_3R_1R_3R_1 = R_3R_1R_3R_1R_3.$$

4.1.2. Embedding in ⋆246.

We now return to Delaney–Dress symbols and the problem of embedding them in the hyperbolic plane. Readers unfamiliar with Delaney–Dress theory should refer to Appendix B for the relevant definitions.

The embedding of D symbols with ⋆246 symmetry is simpler than other subgroups for two reasons. First, the geometry of the fundamental domain is uniquely determined by the angles of the hyperbolic triangle. Second, the three distinct angles permit only one way to embed a D symbol into this fundamental domain.

Consider, for example, the tile-1-transitive glued tiling (shown in Fig. 7b). The D symbol for this tiling is given in Table 5 and its embedding into the ⋆246 fundamental domain is shown in Fig. 11. The D symbol has two chambers labelled a and b that form a triangulation of the ⋆246 orbifold. Embedding this symbol amounts to determining how the two

chambers sit within the ⋆246 triangle. Each triangular chamber has a 0-, 1- and 2-vertex that correspond, respectively, to a vertex of the tiling, an edge midpoint and a tile centroid. First note that the D symbol specifies that a and b are adjacent along their 0-edge, *i.e.* the chamber edge opposite a vertex of the tiling, and that a and b are self-adjacent across their 1- and 2-edges. This means that the 0-edge is internal to the ⋆246 domain and the 1- and 2-edges lie along mirror boundaries. The topological indices m_{12} specify the degree of the tile vertex that sits at the 0-vertex of each chamber. So the 0-vertex of a has degree 6 and the 0-vertex of b has degree 4. This tells us that the 0-vertex of a must sit at the ⋆6 corner and the 0-vertex of b must sit at the ⋆4 corner. Finally, the index
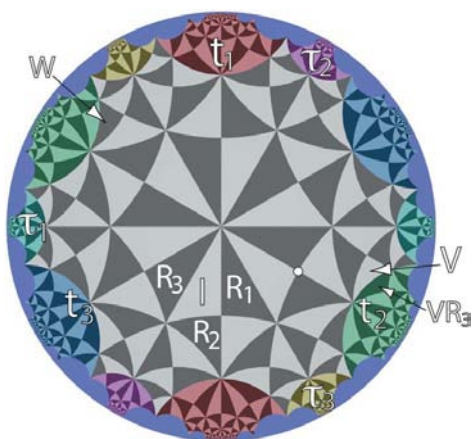
Table 5
The Delaney–Dress symbol for the glued tiling ⋆246 FG 3.

	a	b
0-*nbr*	b	a
1-*nbr*	a	b
2-*nbr*	a	b
m_{01}	4	4
m_{12}	6	4

m_{01} specifies the order of the tile centred at the 2-vertex of a chamber. For this example, the single tile is a quadrilateral and the 2-vertex (of both a and b) sits at the ⋆2 corner.

In general, we know that the chambers of the D symbol must form a triangulation of the ⋆246 fundamental domain. We therefore need to identify the chamber edges that lie along mirror boundaries, and chamber walks that trace a circuit around each corner point of the ⋆246 triangle. This information comes from the adjacency information and the topological indices given by the D symbol. Once the mirror boundaries and corner points are identified, we can deduce the embedding of the remaining chambers into the fundamental domain.

4.1.3. Unfolding into the tritorus. Once we know how a D symbol sits inside the ⋆246 fundamental domain, we can use the ⋆246/T chart to unfold this D symbol to cover the tritorus and thus obtain a larger D symbol with ∘∘∘ symmetry. The process involves a straightforward tile-rewriting procedure, described below for the ⋆246 fundamental tiling, and illustrated in Fig. 12.

Figure 10
The 96 translationally distinct ⋆246 triangles in the ∘∘∘ domain for the *P*, *D* and *G* surfaces are shaded light and dark grey. The 12 domains bordering the central dodecagon are the images of the central polygon after translation by $t_1, t_2, t_3, \tau_1, \tau_2$ and τ_3. The triangles labelled R_1, R_2 and R_3 are the appropriate images of the identity triangle I after reflection. The triangle labelled V is the image of I under a $\pi/2$ rotation about the circled vertex. The R_3 neighbour of triangle V, VR_3, lies outside the central dodecagon and is a translate of the triangle labelled W, *i.e.*, $VR_3 = t_2W$.
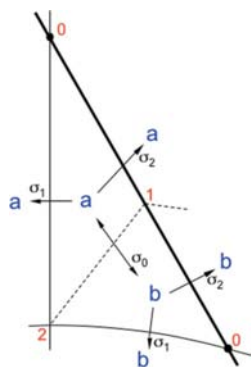

Figure 11
The chamber system for the glued (FG) tiling of ⋆246 (*cf.* Fig. 7b and Table 5). Only a single ⋆246 fundamental domain is shown, with the two chambers a and b labelled in blue, and their 0-, 1-, and 2-vertices marked in red. Tile edges and vertices are shown by heavy black lines and dots. Thin black lines lie between chambers that are 1-neighbours and dashed lines lie between chambers that are 0-neighbours.

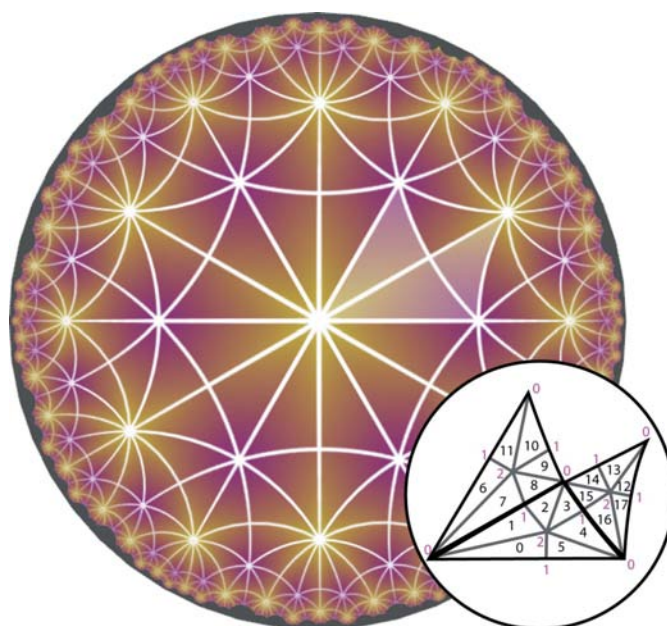**Figure 12**
Unfolding the D symbol of the fundamental ⋆246 tiling into the topologically equivalent tiling with symmetry ∘∘∘ (the unfolded D symbol). The inset shows the chamber labelling (darker labels) of three fundamental domains. Chambers that are equivalent under a ⋆246 symmetry have identical labels modulo 6. The red labels mark the 0-, 1- and 2-vertices of the chamber system.

The D symbol for $\star$246 F contains six chambers labelled 0 to 5. Since the $\star$246/T chart has 96 triangles, the $\circ\circ\circ$ D symbol derived from $\star$246 F will have 576 chambers. These chambers are labelled so that chamber n is symmetrically equivalent to chamber n [modulo(6)], see Fig. 12. In this example, the chamber adjacencies opposite 0- and 1-vertices are exactly those derived from the $\star$246 F symbol. The adjacencies across 2-edges map from one tile into its neighbour, so these relations are derived from the $\star$246/T chart structure. The cuts defined on the $\star$246 chart are also attached to the appropriate D-chamber adjacencies and recorded as information additional to the standard D-symbol form. These cuts effectively define the embedding of the $\circ\circ\circ$ D symbol into the T subgroup of $\star$246.

The *glue* and *split* operations described in §3 preserve the symmetry group, so their chambers cover the same fundamental orbifold domain. Thus, unfolding D symbols associated with tilings generated *via* split and glue operations proceeds in almost identical fashion to that for the fundamental tiling described above.

4.2. Tilings in other kaleidoscopic subgroups

Consider next the construction for D symbols whose symmetries are kaleidoscopic subgroups of $\star$246. The embedding and unfolding of those D symbols proceeds in essentially the same manner as described for $\star$246. However, we shall see that D symbols for subgroups of $\star$246 can embed in more than one way, further complicating the construction, but giving additional surface reticulations.
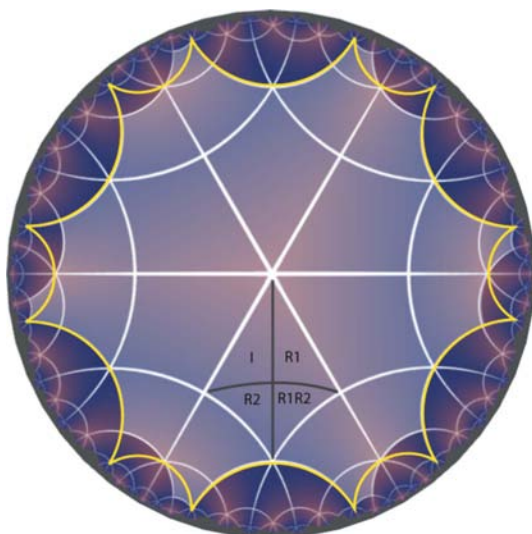

Figure 13
The $\star$2323 subgroup chart. A fundamental domain for $\star$2323 is built from the four $\star$246 triangles labelled I, R_1, R_2 and R_1R_2. Twenty-four of these quadrilateral domains are needed to fill out a translational unit cell. One possible set of 24 quadrilaterals is the highlighted (18-sided) polygonal domain. The dodecagon bounded by yellow edges marks the translational unit cell described in Fig. 10. Note how the highlighted domain differs from the (yellow) dodecagon by displacing a triangle from one edge to the diagonally opposite one.

4.2.1. Building a subgroup chart. The $\star$246 chart is built from knowledge of the quotient-group structure, *i.e.* explicit representation of the t_i, τ_i translations in terms of the R_1, R_2, R_3 reflections. Although we constructed the 131 subgroups to have the t_i, τ_i translations as elements, we do not have explicit expressions for the t_i, τ_i translations as words in the subgroup generators, so it is simpler to use a combinatorial procedure to obtain the subgroup charts. For the 14 kaleidoscopic subgroups this is feasible because each subgroup fundamental domain is built from k whole $\star$246 triangles, where k is the index of the subgroup.

To obtain a fundamental domain for a kaleidoscopic subgroup we use the coset labelling of $\star$246 triangles by the subgroup action. A contiguous set of k distinct coset elements gives us a fundamental domain. For example, $\star$2323, an order-4 subgroup of $\star$246, has a fundamental domain built from the four $\star$246 triangles labelled I, R_1, R_2 and R_1R_2, as shown in Fig. 13. This $\star$2323 domain is replicated by reflections in its boundaries, and we use a cumulative algorithm to acquire one fundamental domain at a time. For example, a neighbouring fundamental domain is the set of triangles labelled $\{R_3, R_3R_1, R_3R_2, R_3R_1R_2\}$, *i.e.*, the image under R_3 of the initial domain. The number of subgroup fundamental domains that covers the tritorus is 96/k, so the $\star$2323/T chart has 24 quadrilaterals.

The final step in building the $\star$2323/T chart is to identify new cut lines that follow $\star$2323 domain boundaries. Starting from an initial $\star$2323 domain, we grow out one $\star$2323 domain at a time with the aim of keeping the region as circular as possible in the hyperbolic plane. This region is shown highlighted in Fig. 13. Eventually, the growing region will meet itself on the genus-3 surface. The edges where this occurs define new cut lines for unwrapping the tritorus into the hyperbolic plane. The translation associated with each cut line is found by tracing a path through the underlying $\star$246/T chart that is constrained to pass through the I triangle and the cut-line edge. The path defines a $\star$246-word that reduces to a $\circ\circ\circ$ translation under *kbmag* word reduction.

4.2.2. Embedding and unfolding in a subgroup chart. The embedding and unfolding of D symbols into a subgroup chart uses the same feature identification and tile-rewriting procedures described earlier for $\star$246. However, fundamental domains for the subgroups are not as constrained as the $\star$246 triangle, and this means some D symbols have more than one embedding compatible with the covering map. There are two routes to generating these multiple embeddings: distinct subgroups with the same orbifold and automorphisms of a subgroup fundamental domain. We illustrate each of these situations below.

First, recall from Table 3 that three orbifolds ($\star$2244, $\star2^5$ and $\star2^6$) appear as pairs of distinct subgroups of $\star$246. The different subgroup structures guarantee that these pairs have distinct subgroup charts. Thus, every D symbol from one of these orbifolds has two distinct embeddings in $\mathbb{H}^2$. An example from $\star2^5$ is illustrated in Fig. 14. Although the two embeddings have identical topology, they are not related by a symmetry of $\star$246, and therefore project to different surface reticulations. The difference is also apparent *via* unfolding in

the respective subgroup charts: a single D symbol embedded in two distinct subgroup charts generates distinct ○○○ D symbols.

Second, automorphisms of the kaleidoscopic subgroups are manifested as abstract symmetries of their fundamental domains, which may or may not be associated with symmetries from ⋆246. Multiple embeddings of a D symbol are possible when the subgroup domain has automorphisms that are not induced by an element of ⋆246. If the tiling does not display the corresponding automorphism of its D chambers, then it has two distinct embeddings that are not related by a ⋆246 symmetry and therefore map to inequivalent surface reticulations. The kaleidoscopic subgroups with automorphisms that can produce distinct embeddings of D symbols are ⋆2224, ⋆2⁵(a), ⋆2⁵(b), ⋆2626, ⋆4444 and ⋆2⁶(a).

We illustrate this situation with an example. The subgroups ⋆2223 and ⋆2224 both have quadrilateral fundamental domains and an automorphism that swaps the two opposing ⋆2 corner points (see Fig. 15). The automorphism of ⋆2223 is defined by conjugacy with the ⋆246 R_3 reflection, so automorphic embeddings of a D symbol will project to equivalent surface reticulations. The geometry of the ⋆2224 subgroup domain shows that the automorphism of its domain is not a ⋆246 conjugacy. Now suppose we have a D symbol that is embedded in the ⋆2224 domain, and consider the effect of the subgroup automorphism. Clearly this induces a new embedding of the D symbol, which may or may not be equivalent to the original embedding. Equivalence occurs only when the D symbol has an automorphism of its chambers that corresponds to the automorphism of the subgroup domain. Thus, a split tiling where the extra edge divides the domain into two quadrilaterals (as in Fig. 16) will have two distinct embeddings, but a split between two opposite vertices will not.

We conclude this section by discussing the effects of a subgroup automorphism on the translation subgroup, T, and the corresponding unfolded ○○○ D symbols. There are two possibilities: the subgroup automorphism may or may not induce an automorphism of T. If T is preserved, then the two embeddings of a subgroup D symbol will unfold to isomorphic ○○○ D symbols. If T is not preserved, then distinct ○○○ D symbols are generated. The only kaleidoscopic subgroup where we see distinct embeddings of a tiling that unfold to equivalent ○○○ D symbols is ⋆2626. This is due to an automorphism of ⋆2626 that swaps the role of the t_i and τ_i generators of T. This automorphism is not generated by a ⋆246 isometry but nonetheless preserves the T subgroup, so the unfolded ○○○ D symbols are isomorphic. An example is shown in Fig. 17.

Figure 14
The orbifold ⋆2⁵ appears as two distinct subgroups of ⋆246. The top row shows the tilings resulting from the subgroup labelled ⋆2⁵(a) in Table 3; the bottom row shows analogous cases for ⋆2⁵(b). Fundamental tilings in the hyperbolic plane for each subgroup are illustrated on the left. In each case we show a single tile broken into ⋆246 triangles and a highlighted translational domain. The central column displays fundamental tilings wrapped onto the P surface using the covering map, forming E-tilings (whose edges and vertices describe e-nets). The irregularly shaped unit cells correspond exactly to the highlighted region in the hyperbolic plane images to the left. The right column shows s-nets derived from the E-tilings by forming the most symmetric Euclidean embeddings of the e-net topologies. Tiling and net labels are described in §§5.1 and 6.3.

5. Surface tilings

Our enumeration so far has taken distinct D symbols from each Coxeter orbifold, found all possible embeddings of these symbols into the kaleidoscopic subgroup domains, and unfolded them into corresponding tritorus charts. The result is 14 lists of subgroup tilings and their unfolded ○○○ D symbols augmented with 'cuts' that determine their embedding in the T subgroup. Since these tilings are obtained *via unfolding* and are embedded in the *universal* cover of the triply periodic minimal surfaces, we call them *U-tilings*. In this section we describe how to determine a unique representative for each distinct U-tiling. These tilings are compatible with the surface covering maps, and so project onto the P, D and G surfaces giving *E-tilings*. Finally, we also consider tilings of the tritorus, or *O-tilings*, obtained from these U-tilings.

5.1. U-tilings

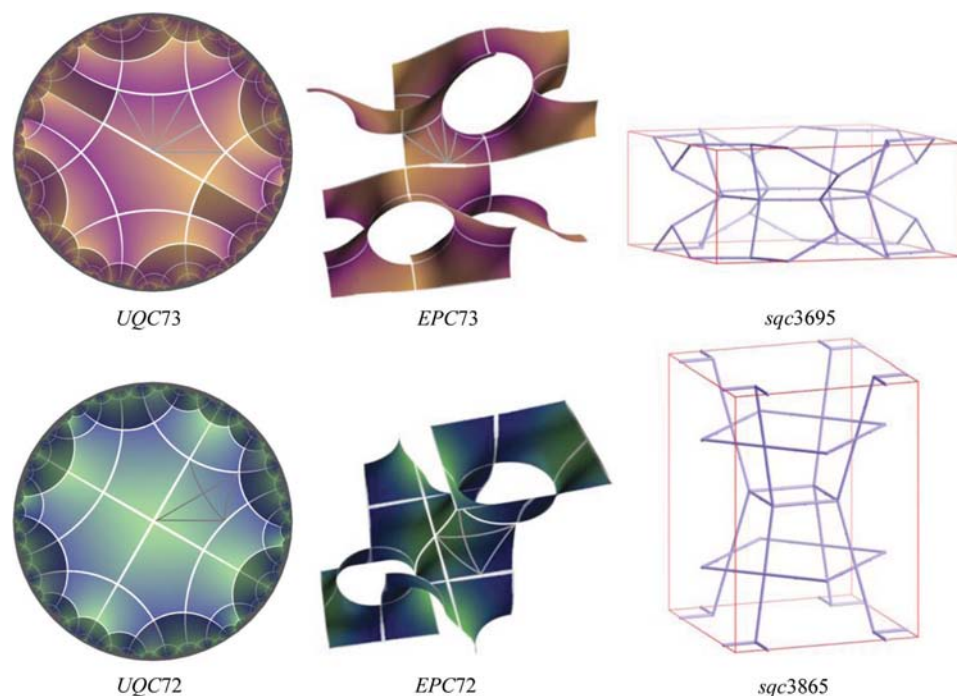In §4 we described how to generate distinct embeddings of D

symbols within each kaleidoscopic subgroup of ⋆246. Each U-tiling is represented by a ◦◦◦ D symbol augmented with cuts and its precursor D symbol. This extra information defines the embedding of the tiling in the ⋆246/T domain. We determine which U-tilings generate equivalent surface reticulations as follows.

Recall from the introduction to §4 that two tilings of a TPMS are equivalent if their chamber systems are related *via* a symmetry of the surface. Thus, two U-tilings are distinct if they have distinct ◦◦◦ D symbols, or if they are distinct embeddings

Figure 15
Fundamental tilings for the ⋆2223 and ⋆2224 subgroups, each shown with a single domain subdivided by ⋆246 triangles. The automorphism swapping the two opposite ⋆2 corners of ⋆2223 is generated by a ⋆246 isometry, but this is not the case for ⋆2224. This leads to multiple embeddings of some D symbols in ⋆2224.
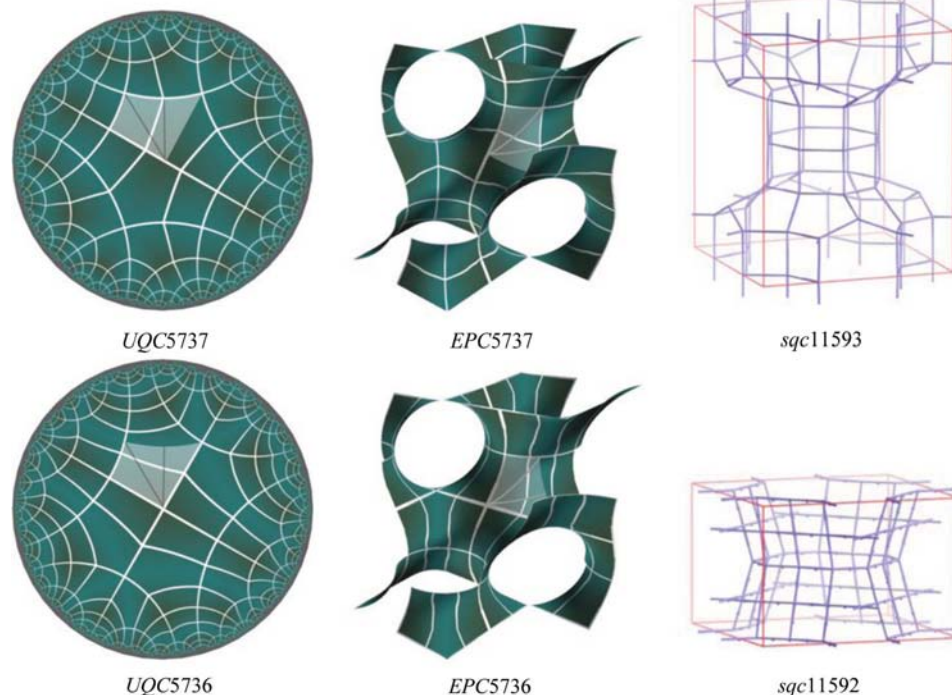
Figure 16
Left: two U-tilings with ⋆2224 symmetry obtained *via* a split from a ⋆2–⋆2 edge to the opposite ⋆2–⋆4 edge, followed by a glue across the other ⋆2–⋆4 edge, derived from *QC*2566. These tilings are distinct embeddings of a single D symbol and therefore generate different surface reticulations. To illustrate this, we show the projections onto the *P* surface (centre) and resulting s-nets (right). Tiling and net labels are described in §§5.1 and 6.3.

of the same ◦◦◦ D symbol. The first situation is easy to test using combinatorial tiling algorithms. The second requires the additional information about the embedding of the U-tilings in the *T* subgroup, either from the cuts, or from an understanding of the embedding and unfolding process. By their construction, the U-tilings generated within a particular subgroup chart are distinct. However, two U-tilings defined *via* different subgroup charts may have the same ◦◦◦ D symbol and equivalent 'cuts'. This occurs when the two precursor D symbols and their embeddings are related by symmetry raising or lowering within the ⋆246 subgroup lattice (see Fig. 6). An example is shown in Fig. 18, where we see a vertex-1-transitive tiling with symmetry ⋆2224, and two different symmetry lowerings to vertex-2-transitive tilings in the subgroups ⋆4444 and ⋆2⁶(b). The three U-tilings have isomorphic ◦◦◦ D symbols and equivalent embeddings.

We next consider the union of the 14 lists of U-tilings, determine equivalence classes by comparing the ◦◦◦ D symbols as described above, and find that there are 6079 distinct U-tilings. We assign them names of the form *UQCn*, where 'Q' and 'C' stand for *cubic* and *Coxeter*. The running index *n* is determined by ranking the distinct U-tilings by their ◦◦◦ D symbols (*cf.* §3). If two distinct U-tilings have the same ◦◦◦ D symbol, as for the ⋆2626 example of Fig. 17, they are listed in arbitrary order. In Table 6, we give the number of U-tilings found within each kaleidoscopic subgroup. We count each equivalence class of U-tilings just once, in the subgroup of the precursor tiling that has the highest symmetry.

5.2. E-Tilings

We now map each U-tiling from the hyperbolic plane onto the *P*, *D* and *G* periodic minimal surfaces, forming *E-tilings*, where 'E' refers to Epinet. Each distinct U-tiling *UQCn* maps to exactly three E-tilings and we give these the directly corresponding names, *EPCn*, *EDCn* and *EGCn*; see Fig. 19. The E-tilings are the bridge between two-dimensional and three-dimensional structure that our enumeration scheme is built around. They are two-dimensional because they are defined almost entirely by the U-tiling; three-dimensional structure is given by the covering map and the geometry of the particular TPMS. More specifically, we use the covering map action on the hyperbolic translations, $t_i, \tau_i \in T$ (as summarized in Table 2 of §2) to map the cuts attached to the ◦◦◦ D symbol into Euclidean translations. The
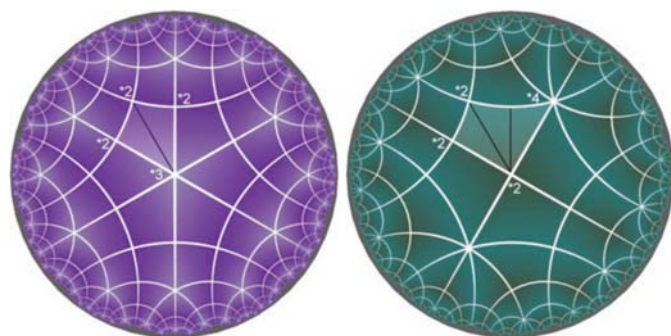
E-tiling, then, is represented by the same pair of subgroup tiling and ∘∘∘ D symbol as the U-tiling, but the cuts are written as Euclidean translations, rather than hyperbolic ones.

The space-group symmetries of the E-tilings can be found by considering the covering map action in more detail. Recall from §2 that the covering map induces a group homomorphism, $\phi : \star 246 \to S$, where S is the space group of the non-oriented surface. The map ϕ is an explicit relationship between the hyperbolic reflections R_1, R_2 and R_3, and Euclidean isometries from S. If we now restrict the action of ϕ to one of the kaleidoscopic subgroups, $\Gamma \subset \star 246$, then the image $\phi(\Gamma)$ is a subgroup of the Euclidean space group S, giving an elegant correspondence between two-dimensional hyperbolic (discrete) groups and three-dimensional Euclidean space groups. This procedure will be explored in detail in a future publication. We give the surface space groups that correspond to each kaleidoscopic subgroup in Table 6.

5.3. O-Tilings

Our hyperbolic tilings also define tilings of the tritorus, which we call O-tilings. These tilings are formed by wrapping U-tilings onto the tritorus, whose structure is explained in §4.1.3. We distinguish O-tilings by the distinct ∘∘∘ D symbols. In our current enumeration we find 5912 O-tilings. This is less than the number of distinct U-tilings (6079) because of the $\star 2626$ examples discussed at the end of §4.2.2.
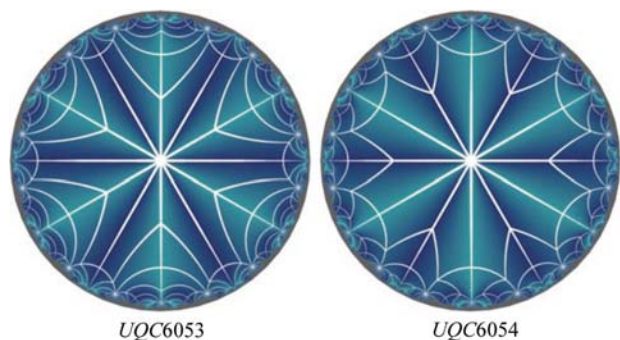

Figure 17
Two embeddings of a split tiling of $\star 2626$, related *via* an automorphism of $\star 2626$ that swaps the role of t_i and τ_i translations in the T subgroup. Tiling labels are described in §5.1.
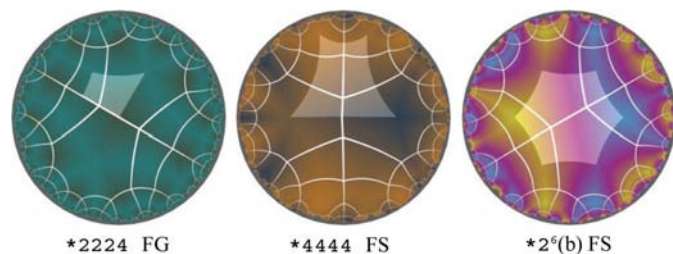

Figure 18
Three subgroup tilings in the equivalence class for the U-tiling *UQC*24. The tilings with symmetry $\star 4444$ and $\star 2^6$(b) are two different symmetry lowerings of the tiling with symmetry $\star 2224$. In each case a single fundamental domain is highlighted.

6. Nets from tilings

In this section we consider the nets derived from the vertices and edges of tilings. Various definitions of the word 'net' can be found in the literature; we use the term here to denote a graph embedded in a metric space. Specifically, we investigate nets embedded in the hyperbolic plane, the tritorus and the periodic minimal surfaces, called *h-nets*, *o-nets* and *e-nets*, respectively. Finally, we discard possible edge entanglements present in the e-nets to study their topology and maximal-symmetry embedding in $\mathbb{E}^3$, and call the resulting structures *s-nets*. h-nets remain embedded in their parent space – the hyperbolic plane – while o-, e- and s-nets are embedded in three-dimensional Euclidean space.

There is a simple canonical form for h-nets that allows different net topologies to be readily distinguished in the hyperbolic plane. Some care must be taken, however, in defining the equivalence class of three-dimensional embedded nets, namely the e-net, o-net and s-net categories. Trivial geometric deformations of a particular embedding are of little interest to us. Nontrivial deformations may include those that change the entanglement of edges within the graph, or the graph topology. We define equivalence classes for *o-nets* and *e-nets* to include all graph embeddings related by deformations that do not involve edges crossing. Accordingly, one of our goals is to determine equivalence classes of nets under ambient isotopy: *i.e.*, two nets are equivalent if their embedding space can be continuously deformed in such a way as to map one onto the other without allowing edges to pass through each other. This leads to the distinction between variously entangled versions of a net as inequivalent o- or e-nets, akin to the distinction between inequivalent entanglements of a loop in space as different knots. Finally, the equivalence class of *s-nets* includes all nets with identical graph topology. As for h-nets, there is a canonical form for s-nets that allows most distinct examples to be readily identified.

6.1. h-Nets

As illustrated in Fig. 20, many tilings can carry the same net topology. Since a Delaney–Dress symbol encodes both topology and symmetry, such tilings have different D symbols. Clearly then, the number of distinct h-net topologies will be less than the number of distinct subgroup tilings enumerated in Table 4 of §3.

We adopt the convention that an h-net should be represented by a tiling with maximal possible symmetry. The process of finding the highest-symmetry version of a given tiling is rather simple using combinatorial tiling theory – it amounts to computing the *minimal image*[3] of the D symbol (Delgado-Friedrichs, 2003). Since minimal-image D symbols classify nets up to homeomorphism in the hyperbolic plane, they are a canonical form for h-nets.

We obtain a list of 2451 distinct h-net topologies by applying the minimal-image algorithm to the 2706 subgroup tilings

---

[3] Here the term 'minimal' refers to the length and ordering of the symbol, and should not be confused with minimal surfaces.

referred to in Table 4. The results are presented in Table 7 and are organized by the orbifolds of the minimal-image D symbols. Although we start with D symbols from kaleidoscopic subgroups of ⋆246, the resulting h-nets have a broader range of symmetry types, for example 2 ⋆ 26, a *hat* orbifold with a single cone point and mirror boundary. Further, the h-net symmetry need not be one of the 131 subgroups of ⋆246, for example, ⋆248. This is because the minimal image of a D symbol may form a tiling whose symmetries are a supergroup of the symmetries of the initial tiling. This catalogue of h-nets will allow us to compare reticulations across different classes of TPMSs.

6.2. e-Nets

We call the edge skeletons of the E-tilings of the three-periodic minimal surfaces *e-nets*, where 'e' again refers to Epinet. Recall that distinct e-nets are characterized by embeddings in three-dimensional space which cannot be deformed to each other without edge crossings or changes in

Table 6
The 14 kaleidoscopic hyperbolic symmetry groups compatible with the P, D and G surfaces.

The number of distinct U-tilings generated from each subgroup is given in the second column (*cf.* Table 4). The P, D and G tiling columns list the Euclidean space groups of the surface reticulations that arise from each hyperbolic symmetry.

Subgroup name	No. of U-tilings	P tiling	D tiling	G tiling
⋆246	204	$Im\bar{3}m$	$Pn\bar{3}m$	$Ia\bar{3}d$
⋆266	92	$Pn\bar{3}m$	$Fd\bar{3}m$	$Ia\bar{3}$
⋆344	108	$Pm\bar{3}n$	$P\bar{4}3m$	$I\bar{4}3d$
⋆2223	336	$Pm\bar{3}m$	$P4_232$	$I4_132$
⋆2224	694	$I4/mmm$	$P4_2/nnm$	$I4_1/acd$
⋆2323	183	$P4_232$	$F\bar{4}3m$	$I2_13$
⋆2244(a)	366	$P4_2/mmc$	$P\bar{4}2m$	$I\bar{4}2d$
⋆2244(b)	370	$P4/nmm$	$I4_1/amd$	$I4_1/a$
⋆2^5(a)	889	$P4/mmm$	$P4_222$	$I4_122$
⋆2^5(b)	893	$Fmmm$	$Cmma$	$Fddd$
⋆2626	396	$R\bar{3}m$	$R\bar{3}m$	$R\bar{3}$
⋆4444	211	$P\bar{4}m2$	$I\bar{4}m2$	$I\bar{4}$
⋆2^6(a)	1026	$Cmma$	$Imma$	$C2/c$
⋆2^6(b)	327	$Pmmm$	$P222$	$I2_12_12_1$

the underlying graph topology. Thus, in particular, different e-nets may have identical topology. The situation is analogous to distinct knottings of a loop, characterized by ambient isotopy (Adams, 2004).

Since an e-net inherits the embedding of an E-tiling in the topologically complex TPMS, its edges may exhibit complex entanglements in three-dimensional space induced by windings about surface channels; for example, e-nets can be self-catenated (as in Fig. 24).

e-nets may also be multigraphs, with more than one edge linking a pair of vertices. These multigraphs can arise from two sources. The distinction between these two scenarios lies in the different types of edge cycles that emerge in E-tilings: those that bound a tile and lie wholly in the surface, called 'null-homotopic rings', and those that do not bound a patch of the surface. The former cycles are found in the two-dimensional universal cover, the latter are not. Therefore the first examples are caused by h-nets that are themselves multigraphs; these are generated as follows. Our Delaney–Dress tiling enumeration allows tilings with adjacent tiles that share more than a single edge; examples are readily generated by

*UQC*473

*EPC*473   *EGC*473   *EDC*473

*sqc*12818   *sqc*14257   *sqc*9141

Figure 19
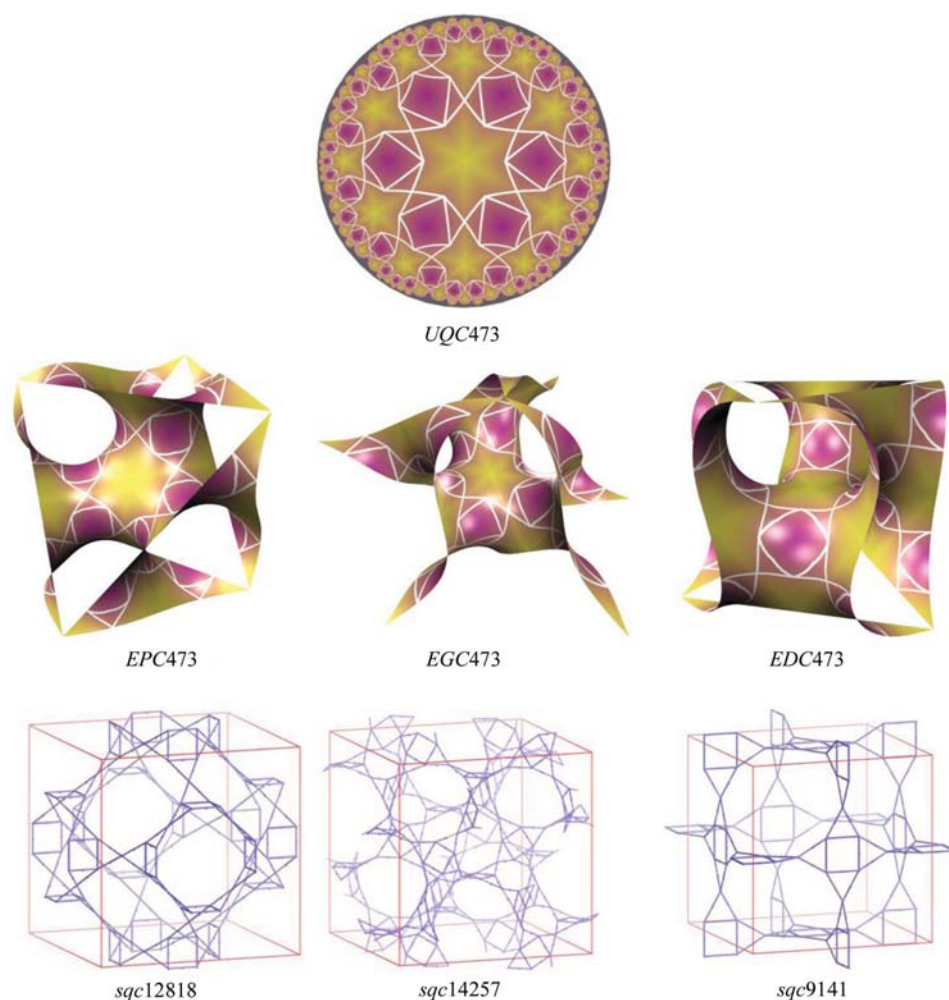A single hyperbolic tiling – here a vertex-2-transitive ⋆246 tiling formed by a sequence of split and glue operations – generates three distinct E-tilings on the P, G and D surfaces. The edges and vertices of E-tilings describe e-nets (*epc*473, *egc*473 and *edc*473, respectively). Canonical embeddings induced by the e-net topologies result in the distinct s-nets, shown below their E-tiling antecedents.

Table 7
The distribution of h-nets formed by symmetrizing the subgroup tilings as much as possible.

Orbifolds are grouped into three classes. The most populated class is that of Coxeter orbifolds from the original 11 orbifolds associated with the *P*, *D* and *G* surfaces. Two further classes are listed (separated by blank rows): Coxeter orbifolds outside the *P*, *D* and *G* class, and hat orbifolds. The total number of h-nets derived from the symbols enumerated in Table 4 is 2451; only 126 of these have symmetries different from the original 11 orbifolds.

Orbifold	No. of h-nets
⋆246	204
⋆266	92
⋆344	108
⋆2223	336
⋆2224	326
⋆2323	167
⋆2244	358
⋆22222	161
⋆2626	168
⋆222222	314
⋆4444	91
⋆248	28
⋆446	14
⋆488	14
⋆2(12)(12)	11
⋆24(12)	10
⋆2226	4
⋆288	3
⋆23(18)	2
⋆238	2
⋆239	2
⋆245	2
⋆334	2
2 ⋆ 23	10
2 ⋆ 26	10
2 ⋆ 222	6
2 ⋆ 44	6
Total	2451

Table 8
List of E-tilings on the *P* and *D* surfaces (Fig. 22) that generate equivalent e-nets, all congruent to the s-net (*sqc947*) illustrated in Fig. 23.

Two distinct two-dimensional hyperbolic nets give rise to these *D* and *P* surface E-tilings: {4, 6} (*hqc5*) and {6, 6} (*hqc10*).

QC No.	(Group)	hqc No.	Tiling type	UQC No.	Surface
1	(⋆246)	5	FG (dual)	UQC1	D
1905	(⋆2626)	5	FSGG	UQC2	P
3	(⋆266)	10	FG	UQC5	P
110	(⋆2626	10	FSGG (dual)	UQC6	D

map action. The simplest examples occur when two edges lying on opposite sides of a channel join the same two vertices. The resulting e-net therefore contains a double bond.

We label e-nets *epcN*, *edcN* and *egcN*, where 'e' stands for Epinet, 'p', 'd' and 'g' denote one of the three TPMSs, and 'c' refers to the Coxeter orbifold family. The value of the index, *N*, depends on the net embedding in $\mathbb{E}^3$, up to edge entanglement. Consider, for example, an e-net generated from the E-tiling *EDCN*. The name of this e-net is determined as follows. First, we compare the candidate e-net with the accumulated list of distinct e-nets generated from simpler tilings, *viz epcJ*, *egcJ* and *edcJ*, where $J \in \{1, \ldots, N-1\}$. If our candidate net is equivalent to a previous e-net, say *epcN'*, where $N' < N$, the tiling does not result in a new e-net. Formally, the e-net name is remapped to *epcN'* and the label *edcN* is unused. If the candidate is distinct from the list of 'lower' e-nets, it is labelled by the new name *edcN*, where *N* is exactly the same index as that of the corresponding E- (and U-) tiling. Thus the maximum index number used is less than or equal to the total number of e-nets, due to unfilled labels wherever the E-tiling skeleton is ambient isotopic to a previous e-net. This schema allows the indexing of e-nets to remain in register with the indexing for E- and U-tilings, thereby explicitly retaining links to precursor tilings of e-nets.

An uncharacteristically degenerate example is given by regular and semi-regular (4, 6) and (6, 6) tilings of the *P* and *D* surfaces. The U-tilings and E-tilings are illustrated in Fig. 22. Those four cases, derived from four distinct subgroup tilings (see Table 8), generate equivalent e-nets whose simplest embedding (identical to the s-net, *sqc947*) is illustrated in Fig. 23. Since the lowest E-tiling index found among these examples is *EDC1*, all four e-nets map to *edc1*. This example demonstrates the phenomenon of collapse of distinct U- and E-tilings to a single e-net.

A major issue remains unresolved here, which we hope to explore in detail later. In contrast to h- and s-nets, we have yet to produce an algorithm for determining equivalence or otherwise of e-nets. The problem is related to that of identification of distinct knots, a central topic of knot theory. We suspect that our problem can be approached by forming a canonical embedding for e-nets. For now, we can only determine 'by inspection' whether e-nets are ambient isotopic. Equivalent e-nets are also necessarily topologically equivalent; we therefore first check for graph isomorphisms by comparing the e-nets' topological structure. In many cases, that can be done by forming a barycentric embedding, giving

gluings across mirror lines incident to ⋆2 vertices. An example of such a tiling is shown in Fig. 21(*a*). The duals of these tilings are necessarily multigraphs, with a pair of edges linking the vertices corresponding to multi-edge sharing faces in the original tiling (Fig. 21*b*).

Alternatively, some h-nets that are themselves simple graphs give rise to e-net multigraphs as a result of the covering

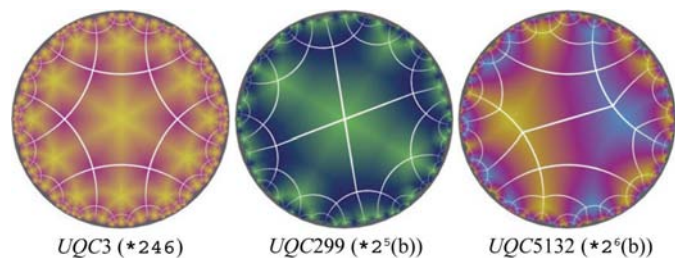UQC3 (⋆246)   UQC299 (⋆2⁵(b))   UQC5132 (⋆2⁶(b))

Figure 20
Three hyperbolic U-tilings with the same topology [Schäfli symbol (6, 4)] but distinct symmetries, highlighted by the background shading. All three examples have the same h-net, *hqc9*, the most symmetric embedding (⋆246) of a hyperbolic (6, 4) net.

an s-net, described in the next section. If the s-nets differ, the e-nets are necessarily distinct. If they are equivalent, we then determine – currently by eye – whether the e-net embeddings are ambient isotopic to the canonical embedding of their common s-net. If they are, both e-nets are equivalent. If one differs, they are inequivalent; if both differ we must look further to establish whether they can be mutually deformed without edge crossings to a common intermediate embedding. Some nets cannot be analysed *via* this method due to vertex 'collisions' that occur in the barycentric embedding. In those cases, comparison with an s-net is impossible. We must then resort to other measures to compare net topologies, including coordination sequences (Brunner & Laves, 1971) and
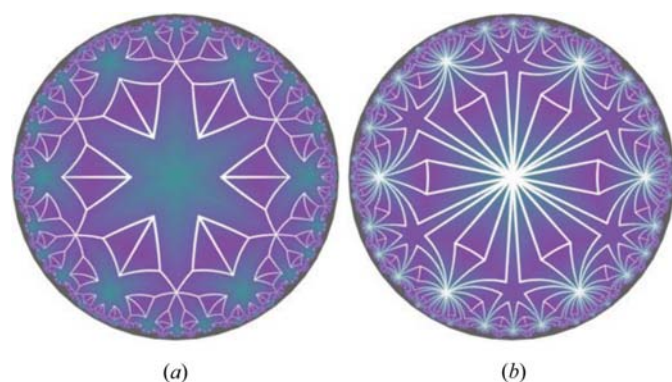
comparison of subgraphs, implemented in the *Topos* package (Blatov, 2006).

We close this section with an example, shown in Fig. 24, that illustrates many of the issues discussed here. The two e-nets shown in Fig. 24(*a*) are generated by tilings of the gyroid (that lie outside the Coxter class) and exhibit self catenation. The e-nets cannot be compared *via* their s-nets, due to vertex collisions in their respective barycentric embeddings. However, they share identical coordination sequences (to shell 24) and certain finite subgraphs strongly pointing to their topological equivalence. Their embeddings are *not* ambient isotopic, evidenced, for example, by the distinct link types highlighted in Fig. 24(*b*). These are the (2, 4) and (2, 6) torus links (Adams, 2004), distinguished by their distinct threadings through each other. These are thus distinct e-nets, despite their topological equivalence.

6.3. s-Nets

We have noted above that different E-tilings can define the same e-net (in Fig. 22) and that different e-nets may share a common periodic net topology (Fig. 24). In general there is no effective algorithm to determine when two crystal nets are topologically identical. However, for a large class of 'collision-free' Euclidean nets, the *Systre* algorithm furnishes a canonical form for the quotient graph of a net (Delgado-Friedrichs & O'Keeffe, 2003). We call these canonical forms *s-nets*.

The *Systre* algorithm [available as part of the *GAVROG* package (Delgado-Friedrichs, 2006)] is based on finding the barycentric embedding (or equilibrium placement) for the net: each vertex is located at the centre of mass of its edge-adjacent neighbours. This produces the highest-symmetry embedding of a periodic net, and permits the computation of the associated space group when there are no collisions. A *collision* occurs when two vertices have the same coordinates in the equilibrium placement (configurations that generate collisions are described later in this section). The final output of the algorithm is a *systre key* – a canonical representation of the labelled quotient graph for the periodic net that uses its smallest translational unit cell. The systre key provides a unique signature for topologically isomorphic crystal nets (Delgado-Friedrichs & O'Keeffe, 2003).

We derive systre keys and barycentric embeddings for our e-nets as a way to identify the distinct periodic net topologies generated by our enumeration. First, a quotient graph for the e-net is derived directly from an E-tiling *via* its cut ◦◦◦ D symbol. The e-net topology is adjusted slightly by

Figure 21
(*a*) An FSGG tiling with symmetry ⋆266. Notice that two edges are common to adjacent tiles, giving edge decorations composed of kites with a single interior edge. (*b*) The dual tiling of (*a*), composed of kites and stars. Edges of adjacent stars sharing two vertices are necessarily kinked to accommodate the interstitial elongated kite-shaped faces dual to the kites of (*a*). Those (vertex-free) edges make the h-net a multigraph.
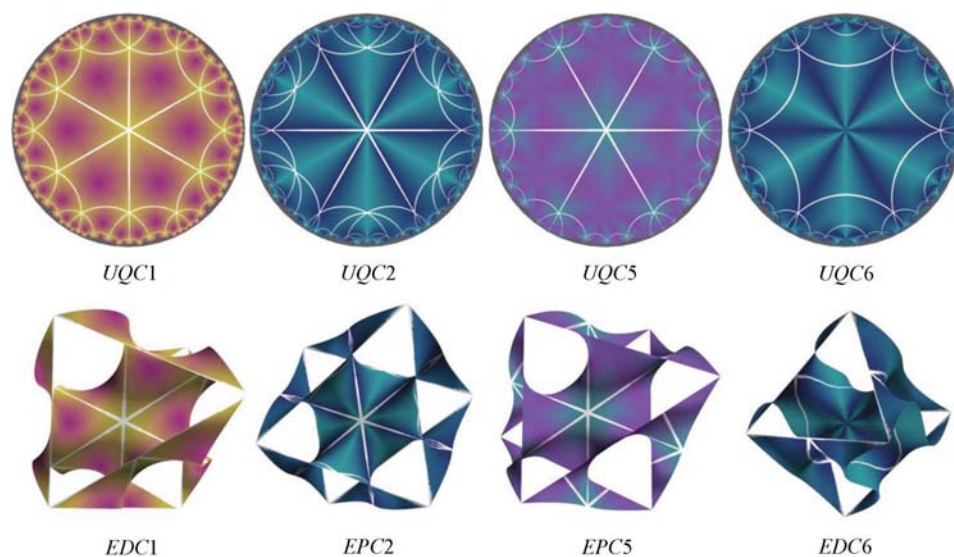

Figure 22
Some E-tilings of the *P* and *D* surfaces that lead to equivalent e-nets, whose simplest embedding is that of Fig. 23. The corresponding U-tilings have (4, 6) and (6, 6) topologies, corresponding to the hyperbolic nets *hqc*5 and *hqc*10, respectively (Table 8). Since *EDC*1 is the lowest index for these cases, all e-nets take the name *edc*1.

coalescing any multigraph edges. We then compute the equilibrium placement and systre key for the net where possible. Most of the nets enumerated *via* our hyperbolic tiling approach are collision-free: from 18 285 E-tilings (6095 U-tilings on each of the *P*, *G* and *D* surfaces) we found 2247 nets with equilibrium collisions. Thus, the systre key is an effective, although occasionally inadequate, tool for identifying the set of distinct network topologies derived from our surface reticulations. Indeed, most s-nets have a single e-net antecedent: of the 14 532 distinct systre keys, only 954 are obtained in more than one way.

The s-nets are ranked by their systre key and assigned distinct names of the form *sqcN*, where *s*, *q* and *c* have the same meanings as above, and the index *N* is derived from the systre-key ranking. This ordering is intuitively appealing, with a gradual increase in net 'complexity' (to first order, the number of vertices and edges in a primitive unit cell) as the index rises. For example, the first-ranked net according to this scheme, *sqc*1, is the standard net of the simple cubic lattice; the diamond net (whose e-net antecedent is a multigraph) is *sqc*6 *etc.*

We finish this section with a discussion of some problems that arise when analysing nets using barycentric embeddings. A minor cautionary example – a degree-6 (4, 6) tiling with orbifold $\star 2^5$(a) (*UQC*42) derived from a reticulation of the D surface, forming the (degree-6) e-net *edc*42 – is shown in Fig. 25. Here, the barycentric embedding of the net is free of vertex collisions but a pair of edges intersect, inducing apparent vertices of degree four: a situation that is due solely to the edge embedding in $\mathbb{E}^3$. Barycentric embedding of *edc*42 gives the s-net *sqc*900, which contains virtual degree-4 edge crossings in addition to the degree-6 vertices. These edge crossings in the s-net are problematic when determining the ambient isotopy classes of associated e-nets.

Vertex collisions are even more problematic, and can involve either small subsets of the net or triply periodic subgraphs. An example of the former are tilings that contain kites with a single interior edge like those in Fig. 21. The barycentric embedding of an associated e-net collapses the interior edge of the kite and places its two vertices at the same point. A more serious collapse is induced by 'ladder' graphs: examples composed of identical periodic nets linked by rung-

like edges. In those cases, the identical components collapse onto themselves entirely. An example is the graph illustrated in two distinct entangled forms in Fig. 24. The graph topology of both nets is a ladder graph built from a pair of regular degree-3 triply periodic graphs. Each component of the ladder is a graph known to chemists as the *srs* graph (O'Keeffe, 2008). As mentioned in §6.2 we must resort to other topological signatures such as coordination sequences to characterize nets that contain vertex collisions.

6.4. o-Nets

Our enumeration of tilings of the tritorus, known as O-tilings, can be extended to produce embeddings of finite graphs in $\mathbb{E}^3$, just as E-tilings induce e-nets. Recall from the introduction to this section that – like e-nets – o-nets are distinguished by their ambient isotopy class. Unlike the E–e map, a single O-tiling can induce an infinite number of distinct o-nets, due to the flexibility in forming a tritorus in $\mathbb{E}^3$. The nuances of the construction of o-nets from O-tilings are complex and remain to be explored in detail. Some discussion

Figure 24
(*a*) A pair of rhombohedral e-nets, with degree-3 and degree-4 vertices. (*b*, *c*) Subgraphs of these e-nets give the distinct two-component (2, 4) and (2, 6) torus links, shown as coloured cycles.

*edc*1

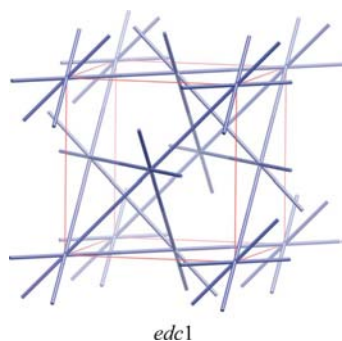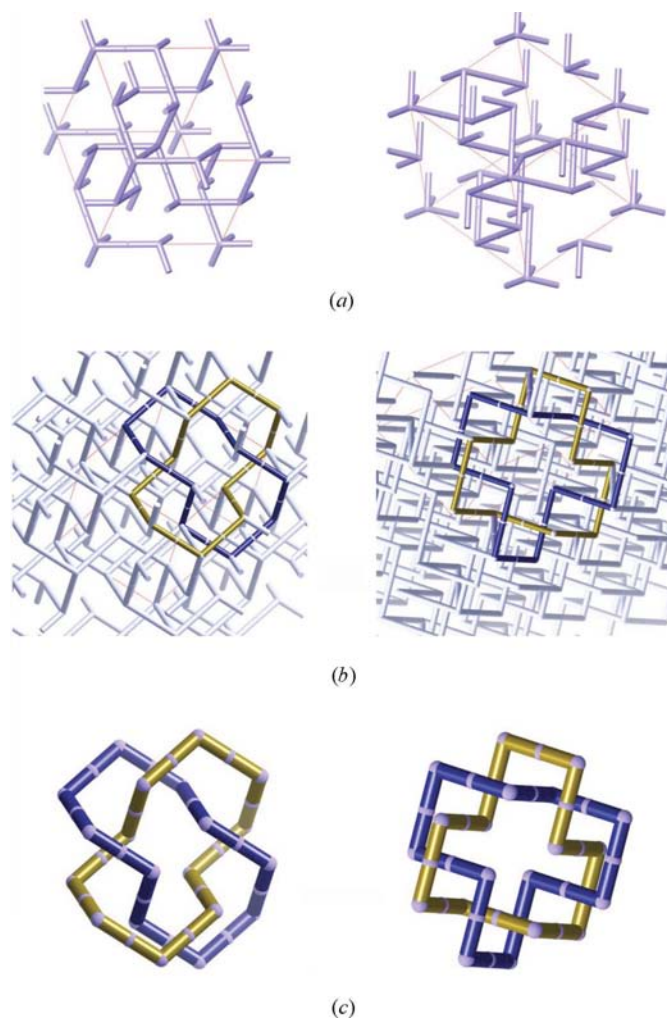**Figure 23**
The e-net formed by the E-tilings depicted in Fig. 22 – a cubic degree-6 structure with all edges and angles equal.

of the problem will be presented in a forthcoming paper which explores the construction of entangled polyhedral nets from tilings of the (genus-1) torus, to which we direct the interested reader (Castle *et al.*, 2009). In addition to the one-to-many map in going from O-tilings to o-nets, a many-to-one collapse is also possible: o-nets derived from distinct O-tilings may be equivalent, in the sense that one can be deformed into the other without any edge crossings. The O–o map is therefore many-to-many. We have yet to explore the taxonomy of o-nets in detail, but flag their presence here due to their strong links to the other surface reticulations.

7. A worked example

The path described in this paper – from Delaney–Dress symbols and their embedding within subgroups of $\star 246$ (U-tilings), then to E-tilings on the TPMS, giving e-nets and their canonical $\mathbb{E}^3$ embeddings as s-nets – is one that traverses aspects of geometry, group theory and tiling theory. It is therefore rather tortuous to navigate. With the formalities in place, we offer a detailed worked example – starting with a D symbol, and finishing with the triply periodic Euclidean nets obtained from reticulations of the P, D and G surfaces – to illustrate the connections between structures generated by our enumeration procedure.

We begin with a sequence of hyperbolic tilings in Fig. 26 that show the progression from the fundamental tile-transitive tiling for $\star 2244$, to a tile-2-transitive split and glued tiling, to its vertex-2-transitive dual (*cf.* §3). We adopt the vertex-2-transitive tiling as the starting point for our example. This tiling contains one vertex of degree 3 and one of degree 5, with four distinct tiles (three quadrilaterals and an octagon) arranged around each vertex to give the two-dimensional Schläfli symbol (4.8.8), (4.4.4.8.8). The combinatorial description of this tiling requires eight chambers and is given by the Delaney–Dress symbol in Table 9 (*cf.* Appendix *B*). This symbol is given the label $QC643$ according to the algorithm explained at the end of §3. Since this symbol is minimal –
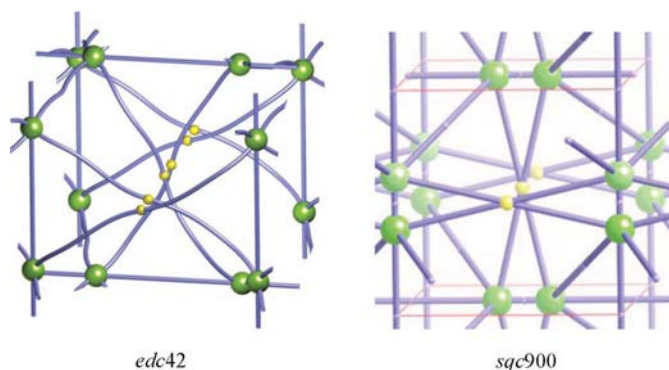

edc42   sqc900

Figure 25
The e-net *edc*42 embedded according to its edge geometry in the D surface. The midpoints of curved edges are marked by non-coincident yellow spheres. Barycentric embedding of *edc*42 in $\mathbb{E}^3$ gives the s-net *sqc*900. Pairs of curved edges of *edc*42 collide, giving edge crossings (yellow spheres) in addition to the degree-6 (green) vertices.

Table 9
The Delaney–Dress symbol for the vertex-2-transitive tiling $UQC1346$ shown in Fig. 26(*c*).

	a	b	c	d	e	f	g	h
0-*nbr*	c	b	a	d	e	f	g	h
1-*nbr*	b	a	d	c	e	g	f	h
2-*nbr*	a	f	c	e	d	b	h	g
m_{01}	8	8	8	8	4	4	4	4
m_{12}	5	5	3	3	3	5	5	5

Table 10
Periodic graph representations derived from our example tiling $UQC1346$ from the subgroup $\star 2244$(a).

The 16 translationally distinct vertices are labelled 0–15, with the choice of representatives as shown in Fig. 28. Each row defines one of the 32 distinct edges. The first two columns specify the vertex pairs that define an edge. The next four columns give the translation that needs to be applied to the second vertex to put it in the correct neighbouring cell. These translations are given for the hyperbolic net, then the e-nets on the P, D and G surfaces. The Euclidean translations are given with respect to the $\{\mathbf{a}, \mathbf{b}, \mathbf{c}\}$ bases defined in Table 2.

v_1	v_2	$\mathbb{H}^2$	P	D	G
0	1	I	000	000	000
0	2	I	000	000	000
0	4	I	000	000	000
0	8	I	000	000	000
0	8	t_2^{-1}	$0\bar{1}0$	$1\bar{1}0$	001
1	3	t_1	100	$10\bar{1}$	$0\bar{1}0$
1	5	t_3^{-1}	$00\bar{1}$	$01\bar{1}$	100
2	3	I	000	000	000
2	6	$\tau_2^{-1}t_3^{-1}$	$\bar{1}00$	010	$00\bar{1}$
2	10	I	000	000	000
2	10	$t_1^{-1}\tau_3^{-1}$	$0\bar{1}0$	001	$\bar{1}00$
3	7	$t_1^{-1}t_3^{-1}$	$\bar{1}0\bar{1}$	$\bar{1}10$	110
4	5	I	000	000	000
4	6	$\tau_2^{-1}t_3^{-1}$	$\bar{1}00$	010	$00\bar{1}$
4	12	I	000	000	000
4	12	$t_3\tau_1^{-1}$	010	001	$\bar{1}\bar{1}\bar{1}$
5	7	I	000	000	000
6	7	I	000	000	000
6	14	I	000	000	000
6	14	$t_3t_1t_3t_2^{-1}\tau_1^{-1}$	010	$1\bar{1}0$	$0\bar{1}0$
8	9	I	000	000	000
8	10	I	000	000	000
8	12	t_2	010	$\bar{1}10$	$00\bar{1}$
9	11	$t_2\tau_3^{-1}$	100	010	$\bar{1}\bar{1}\bar{1}$
9	13	τ_1^{-1}	$01\bar{1}$	010	$01\bar{1}$
10	11	I	000	000	000
10	14	$\tau_3t_1\tau_2^{-1}t_3^{-1}$	$\bar{1}10$	$01\bar{1}$	$10\bar{1}$
11	15	$\tau_3t_2^{-1}\tau_1^{-1}$	$\bar{1}1\bar{1}$	000	100
12	13	I	000	000	000
12	14	$t_2^{-1}\tau_3t_1\tau_2^{-1}t_3^{-1}$	$\bar{1}00$	$10\bar{1}$	100
13	15	I	000	000	000
14	15	I	000	000	000

i.e. it has the highest possible symmetry for a tiling with this topology – it defines an h-net, labelled *hqc*583 (*cf.* §6.1).

Our tiling can be embedded into two distinct subgroups of $\star 246$, as shown in Fig. 27. The next step is to unfold these tilings in their respective subgroup charts to obtain the corresponding U-tilings: $UQC1346$ and $UQC1345$ (*cf.* §5.1). The translational domain for the $\star 2244$(a) example is shown in Fig. 28.

We pause for a moment to look for other D symbols from our enumeration that generate the same U-tilings as our example. We find that $UQC1346$ is also the unfolding of a distinct D symbol, labelled $QC1442$, in the subgroup $\star2^6$(b). Further, $UQC1345$ is generated by $QC1442$, but *via* an embedding in $\star2^6$(a). The initial tiling ($QC643$, with symmetry $\star2244$) is related to the tiling $QC1442$, of symmetry $\star2^6$, by an additional symmetry that halves the area of the fundamental domain. Thus, $QC1442$ has the same h-net as $QC643$. If we now step sideways, we find that $QC1442$ has two distinct automorphic embeddings in the subgroup $\star2^6$(a), giving two additional U-tilings. And so we see that this h-net topology ($hqc583$) is in fact associated with two subgroup tilings, five embeddings in four kaleidoscopic subgroups and three U-tilings. These relationships are illustrated in Fig. 29.

We now return to the $\star2244$(a) example ($UQC1346$) and the periodic nets generated by its projection onto the P, D and G surfaces (*cf.* §6). First look at the periodic net as it sits in the hyperbolic plane. The net carried by $UQC1346$ has 16 translationally distinct vertices and 32 distinct edges, as shown in Fig. 28. With the vertex labels marked as shown in this figure,

we build the labelled quotient-graph description of the net given in Table 10. Edges between vertices that lie in different copies of the translational unit cell are labelled by the relative translation between the cells. We map this hyperbolic periodic net directly onto three e-nets (one on each TPMS) using the covering map actions defined in Table 2 of §2. The resulting triply periodic nets are also given in Table 10. Translational unit cells for the P and D surface tilings are shown in Fig. 30.

Lastly, we apply the *Systre* algorithm to the three e-nets to find a canonical form for the associated s-nets (*cf.* §6.3). Both the P and D e-nets reduce to crystal nets with only eight vertices per unit cell, rather than the 16 found in the e-nets. This means that the crystal nets can be symmetrized to display an additional translational symmetry, absent in the surface embedding (in fact, this translation is one that swaps sides of the surface). In contrast, the s-net derived from the e-net embedded in the G surface retains the full 16 vertices (because the G surface does not have an extra translational symmetry that swaps sides of the surface). Crystallographic descriptions of the resulting s-nets are given in Tables 11–13 and unit-cell images are shown in Figs. 31 and 32. Note that each s-net has two symmetrically distinct vertices and four distinct edges – an equivalent multiplicity to that of the $\star2244$(a) subgroup tiling.

Figure 26
(*a*) The fundamental tile-transitive tiling for the orbifold $\star2244$. (*b*) A tile-2-transitive tiling obtained *via* a split from a $\star4$ vertex to the opposite edge of the fundamental tile and two edge-glue operations, resulting in deletion of the edges in the fundamental tiling joining adjacent $\star6$ sites. (*c*) The corresponding vertex-2-transitive dual.
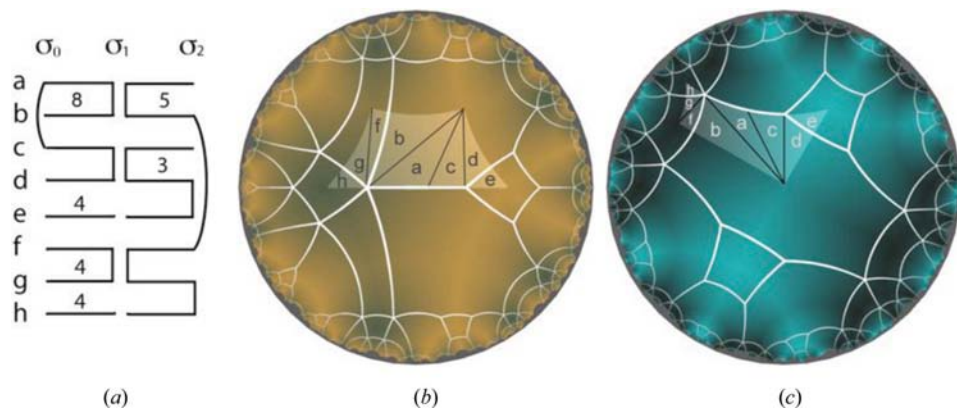
Figure 27
(*a*) The Conway crankshaft diagram (described in Appendix B) for the D symbol of Table 9, labelled $QC643$. (*b*) An embedding of this tiling in the $\star2244$(a) subgroup of $\star246$, the U-tiling $UQC1346$. (*c*) The same symbol embedded in $\star2244$(b), $UQC1345$. Single orbifold domains are highlighted and the chambers are labelled exactly as specified in the crankshaft diagram.

8. Future directions

The work presented here has focused on tilings and surface reticulations derived from the kaleidoscopic subgroups that are compatible with the genus-3 translational unit cells of the P, D and G minimal surfaces. There are many directions to extend our enumeration, including other subgroups compatible with the P, D and G surfaces; other triply periodic minimal surfaces; and further generalizations discussed below.

In the near future we intend to study nets derived from non-kaleidoscopic subgroups compatible with the P, D and G surfaces. These orbifold families include mixed reflection–rotation examples that we call 'hat' orbifolds, and pure rotational 'stellate' orbifolds (Hyde *et al.*, 2009). Given that there are 29 and 21 distinct subgroups of $\star246$ with hat and stellate orbifolds, respectively, the profusion of results is expected to be significant. Other orbifold classes, including the non-orientable 'projective' examples, may be explored at a
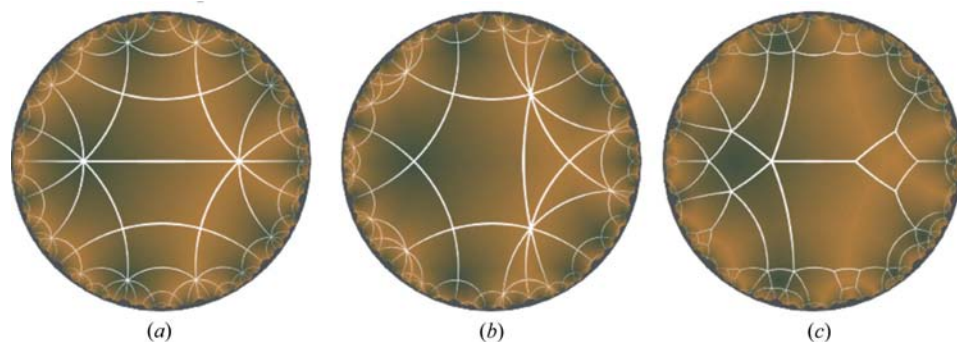
later date. We also plan to extend the project to explore nets generated as reticulations of other TPMSs, specifically the remaining genus-3 TPMSs: the hexagonal *H* and tetragonal *CLP* surfaces. We also plan to investigate this process on the
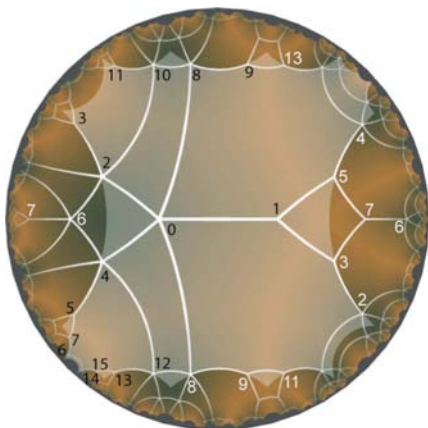
Figure 28
Our example tiling embedded in ★2244(a) is shown with a translational domain highlighted to represent the U-tiling *UQC*1346. The vertices in the translational domain are labelled 0–15 in black, consistent with the periodic graph description in Table 10. The vertices labelled in white are translated copies of those in the translational unit cell.

genus-4 cubic *I-WP* surface. The complications associated with extending to genus-4 examples are likely to be outweighed by the novelty of the examples. We have already determined the hyperbolic crystallography and covering maps for these surfaces and derived the relevant compatible orbifolds (Robins *et al.*, 2004*b*; Robins, 2006).

A parallel effort will be directed at generalizing the remarkably powerful Delany–Dress apparatus to allow enumeration of other tilings that are commonplace in the hyperbolic plane but which have no analogue in the Euclidean plane. Such tilings contain infinite-sided hyperbolic polygons arranged in ribbon- or tree-shaped patterns which we propose to call 'free tilings'. Specific examples that are commensurate with the *P*, *D* and *G* surfaces have been shown to lead to multiple intergrown nets (Hyde & Oguey, 2000; Hyde *et al.*, 2003); other examples are summarized online (Ramsden *et al.*, 2004). We intend to explore free tilings using an extension to Delaney–Dress tiling theory and thus achieve a systematic enumeration of complex net intergrowths.

In addition to the triply periodic structures described here, our techniques readily adapt to the enumeration of *finite* (*molecular*) *nets via* reticulations of compact surfaces. We have already mentioned these O-tilings and o-nets in §§5.3 and 6.4. Specific embeddings of these nets are a rich topic, barely
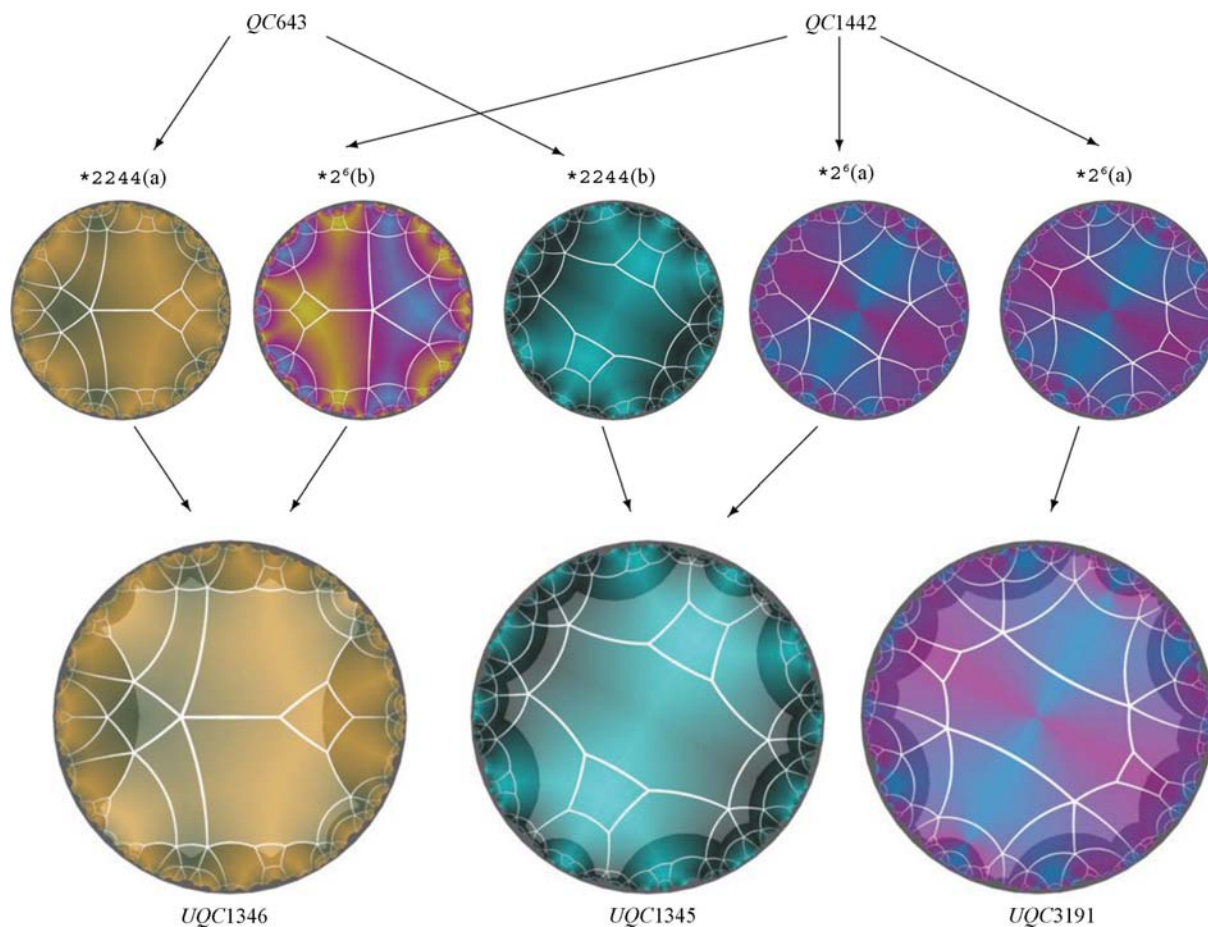

Figure 29
This diagram shows the relationships between two subgroup tilings *QC*643 and *QC*1442 that lift to the same h-net, and their various embeddings and unfoldings in ★246 subgroups. The three U-tilings are *UQC*1346, *UQC*1345 and *UQC*3191.

explored to date. By defining a covering map from the hyperbolic plane onto an explicit embedding of the tritorus in Euclidean space, we can use the structure of the surface reticulation to define an embedding of the o-net. Generic examples will be knotted, linked and ravelled, governed by the wealth of possible cycle homotopies on the tritorus. A sample of the possibilities can be seen in a study of knotted toroidal polyhedra, generated *via* reticulations of a torus (Hyde & Schröder-Turk, 2007; Castle *et al.*, 2009), as well as a recent exploration of ravelled graphs (Castle *et al.*, 2008).

9. Discussion of results

As discussed in §6, our approach enumerates a range of nets from infinite symmetric two-dimensional hyperbolic h-nets and infinite triply periodic Euclidean e- and s-nets, to finite o-nets.
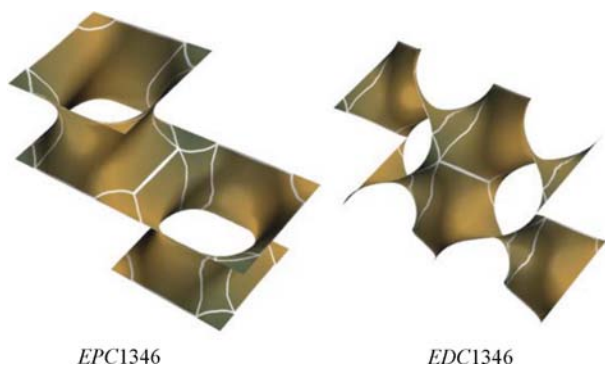

Figure 30
Our example tiling *UQC*1346 projected onto the *P* and *D* surfaces. The surface patches are (nonstandard) translational unit cells that correspond to the highlighted region in the hyperbolic plane shown in Fig. 28.
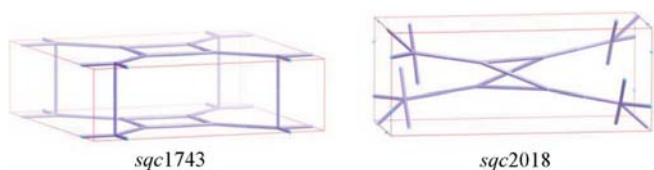

Figure 31
The two s-nets derived from projections onto the *P* and *D* surfaces (left and right, respectively) of the ⋆2244(a) U-tiling *UQC*1346.
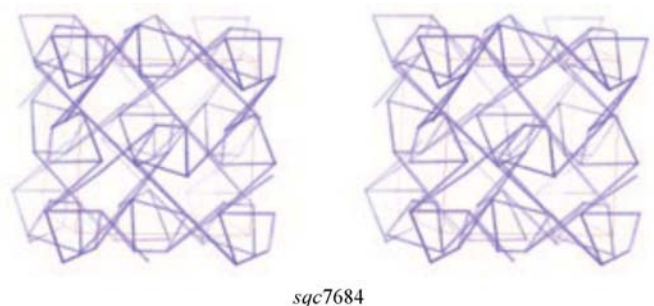

Figure 32
Stereo pair for the s-net derived from the ⋆2244(a) U-tiling *UQC*1346 projected onto the *G* surface.

Table 11
Crystallographic description of the s-net derived from the *UQC*1346 tiling projected onto the *P* surface, forming the *EPC*1346 tiling.

This information can be accessed online at http://epinet.anu.edu.au/sqc1743. Symmetry group *P*4/*mmm*, *a* = 3.41421, *c* = 1.0 Å.

Vertex	Degree	x	y	z
v_1	5	0.14645	0.14645	0.00000
v_2	3	0.35355	0.35355	0.00000

Edges Start x, y, z			End x, y, z		
0.1464	0.1464	0.0000	−0.1464	0.1464	0.0000
0.1464	0.1464	0.0000	0.1464	0.1464	−1.0000
0.1464	0.1464	0.0000	0.3535	0.3535	0.0000
0.3535	0.3535	0.0000	0.3535	0.6464	0.0000

Table 12
Crystallographic description of the s-net derived from the *UQC*1346 tiling projected onto the *D* surface, forming the *EDC*1346 tiling.

This information can be accessed online at http://epinet.anu.edu.au/sqc2018. Symmetry group $P\bar{4}2m$, *a* = 1.41228, *c* = 3.60072 Å.

Vertex	Degree	x	y	z
v_1	5	0.23691	0.23691	0.10319
v_2	3	0.36278	0.36278	0.37199

Edges Start x, y, z			End x, y, z		
0.2369	0.2369	0.1032	−0.2369	0.7631	0.1032
0.2369	0.2369	0.1032	0.2369	−0.2369	−0.1032
0.2369	0.2369	0.1032	0.3628	0.3628	0.3720
0.3628	0.3628	0.3720	0.6372	0.3628	0.6280

Table 13
Crystallographic description of the s-net derived from the *UQC*1346 tiling projected onto the *G* surface, forming the *EGC*1346 tiling.

This information can be accessed online at http://epinet.anu.edu.au/sqc7684. Symmetry group $I\bar{4}2d$, *a* = 2.30597, *c* = 2.31613 Å.

Vertex	Degree	x	y	z
v_1	5	0.10435	0.14742	0.17446
v_2	3	0.14603	0.40658	0.92807

Edges Start x, y, z			End x, y, z		
0.1043	0.1474	0.1745	−0.1474	0.3956	0.4245
0.1043	0.1474	0.1745	0.1474	−0.1043	−0.1745
0.0934	0.8540	0.6781	0.3956	1.14742	0.5755
0.0934	0.6460	0.5719	0.1460	0.40658	0.9281

The Euclidean s-nets formed from low-transitivity tilings of the hyperbolic plane afford an interesting variety of canonical, symmetric embeddings of distinct network topologies in three dimensions. These data are directly comparable with complementary enumeration schemes pursued by (principally) structural chemists, including O'Keeffe's Reticular Chemistry Structural Resource (RCSR) database of three-

dimensional nets and tilings (O'Keeffe, 2008; O'Keeffe *et al.*, 2008) and the Hypothetical Zeolite database of Treacy and colleagues (Foster & Treacy, 2008). These collections comprise nets in three-dimensional Euclidean space constructed using three-dimensional tilings and symmetry principles. We note that our 14 532 embedded s-nets derived from tilings of Coxeter orbifolds include *circa* 150 examples found in the RCSR database, which contains more than 1500 distinct structures. Among those common examples are 14 known zeolite structures. Therefore, this first pass has already furnished around 10% of nets considered to be of relevance to reticular chemistry. More than 200 examples of our s-nets also appear in the Hypothetical Zeolite database (Foster, 2008).

Many more structures of actual and/or potential chemical interest are certain to emerge by moving beyond Coxeter orbifolds and by reticulating noncubic surfaces. However, we hasten to point out that our motivation for this work goes beyond enumeration of chemically interesting patterns; rather we are primarily interested in exploring the variety of nets that emerge without imposition of chemical constraints. Evidently, we are forced to impose filters to avoid the inevitable combinatorial explosion associated with such a search; we have opted to filter the nets according to their two-dimensional hyperbolic symmetries.

Our approach is governed by the simplicity of tiling enumeration in two dimensions as opposed to three, along with the suspicion that symmetric patterns in two-dimensional hyperbolic space yield symmetric patterns in three-dimensional space. The latter feature has been (and will be) discussed elsewhere; some statistics on the variety of three-dimensional symmetries and vertex transitivity of s-nets can be read in Hyde *et al.* (2006). Our goal, however, has not been to offer a circuitous route to what is already enumerated elsewhere. Rather, we hope that this route generates new examples not readily deduced by more conventional three-dimensional approaches.

We find, for example, a vertex-2-transitive (12, 4) tiling generated in the ⋆2626 orbifold within the FSGG class of tilings. Projection of this subgroup tiling onto the gyroid, *via*

the U-tiling *UQC*104, gives an e-net, *egc*104, whose barycentric form is a rhombohedral ($R\bar{3}c$) vertex-1-transitive (uninodal) s-net in three-dimensional Euclidean space, *sqc*906, labelled *usf* by O'Keeffe (2008). This structure has been identified in molecular frameworks (Moulton *et al.*, 2003), yet three-dimensional tiling theory cannot find a 'proper' tiling (Blatov *et al.*, 2007) for this pattern.

A second example is the sphere packing 3/10/*h*4 (Sowa & Koch, 2006), named *wiw* by O'Keeffe (2008). This structure too emerges as an s-net, *sqc*3054 *via* the degree-3 FSGG subgroup tiling with ⋆2626 symmetry illustrated in Fig. 33. That tiling unfolds to the U-tiling *UQC*262, which maps onto the gyroid to form an E-tiling, whose e-net, *egc*262, is topologically equivalent to *sqc*3054, also illustrated in Fig. 33. The three-dimensional tiling algorithm *TOPOS* (Blatov, 2006) fails to find a simple tile for this structure (Blatov, 2007), due to the threading of short rings in the net by other edges. This threading precludes the possibility of those rings spanning faces of a simple three-dimensional tile.

We expect the most powerful aspect of our enumeration technique to emerge over time: namely the possibility of finding distinctly tangled embeddings of topologically isomorphic nets. It is that possibility that has encouraged us to retain *e-nets* as a distinct class, since their edges wind according to the tiling of the TPMS, in contrast to barycentric *s-net* embeddings. The latter nets relax to a canonical embedding that may lose the knottedness and generic entanglements of the original e-net.

The complete results of our enumeration are collated in an online database, accessible at http://epinet.anu.edu.au, which we urge interested readers to explore at their leisure. Ultimately, we expect this growing database of nets to provide a substantial foundation for a range of investigations into the physical features of nets, extending current work on percolation, transport and elastic responses (Gibson & Ashby, 1997; Roberts & Garboczi, 2002; Durand, 2005; Durand & Weaire, 2004). This suite of examples will allow detailed exploration of possible correlations between physical, topological and geometric features of crystal nets.

APPENDIX *A*
Orbifolds

An *orbifold* encodes the symmetric properties of an infinitely repeating two-dimensional pattern with the topological features of a compact, connected surface. The concept has its roots in the theory of discrete groups, but the name was coined by Thurston and his students by combining *orbit* and *manifold*, since each point of an orbifold represents the entire orbit of a point under the group of symmetry operations of the pattern. Alter-
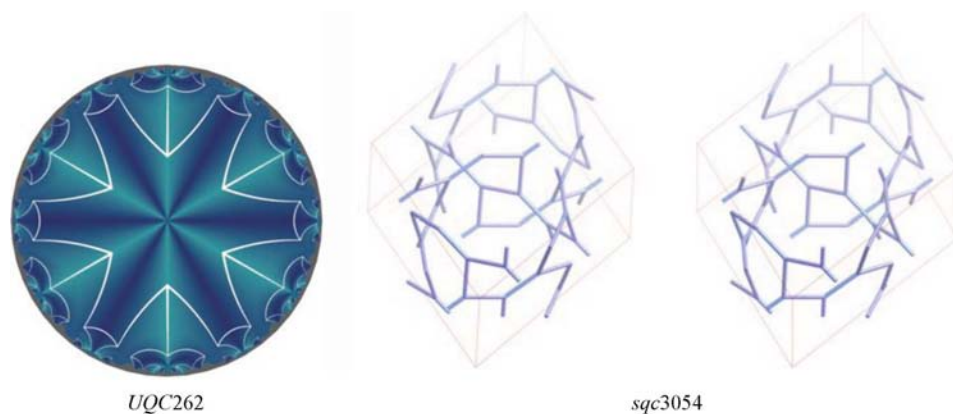

Figure 33
Left: hyperbolic tiling with orbifold symmetry ⋆2626. Right: stereo pair of the tiling projected onto the gyroid, relaxed to form the s-net *sqc*3054, also known as *wiw* (O'Keeffe, 2008), whose vertices define the centres of equal spheres in the rhombohedral 3/10/*h*4 sphere packing.

*UQC*262   *sqc*3054

natively, one can invert that process and generate a symmetric pattern in the universal cover of the orbifold by 'rolling' the inked orbifold in all possible directions on the relevant two-dimensional homogeneous space – the sphere, $\mathbb{S}^2$, Euclidean, $\mathbb{E}^2$, or hyperbolic, $\mathbb{H}^2$, planes – printing the pattern as it goes. Technically, the orbifold is the *quotient space* of the pattern by its symmetry group, and the pattern is the *universal cover* of the orbifold.

Two-dimensional symmetries include rotation about a point, reflection in a mirror line, translation and glide reflection. In the orbifold a reflection produces a boundary component, a rotation induces a cone point, while translations and glides that do not arise from other symmetries of the pattern are encoded by global topological features such as rings, handles and cross-caps. These relationships are described in more detail below.

A *cone point* is just what its name suggests. Rolling a cone about its apex will generate a pattern with rotational symmetry at that point. For the pattern to repeat exactly, the angle swept out by the cone must be an integer subdivision of 2π. An orbifold containing just three cone points looks like a samosa, or triangular pillow. If the only orbifold features are two cone points, then each must be of the same order, as they represent opposite poles of a spherical pattern. There are no compact orbifolds consisting of a single cone point.

Now suppose the orbifold has a puncture, and therefore a boundary edge. When rolling the orbifold surface over its universal cover, in order to continue beyond the edge the orbifold side in contact with the covering space must flip so the pattern is locally reversed. Boundaries therefore encode mirror reflections. A boundary must be a closed loop, which may contain corners. A *corner* in the orbifold is formed by two mirror lines meeting at an angle less than π. As with cone points, for the pattern to unwrap and meet itself exactly, only corner angles that are integer subdivisions of π are allowed. Note that cone and corner points are local geometric features that distinguish orbifolds from true manifolds.

Translational symmetries may be induced by combinations of reflections or rotations. Those translations that are not induced by other symmetries are encoded by a nontrivial loop in the orbifold such as formed by a ribbon with its ends glued together, or around the two axes of a torus. For example, a finite cylinder with two boundary loops can be rolled across the Euclidean plane mapping out a translationally periodic pattern in one direction normal to its axis. Its boundaries represent mirror lines so extension in the other direction of the Euclidean plane is induced by parallel reflection lines.

A glide reflection is the (irreducible) composition of a translation with a reflection, encoded within an orbifold by a non-orientable surface feature such as a Möbius strip or cross-cap. A cross-cap has the topology of the real projective plane, and can be modelled by a disc with opposite points on the boundary identified. For example, one traversal along the centre of a Möbius strip brings you to the same point on the opposite side of the ribbon (the flipped image of the initial pattern); a second traversal is needed to come back to the starting point on the original side. This doubled traversal is identical to the translation induced by applying a glide reflection twice.

We can encode independent symmetry operations of any two-dimensional pattern on the sphere, Euclidean and hyperbolic planes by distinct surface features of its orbifold; namely boundaries, cone and corner points, cross-caps and handles. This result depends on the central theorem of 2-manifold topology that every closed, compact, two-dimensional manifold is topologically equivalent to either the sphere, a sphere with n handles attached, or a sphere with n cross-caps.

A simple notation for orbifolds is due to John Conway (Conway, 1992; Conway & Huson, 2002). From it we can reconstruct the orbifold surface topology and all its associated symmetry features. The notation is unique up to certain rearrangements of elements. The global topological features are the handles, cross-caps and boundaries, and are denoted by the symbols $\circ$, $\times$ and $\star$. A handle transforms to two cross-caps in the presence of another cross-cap ($\circ\times = \times\times\times$). Thus, by convention, handles and cross-caps are quarantined and an orbifold symbol will either be prefixed by $\circ$'s or suffixed by $\times$'s. The nonmanifold features (cone and corner points) have specific angles that are represented by their integer fractions of 2π and π, respectively.

As there is no inherent ordering to nonboundary points of a manifold, the cone points may be listed in any order. Thus, the orbifold symbol 457 is equivalent to 754 or even 475 – in fact any permutation is valid. We typically use lexical ordering to make a nice canonical form, but this is not intrinsic to the notation. Corner points, however, have a distinct ordering based on their sequence around a boundary component. A *mirror string* is given by a single $\star$, representing the boundary component, followed by a list of its corner points (if any). Here the order matters, $\star$4567 is different to $\star$5467. However, because there is no intrinsic start or end to a mirror boundary, any cyclic permutation is equivalent, *i.e.* $\star$5674 and $\star$6745 represent the same orbifold as $\star$4567. Also, because a punctured surface has no intrinsic inside or outside, one can reverse the order of the corner points: *i.e.* $\star$4567 is equivalent to $\star$7654. Entire mirror strings correspond to punctures and, like cone points, there is no way to intrinsically order these features. Therefore the order of the mirror strings does not matter.

The kaleidoscopic symmetry groups considered in this paper are generated by reflections only, and are therefore examples of Coxeter groups. Their orbifolds have the form of a single mirror string, $\star m_1 m_2 \ldots$.

We give the complete lexical specification of an orbifold symbol to summarize the above discussion: [some number of $\circ$'s][cone points in any order][mirror strings in any order][some number of $\times$'s] written as

$$\circ \circ \circ \ldots c_1 c_2 c_3 \ldots \star m_1 m_2 \ldots [\star m_4 m_5 \ldots] \ldots \times \times \times.$$

Note that any string that meets the above specification corresponds to a genuine orbifold except for those of the form c, $c_1 c_2$ where $c_1 \neq c_2$, $\star$m, and $\star m_1 m_2$ where $m_1 \neq m_2$. Note also

that we contract strings of identical numerals using power notation to improve readability: *e.g.* 22222 will be written 2^5.

Orbifolds have an associated scalar value, computable from the symbol elements, called the *curvature index* (also called the cost or characteristic). For kaleidoscopic groups it is

$$\chi(\star m_1 m_2 \ldots) = 1 - \sum_i (m_i - 1)/2m_i.$$

The general formula is given in Conway & Huson (2002). Its value determines which of the three two-dimensional geometries the symmetry pattern must belong to: positive implies a spherical symmetry group, zero implies $\mathbb{E}^2$ and negative implies $\mathbb{H}^2$. In the case of spherical or hyperbolic geometry, the curvature also quantifies the area of a single copy of the orbifold in the relevant plane of unit Gaussian curvature. Euclidean space is unique in that shapes may be scaled arbitrarily whilst preserving angles, so a Euclidean orbifold has no associated area. The majority of all two-dimensional orbifolds are hyperbolic, thus the variety of hyperbolic two-dimensional tilings far exceeds that of Euclidean or spherical space.

Table 14
The Delaney–Dress symbol for the tiling shown in Fig. 34.

	a	b	c	d	e	f	g	h	i	j
0-*nbr*	b	a	d	c	e	g	f	i	h	j
1-*nbr*	a	c	b	e	d	f	h	g	j	i
2-*nbr*	a	b	g	f	e	d	c	h	i	j
m_{01}	5	5	5	5	5	5	5	5	5	5
m_{12}	4	4	4	3	3	3	4	4	8	8

APPENDIX *B*
Delaney–Dress symbols

We discuss and illustrate tilings of the hyperbolic plane here, see Fig. 34 for example, but the concepts apply to the sphere and Euclidean plane with only minor changes, and the underlying theory generalizes to higher-dimensional spaces.

The following four conditions form the definition of a tiling of the hyperbolic plane:

(1) The tiles are closed topological discs.

(2) Tiles intersect only along their boundaries. The intersection of two tiles defines an edge, the intersection of three or more tiles defines a vertex.

(3) The tiles are uniformly bounded in size.

(4) The tiles cover the whole hyperbolic plane.

To describe a tiling pattern we subdivide the tiles into triangles, called *flags* or *chambers*, and then record the neighbour relations of the different symmetry classes of these chambers. To generate a chamber system from a tiling, we make a barycentric subdivision by placing a 2-vertex in the centre of each tile, a 1-vertex at the midpoint of each edge and a 0-vertex at each tiling vertex, then form 0-1-2-triangles within each tile. The edges of the chambers are also labelled 0-, 1- and 2-edges according to the type of vertex they face. The neighbour relations are formally described by the action of three maps, σ_0, σ_1 and σ_2, that map each chamber to its neighbour across the corresponding edge (or equivalently, opposite the corresponding vertex). For obvious geometric reasons, these maps are involutions (they are their own inverse).

The topology of the tiling is encoded by describing what happens on repeated application of pairs of the neighbour maps. There are three orbits to consider:
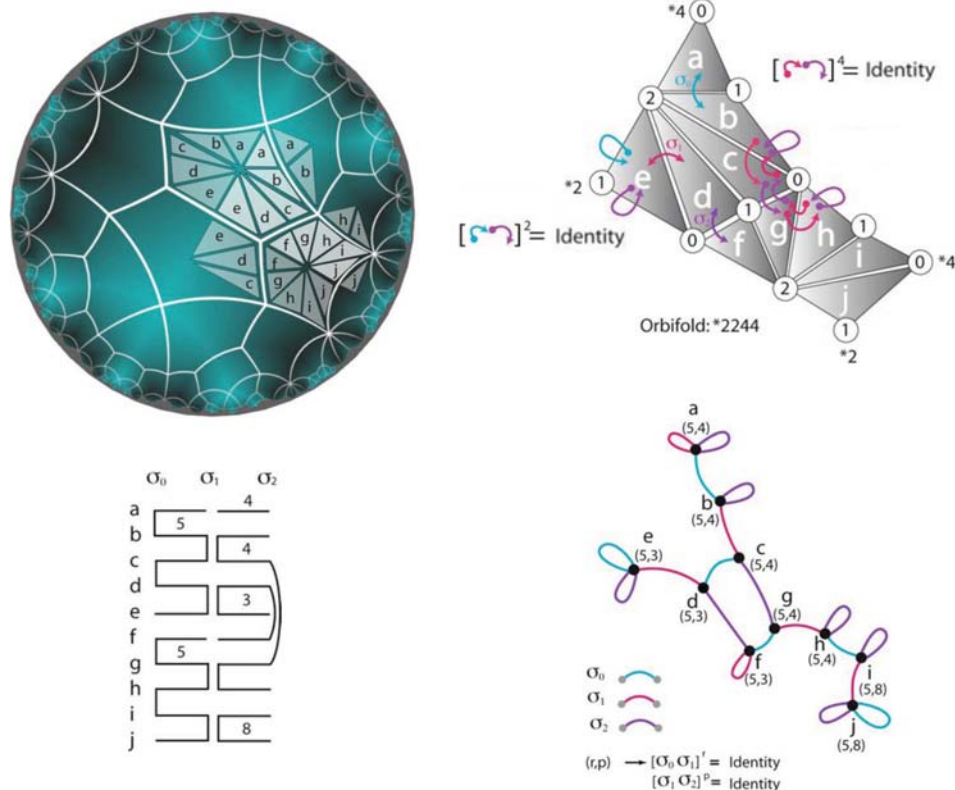
Figure 34
An example hyperbolic tiling and chamber system illustrating Delaney–Dress tiling theory. At the upper left is a tiling of the hyperbolic plane with symmetry ⋆2244, consisting of two types of pentagonal tile. Each symmetrically distinct chamber appears once in the orbifold domain of ⋆2244, shown at the upper right. The vertices of the chambers are labelled 0, 1 or 2, according to whether they lie on a tile vertex, edge or centre, respectively, and the neighbour maps, σ_i, encode the adjacencies of chambers opposite a vertex of type *i*. A Delaney–Dress symbol can be depicted (lower right) as an edge-coloured graph with topological index pairs (r, p) added to each node. An alternative visualization is a Conway crankshaft diagram, shown at the lower left, where the connected components indicate distinct tile or vertex orbits, and orbit indices need only be listed once for each component.

Table 15
The $\star 246/\!\circ\!\circ\!\circ$ chart.

No.	Label	R_1-nbr	R_3-nbr	R_2-nbr
0	I	1	2	3
1	R_1	0	4	5
2	R_3	6	0	7
3	R_2	5	8	0
4	R_1R_3	9	1	10
5	R_1R_2	3	11	1
6	R_3R_1	2	12	13
7	R_3R_2	13	14	2
8	R_2R_3	15	3	16
9	$R_1R_3R_1$	4	17	18
10	$R_1R_3R_2$	18	19	4
11	$R_1R_2R_3$	20	5	21
12	$R_3R_1R_3$	22	6	23
13	$R_3R_1R_2$	7	24	6
14	$R_3R_2R_3$	25	7	26
15	$R_2R_3R_1$	8	27	28
16	$R_2R_3R_2$	28	26	8
17	$R_1R_3R_1R_3$	29	9	30
18	$R_1R_3R_1R_2$	10	31	9
19	$R_1R_3R_2R_3$	32	10	33
20	$R_1R_2R_3R_1$	11	34	35
21	$R_1R_2R_3R_2$	35	33	11
22	$R_3R_1R_3R_1$	12	36	37
23	$R_3R_1R_3R_2$	37	38	12
24	$R_3R_1R_2R_3$	39	13	40
25	$R_3R_2R_3R_1$	14	41	42
26	$R_3R_2R_3R_2$	42	16	14
27	$R_2R_3R_1R_3$	43	15	44
28	$R_2R_3R_1R_2$	16	45	15
29	$R_1R_3R_1R_3R_1$	17	46	47
30	$R_1R_3R_1R_3R_2$	47	48	17
31	$R_1R_3R_1R_2R_3$	49	18	50
32	$R_1R_3R_2R_3R_1$	19	51	52
33	$R_1R_3R_2R_3R_2$	52	21	19
34	$R_1R_2R_3R_1R_3$	53	20	54
35	$R_1R_2R_3R_1R_2$	21	55	20
36	$R_3R_1R_3R_1R_3$	46	22	56
37	$R_3R_1R_3R_1R_2$	23	57	22
38	$R_3R_1R_3R_2R_3$	$t_2^{-1}51$	23	58
39	$R_3R_1R_2R_3R_1$	24	59	60
40	$R_3R_1R_2R_3R_2$	60	58	24
41	$R_3R_2R_3R_1R_3$	$t_3 48$	25	61
42	$R_3R_2R_3R_1R_2$	26	62	25
43	$R_2R_3R_1R_3R_1$	27	$t_1^{-1}56$	63
44	$R_2R_3R_1R_3R_2$	63	64	27
45	$R_2R_3R_1R_2R_3$	65	28	66
46	$R_1R_3R_1R_3R_1R_3$	36	29	67
47	$R_1R_3R_1R_3R_1R_2$	30	68	29
48	$R_1R_3R_1R_3R_2R_3$	$t_3^{-1}41$	30	69
49	$R_1R_3R_1R_2R_3R_1$	31	70	71
50	$R_1R_3R_1R_2R_3R_2$	71	69	31
51	$R_1R_3R_2R_3R_1R_3$	$t_2 38$	32	72
52	$R_1R_3R_2R_3R_1R_2$	33	73	32
53	$R_1R_2R_3R_1R_3R_1$	34	$t_1^{-1}67$	74
54	$R_1R_2R_3R_1R_3R_2$	74	75	34
55	$R_1R_2R_3R_1R_2R_3$	76	35	77
56	$R_3R_1R_3R_1R_3R_2$	67	$t_1 43$	36
57	$R_3R_1R_3R_1R_2R_3$	$t_2^{-1}70$	37	78
58	$R_3R_1R_3R_2R_3R_2$	$t_2^{-1}72$	40	38
59	$R_3R_1R_2R_3R_1R_3$	$t_3 68$	39	79
60	$R_3R_1R_2R_3R_1R_2$	40	80	39
61	$R_3R_2R_3R_1R_3R_2$	$t_3 69$	81	41
62	$R_3R_2R_3R_1R_2R_3$	$\tau_2^{-1}t_1 75$	42	82
63	$R_2R_3R_1R_3R_1R_2$	44	$t_1^{-1}78$	43
64	$R_2R_3R_1R_3R_2R_3$	$t_1^{-1}\tau_3^{-1}73$	44	83
65	$R_2R_3R_1R_2R_3R_1$	45	$\tau_2^{-1}t_3^{-1}79$	84
66	$R_2R_3R_1R_2R_3R_2$	84	83	45
67	$R_1R_3R_1R_3R_1R_3R_2$	56	$t_1 53$	46
68	$R_1R_3R_1R_3R_1R_2R_3$	$t_3^{-1}59$	47	85
69	$R_1R_3R_1R_3R_2R_3R_2$	$t_3^{-1}61$	50	48
70	$R_1R_3R_1R_2R_3R_1R_3$	$t_2 57$	49	86

Table 15 (continued)

No.	Label	R_1-nbr	R_3-nbr	R_2-nbr
71	$R_1R_3R_1R_2R_3R_1R_2$	50	87	49
72	$R_1R_3R_2R_3R_1R_3R_2$	$t_2 58$	88	51
73	$R_1R_3R_2R_3R_1R_2R_3$	$\tau_3 t_1 64$	52	89
74	$R_1R_2R_3R_1R_3R_1R_2$	54	$t_1^{-1}85$	53
75	$R_1R_2R_3R_1R_3R_2R_3$	$t_1^{-1}\tau_2 62$	54	90
76	$R_1R_2R_3R_1R_2R_3R_1$	55	$\tau_3 t_2^{-1}86$	91
77	$R_1R_2R_3R_1R_2R_3R_2$	91	90	55
78	$R_3R_1R_3R_1R_2R_3R_2$	$t_2^{-1}86$	$t_1 63$	57
79	$R_3R_1R_2R_3R_1R_3R_2$	$t_3 85$	$t_3 \tau_2 65$	59
80	$R_3R_1R_2R_3R_1R_2R_3$	$\tau_1 87$	60	92
81	$R_3R_2R_3R_1R_3R_2R_3$	$t_3 \tau_1^{-1}t_2^{-1}88$	61	93
82	$R_3R_2R_3R_1R_2R_3R_2$	$\tau_2^{-1}t_1 90$	93	62
83	$R_2R_3R_1R_3R_2R_3R_2$	$t_1^{-1}\tau_3^{-1}89$	66	64
84	$R_2R_3R_1R_2R_3R_1R_2$	66	$\tau_2^{-1}t_3^{-1}92$	65
85	$R_1R_3R_1R_3R_1R_2R_3R_2$	$t_3^{-1}79$	$t_1 74$	68
86	$R_1R_3R_1R_3R_1R_3R_2$	$t_2 78$	$t_2 \tau_3^{-1}76$	70
87	$R_1R_3R_1R_2R_3R_1R_2R_3$	$\tau_1^{-1}80$	71	94
88	$R_1R_3R_2R_3R_1R_3R_2R_3$	$t_2 \tau_1 t_3^{-1}81$	72	95
89	$R_1R_3R_2R_3R_1R_2R_3R_2$	$\tau_3 t_1 83$	95	73
90	$R_1R_2R_3R_1R_3R_2R_3R_2$	$t_1^{-1}\tau_2 82$	77	75
91	$R_1R_2R_3R_1R_2R_3R_1R_2$	77	$\tau_3 t_2^{-1}94$	76
92	$R_3R_1R_2R_3R_1R_2R_3R_2$	$\tau_1 94$	$t_3 \tau_2 84$	80
93	$R_3R_2R_3R_1R_3R_2R_3R_2$	$t_3 \tau_1^{-1}t_2^{-1}95$	82	81
94	$R_1R_3R_1R_2R_3R_1R_2R_3R_2$	$\tau_1^{-1}92$	$t_2 \tau_3^{-1}91$	87
95	$R_1R_3R_2R_3R_1R_3R_2R_3R_2$	$t_2 \tau_1 t_3^{-1}93$	89	88

(1) The $(\sigma_0\sigma_1)$ orbit maps around a 2-vertex, and so visits the chambers in a single tile. If the tile has r edges, then $(\sigma_0\sigma_1)^r$ is the identity map on each chamber of the tile. The index r is also called m_{01}.

(2) The $(\sigma_1\sigma_2)$ orbit maps around a 0-vertex, so walks around the chambers that meet at a vertex of the tiling. If the vertex has degree p then $(\sigma_1\sigma_2)^p$ is the identity for each chamber incident at that vertex. The index p is also called m_{12}.

(3) Finally, $(\sigma_2\sigma_0)$ maps around a 1-vertex, *i.e.* an edge of the tiling. Since exactly four chambers meet at a 1-vertex, we have that $(\sigma_2\sigma_0)^2$ is the identity for every chamber.

The chamber system described so far is an infinite complex because infinitely many bounded tiles are needed to cover the hyperbolic plane. We obtain a finite description of the tiling by forming equivalence classes of chambers under the action of the symmetry group. If the group has a compact orbifold, then there will be a finite number of chamber classes. The σ_i maps preserve these symmetry classes, so only a finite number of neighbour relations need to be recorded. These chamber classes, their neighbour maps, and the topological indices r and p defined above are all the information needed to form the Delaney–Dress symbol.

We illustrate the above definitions with an example, shown in Fig. 34. This hyperbolic tiling is built from two types of pentagonal tile, with vertices of degree 3, 4 and 8, and has symmetry $\star 2244$. On the upper left of Fig. 34 each pentagonal tile is shown subdivided into triangular chambers. Labels are assigned so that symmetrically equivalent chambers have the same letter. There are ten symmetrically distinct chambers and these cover a single copy of the $\star 2244$ orbifold (shown upper right). The chamber vertices are labelled 0, 1, 2, according to whether they lie on a tile vertex, edge or centre, respectively. The neighbour maps σ_i that encode the adjacency of chambers

opposite a vertex of type *i* are also shown in Fig. 34, upper right.

The complete set of adjacency relations for the chambers and the topological indices may be given in tabular form, see Table 14, or in two visual formats shown in Fig. 34. At the lower right of Fig. 34 the D symbol is shown as a graph where the nodes represent each chamber class and coloured edges denote the σ_i involutions. The topological indices (r, p) for each chamber are also attached to each node.

A more concise depiction is a *crankshaft diagram* shown in Fig. 34, lower left, another elegant notation due to John Conway. In this representation, each chamber is a broken horizontal line, and the adjacencies between chambers are given by vertical connections in the appropriate σ_i column (self-adjacency is indicated by an open endpoint). The topology of the $\sigma_i\sigma_j$ chamber orbits is immediately apparent from this crankshaft diagram. The $\sigma_0\sigma_1$ tile orbits are represented by connected components in the left part of the crankshaft, and the $\sigma_1\sigma_2$ vertex orbits on the right. The $\sigma_0\sigma_2$ edge orbits are only implicitly represented in this diagram. Each symmetrically distinct tile and vertex is defined by a single connected component of the crankshaft, so the corresponding orbit numbers (r, p) are listed only once. The transitivity is visually clear – in our tile-2-transitive example there are two connected $\sigma_0\sigma_1$ components, while the four different orbit components in the $\sigma_1\sigma_2$ column tell us the tiling is vertex-4-transitive. A further advantage of the crankshaft diagram is that the dual tiling is obtained simply by reflecting the diagram about the central axis (in our two-dimensional case, exchanging the σ_0 and σ_2 columns).

The power of Delaney–Dress symbols is that any two tilings with identical topology and symmetry will have isomorphic symbols, and the tiling can be completely reconstructed from this finite amount of information. In addition, the orbifold symbol and curvature index can be computed directly from the Delaney–Dress symbol; see Delgado-Friedrichs (2003) for further details.

APPENDIX C
The $\star 246/\circ\circ\circ$ chart

Here we give the complete description of the $\star 246/\circ\circ\circ$ chart that forms the basis from which we derive all the unfolded $\circ\circ\circ$ D symbols; see §4 for further details.

The full labelling of triangles and their neighbours in the $\star 246/\circ\circ\circ$ chart is given in Table 15. This list is generated from the *kbmag* word enumeration with short-lex ordering. The generator order here is R_1, R_3, R_2, since this gave a slightly more balanced region in the hyperbolic plane than the ordering R_1, R_2, R_3. The domain covered by the elements given here is not the semi-regular dodecagon illustrated in Fig. 10, but is equivalent to it. The triangles with the 96 words from Table 15 define an irregular shape illustrated in Fig. 35. The underlying group and triangulation structure on the tritorus is exactly the same as that given by the dodecagon with opposite sides glued.
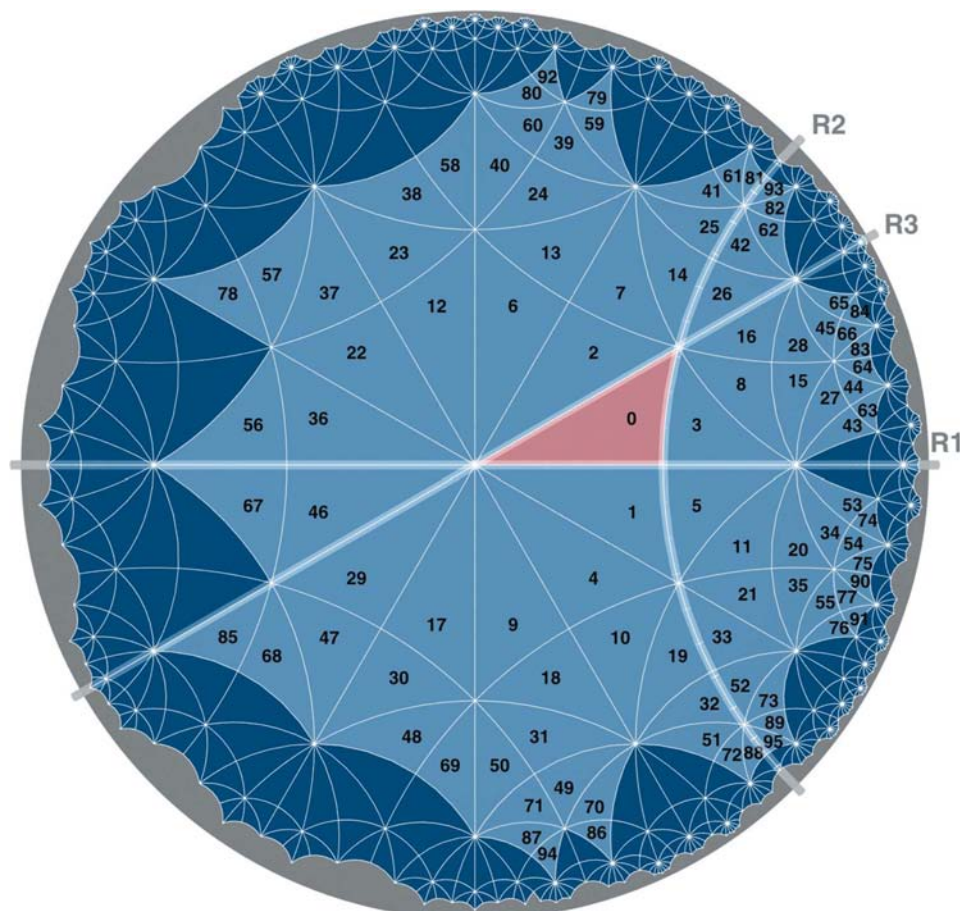

Figure 35
The 96 translationally distinct triangles in the $\star 246/\circ\circ\circ$ chart. The labels refer to the words in Table 15.

References

Adams, C. C. (2004). *The Knot Book: An Elementary Introduction to the Mathematical Theory of Knots.* Providence: American Mathematical Society.

Balke, L. & Huson, D. H. (1996). *Geom. Dedic.* **60**, 89–106.

Banerjee, R., Phan, A., Wang, B., Knobler, C., Furukawa, H., O'Keeffe, M. & Yaghi, O. M. (2008). *Science*, **319**, 939–943.

Barabasi, A. L. (2002). *Linked: The New Science of Networks.* Cambridge: Perseus.

Batten, S. R. (2001). *Cryst. Eng. Commun.* **3**, 67–73.

Blatov, V. A. (2006). *Topos*, http:/// www.topos.ssu.samara.ru. Described

in the online newsletter of the International Union of Crystallography Commission on Crystallographic Computing, No. 7, http://www.iucr.org/resources/commissions/crystallographic-computing/newsletters.

Blatov, V. A. (2007). Personal communication.

Blatov, V. A., Carlucci, L., Ciani, G. & Proserpio, D. M. (2004). *Cryst. Eng. Commun.* **6**, 377–395.

Blatov, V. A., Delgado-Friedrichs, O., O'Keeffe, M. & Proserpio, D. M. (2007). *Acta Cryst.* A**63**, 418–425.

Brunner, G. O. & Laves, F. (1971). *Wiss. Z. Tech. Univ. Dresden*, **20**, 387–390.

Castle, T., Evans, M. E. & Hyde, S. T. (2008). *New J. Chem.* **32**, 1484–1492.

Castle, T., Robins, V. & Hyde, S. T. (2009). In preparation.

Conway, J. H. (1992). *Groups, Combinatorics, and Geometry*, Vol. 165, pp. 438–447. Cambridge University Press.

Conway, J. H. & Huson, D. H. (2002). *Struct. Chem.* **13**, 247–256.

Delgado-Friedrichs, O. (1994). *Euclidicity criteria for three-dimensional branched triangulations.* Doctoral thesis, University of Bielefeld, Germany.

Delgado-Friedrichs, O. (2003). *Theor. Comput. Sci.* **303**, 431–445.

Delgado-Friedrichs, O. (2006). *GAVROG*, http://www.gavrog.org.

Delgado-Friedrichs, O., Dress, A. W. M., Huson, D. H., Klinowski, J. & Mackay, A. L. (1999). *Nature (London)*, **400**, 644–647.

Delgado-Friedrichs, O. & Huson, D. H. (1999). *Discrete Comput. Geom.* **21**, 299–315.

Delgado-Friedrichs, O. & O'Keeffe, M. (2003). *Acta Cryst.* A**59**, 351–360.

Dress, A. W. M. (1987). *Adv. Math.* **63**, 196–212.

Dress, A. W. M. & Huson, D. H. (1987). *Geom. Dedic.* **24**, 295–310.

Durand, M. (2005). *Phys. Rev. E*, **72**, 011114.

Durand, M. & Weaire, D. (2004). *Phys. Rev. E*, **70**, 046125.

Foster, M. D. (2008). Personal communication.

Foster, M. D. & Treacy, M. J. (2008). *Atlas of Prospective Zeolite Frameworks*, http://www.hypotheticalzeolites.net/.

Gibson, L. J. & Ashby, M. F. (1997). *Cellular Solids – Structure and Properties*, 2nd ed. New York: Cambridge University Press.

Grünbaum, B. & Shephard, G. C. (1987). *Tilings and Patterns*. New York: W. H. Freeman.

Holt, D. F. (1998). *KBMAG – Knuth–Bendix for Monoids and Automatic Groups* (http://www.maths.warwick.ac.uk/~dfh). Version 2.4. University of Warwick, Coventry, UK.

Huson, D. H. (1993). *Geom. Dedic.* **47**, 269–296.

Hyde, S. T., Delgado-Friedrichs, O., Ramsden, S. J. & Robins, V. (2006). *Solid State Sci.* **8**, 740–752.

Hyde, S. T., Larsson, A.-K., Di Matteo, T., Ramsden, S. J. & Robins, V. (2003). *Austr. J. Chem.* **56**, 981–1000.

Hyde, S. T. & Oguey, C. (2000). *Eur. Phys. J. B*, **16**, 613–630.

Hyde, S. T., O'Keeffe, M. & Proserpio, D. M. (2008). *Angew. Chem. Intl Ed. Engl.* **47**, 7996–8000.

Hyde, S. T., Robins, V. & Ramsden, S. J. (2009). In preparation.

Hyde, S. T. & Schroeder, G. E. (2003). *Curr. Opin. Collect. Interf. Sci.* **8**, 5–14.

Hyde, S. T. & Schröder-Turk, G. E. (2007). *Acta Cryst.* A**63**, 186–197.

International Zeolite Association (2008). *Database of Zeolite Structures*, http://www.iza-structure.org/databases/.

Klee, W. E. (2004). *Cryst. Res. Technol.* **39**, 959–968.

Li, Y., Yu, J. & Xu, R. (2008). *AlPO Database*, http://mezeopor.jlu.edu.cn/alpo/.

Molnar, E. (2002). *Struct. Chem.* **13**, 267–275.

Moulton, B., Abourahma, H., Bradner, M. W., Lu, J., McManus, G. J. & Zaworotko, M. J. (2003). *Chem. Commun*. pp. 1342–1343.

Ockwig, N. W., Delgado-Friedrichs, O., O'Keeffe, M. & Yaghi, O. M. (2005). *Acc. Chem. Res.* **38**, 176–182.

Öhrström, L. & Larsson, K. (2005). *Molecule-Based Materials: The Structural Network Approach.* Amsterdam, Boston: Elsevier Science.

O'Keeffe, M. (2008). *RCSR: Reticular Chemistry Structure Resource*, http://rcsr.anu.edu.au.

O'Keeffe, M. & Hyde, B. G. (1996). *Crystal Structures I. Patterns and Symmetry*. Washington: Mineralogical Society of America.

O'Keeffe, M., Peskov, M. A., Ramsden, S. J. & Yaghi, O. M. (2008). *Acc. Chem. Res.* **41**, 1782–1789.

Ramsden, S. J., Robins, V. & Hyde, S. T. (2004). *2D Hyperbolic tilings by infinite tiles*, http://epinet.anu.edu.au/infinite_tiles.

Ramsden, S. J., Robins, V. & Hyde, S. T. (2005). *EPINET: Euclidean Patterns In Non-Euclidean Tilings*, http://epinet.anu.edu.au.

Roberts, A. P. & Garboczi, E. J. (2002). *J. Mech. Phys. Solids*, **50**, 33–55.

Robins, V. (2006). Unpublished.

Robins, V., Hyde, S. T. & Ramsden, S. J. (2004a). *Eur. Phys. J. B*, **39**, 365–375.

Robins, V., Hyde, S. T. & Ramsden, S. J. (2004b). *Physica A*, **339**, 173–180.

Robins, V., Ramsden, S. J. & Hyde, S. T. (2005). *Eur. Phys. J. B*, **48**, 107–111.

Sadoc, J.-F. & Charvolin, J. (1989). *Acta Cryst.* A**45**, 10–20.

Sowa, H. & Koch, E. (2006). *Acta Cryst.* A**62**, 379–399.

Stillwell, J. (1989). *Mathematics and Its History.* New York, Berlin: Springer-Verlag.

Strong, R. T., Pickard, C. J., Milman, V., Thimm, G. & Winkler, B. (2004). *Phys. Rev. B*, **70**, 045101.

The GAP Group (2002). *GAP - Groups, Algorithms, and Programming*, Version 4.3, http://www.gap-system.org.

Treacy, M. M. J., Randall, K. H., Rao, S., Perry, J. A. & Chadi, D. J. (1997). *Z. Kristallogr.* **212**, 768–791.

Watts, D. J. (2003). *Six Degrees: The Science of a Connected Age.* New York: Norton.

Wells, A. F. (1977). *Three-dimensional Nets and Polyhedra.* New York: Wiley.

Zou, X., Conradsson, T., Klingstedt, M., Dadachov, M. S. & O'Keeffe, M. A. (2008). *Nature (London)*, **202**, 69–88.