

Intricate Isohedral Tilings of 3D Euclidean Space

Carlo H. Séquin
CS Division, University of California, Berkeley
E-mail: sequin@cs.berkeley.edu

Abstract

Various methods to create intricate tilings of 3D space are presented. They include modulated extrusions of 2D Escher tilings, free-form deformations of the fundamental domain of various 3D symmetry groups, highly symmetrical polyhedral toroids of genus 1, higher-genus cage structures derived from the cubic lattice as well as from the diamond and triamond lattices, and finally interlinked tiles with the connectivity of simple knots.

1. Introduction

In celebration of the Bridges 2008 conference at the birthplace of M.C. Escher, this paper examines several topics concerning isohedral tilings in 3-dimensional space. (These are tilings where any tile A can be taken to any other tile location B with a symmetry operation on the whole tiling.) The first and most direct extension of Escher's work is to design shapes of genus zero that tile 3-dimensional space without voids, and which may resemble some “natural” object, such as a fish or a flower. Construction and visualization of such shapes is quite a bit more challenging than generating 2D tiles; thus we rely on computer assistance.

In a second part, we study tiles of simple polyhedral geometry, but of genus larger than zero. First, we look for toroidal tiles (genus 1) that are derived from some crystal lattices with a high degree of symmetry. These ring-tiles topologically interlock so as to “knit together” 3D space. Solutions with as few as 3 tiles and as many as 20 tiles interlinked with a single toroidal ring have been found.

Subsequently we construct tiles of even higher genus with a high degree of symmetry, ranging from simple cubic frame structures to “cages” made by assembling multiple toroidal ring-tiles. Finally we look for tiles that have the connectivity of knots that interlink with several of their neighbors.

We discuss several different paradigms to design these various 3D tiles and the prototype tools that we have built for that purpose. The application of these methods has resulted in many different types of tiles, which have been fabricated on a rapid prototyping machine (Fig.1).



Figure 1: Various types of isohedral 3D space tiles: (a) flat extruded 2.5D tiles, (b) morphing between two tile shapes, (c) natural forms grown by simulated annealing, (d) interlinked toroidal tiles, (e) tiles of higher genus, (f) a tile with the geometry of a trefoil knot.

2. Tiles of Genus 0

The first idea that comes to mind when pondering the idea how to make “3D Escher Tiles,” is to take one of M.C. Escher’s beautiful tiling patterns, extrude it into the 3rd dimension (Fig.1a), and then stack several such layers on top of one another, – as inspired, for instance, by the “Sardines” puzzle [6]. However, this rates as too simplistic a solution in our context, and such a tiling may be counted as a “2.5D Escher tiling” at best. We expect all our tile surfaces to be sculpted in non-trivial ways in all 3 dimensions. Clearly this is a much harder task in 3D than in 2D, and no good 3D design tools seem to exist. In order to make our initial domain of exploration not overwhelmingly large, we limit ourselves to isohedral tilings, so that we have to deal only with a single tile shape and a single set of adjacency relationships.

Conceptually it is clear how one should create such shapes. One needs to find the fundamental domain of a suitable tiling; identify corresponding edge segments that transform into one another under the symmetry operations allowed for the chosen tiling; edit the unique representatives of all such edges and faces to produce some interesting shape; and then apply the same edits to all sibling boundary elements that emerge from the primary ones by the allowed symmetry operations.

There are dozens of computer programs, some with very nice and convenient user interfaces, that assist a tile designer for the 2D case; but I am not aware of any nice and convenient tools that do this job in 3D space. Why is the 3D case so much harder? When we display a 2D tiling pattern on a computer screen, then everything we need to see is right in front of us. When we are in “edit mode” and grab a point on the tile boundary and move it around on the screen by using the mouse, there is a direct unambiguous correlation between our mouse movements and the action we expect to see on the screen. This makes editing and visualization of a 2D tile easy.

It is much harder to visualize the complete geometric shape of a 3D tile displayed on a computer screen; we either have to use multiple projections of it, or need to be able to interact with it in real time so that we can quickly inspect it from all sides. Editing 3D shapes is also non-trivial, since we have to control the movement of vertices or other surface elements in 3 dimensions, typically using only a 2-dimensional input device, and view the results on a 2-dimensional display. In order to see how the tiles fit together, we have to display several of them in close proximity; but this will then obscure exactly those portions where we might need to make some fine-tuned adjustments. Clearly, in this task, a computerized tool not only reduces the chore of drawing repeatedly the same shapes occurring in a regular tiling pattern, but it is also absolutely crucial for maintaining the constraints on the boundary of the 3D tile that guarantee that the end-result will indeed fill 3D space without voids. In the next few subsections we discuss some approaches and prototype tools that we have used to create such 3D tiles.

Modulated Extrusions

Since it is rather difficult to create nice looking shapes via a direct deformation of a 3-dimensional fundamental domain, we first discuss an approach that builds directly on the creative work of M.C. Escher. We start from one of his isohedral 2D tilings, and explore whether they are suitable to make non-trivial 3D shapes with modulated thickness, if additional layers are added but with some translatory offset or suitable rotations. For this purpose we analyze the “natural” thickness of the indicated shapes at various locations, and then design a “height map” overlaid on the 2D pattern (Fig.2a). These fish clearly want to be thickest in the region just behind the eye and in the middle of the body (dark green areas), and much thinner in the region of the tail fins (light pink areas). We then try to find a suitable rigid-body transformation for the next layer that can gracefully match the thick and thin regions in a complementary manner. In this case it requires a shift of half a fish-length from one layer to the next. With the thick areas nicely accommodated in the thin areas of the two adjacent layers, we then use a simple height-editing tool to form smooth transitions from one thickness region to another, so that a recognizable 3D fish-form results. Using layered manufacturing we can create physical fish-tiles, which, of course, may also be decorated further to make them look more attractive (Fig.2b).




Figure 2: (a) A 2D array of fish with overlaid height map; (b) extruded fish shapes, without and with surface modulation; (c) possible thick (red) and thin (cyan) regions in pattern #119 by M.C. Escher.

The same principle can also be applied to more intricate 2D starting patterns. If we start with Escher’s fish pattern #119 with 4-fold rotational symmetry [7, p.214] (Fig.2c), we might cluster islands of maximum thickness around the places where 4 fish heads and 4 tail fins join (if we make the tail fins mostly vertical). These areas can then be overlaid onto the areas with minimal thickness around the points where four right-hand fins join. The regions where two left-hand fins join would be assigned a “medium” thickness and would then be positioned on top of one another (with a 90° rotation).

Layered Tiles

To move another step closer to true 3D tiles, we might build layered tiles like wedding cakes. In the simplest case, we start with a 2D tiling pattern that uses two different tiles (Fig.3a), say, a larger one (e.g. an octagon) and a smaller one (e.g. a square that fills the holes between the octagons). We can form a single isohedral tile by gluing the smaller (extruded) shape on top of the larger one. Matthias Goerner, a student in my CS285 class in Fall 2007, came up with a fancy variant of this approach by adding twisted dove-tail joints. The larger shape is a set of mutually touching circles in a square array; the smaller shape is the remaining 4-arc shape between them. The result is a nifty serrated cog-wheel tile, which can be assembled, layer by layer, using a helical screw motion (Fig.3b-d). 2D tiles, even if they have heavily serrated, convoluted boundaries, can always be assembled by moving them through 3D space and “dropping” them into place. In 3 dimensions it is an extra challenge to design tiles that can physically hold together in 3D space, but which nevertheless can be assembled when added in the right order.




Figure 3: (a) Two-tile pattern. M. Goerner’s cog-wheel tiles: (b) a single tile; (c) assembly of four; (d) inserting the tenth tile into an assembly of nine tiles.

Metamorphoses

More “integrated” 3D tiles can be obtained, if we don’t just glue extruded 2D tile shapes together, but instead smoothly change the geometry from one tile shape into another. Again we take two complementary tile shapes that together tessellate the whole 2D plane; but now we apply a sweep-morph operation that gradually transforms one shape into the other. The object obtained by a full morph cycle, in which shape A transforms into shape B and then back again into shape A, forms an isohedral tile of 3D

space (Fig.1b and Fig.4a-c). Of course, for all the tiles discussed in the last two subsections, we always have the option to modulate the top and bottom surfaces with some compatible height field.




Figure 4: (a) *Morphed test shapes;* (b) *assembly of 8 such morph tiles;* (c) *metamorphosis of a bird into a fish.* (d) *Fish tile generated by simulated annealing (A. Megacz).*

3. True 3D Tiles

The most general way to create a 3D isohedral tile is to take the fundamental domain of some chosen tiling and deform its boundary, while observing all the symmetries associated with the chosen tiling. To explore this approach, Mark Howison helped me to build a few simple prototype tools.

The most generic 3D fundamental domain is based on a deformed truncated octahedron (Fig.5a). Each tile shares a face with 14 neighbors. If we assume a tiling with only triclinic translational symmetries, we obtain the most generic isohedral tiling of 3D space. It corresponds to Grünbaum’s *IH1* for planar tilings [4]. In 2D we can identify 93 different marked isohedral tilings. Kaplan and Salesin present parametrizations for all of them [5]. In 3D there are many more possibilities, given that there are 230 lattice groups in 3D, compared to only 17 in 2D. We did not attempt to define parametrizations for all of them, but just picked two very general examples cases.




Figure 5: *Editing the basic cell of a lattice with only translational symmetries: (a) based on a truncated octahedron; (b) deformed fundamental domains; (c) domain based on a rhombic dodecahedron.*

Rather than giving the user the minimal number of parameters to control, we provide a more hierarchical editing approach with a few higher-level, but somewhat redundant controls. These first define the translational symmetries of the lattice; then adjust the remaining tiling vertices that fully define the fundamental domain (Fig.5b); then the user can add extra vertices and fine tune their positions to create the desired tile geometry; finally “decorations” can be added. In Escher’s drawings these decorations are just line patterns or colors that facilitate the semantic interpretation of the tile shape. When making physical 2D tiles (Fig.1a), or when mapping the tiling onto a sphere [11], one can use embossing, specified by a height field, to create suitable decorations. Such geometrical decorations are more difficult on the surface of 3D tiles that have to fit together snugly. For instance, a protruding eye ball in a fish will leave a dimple at that location in the neighbor tile. We have not yet spent much time looking for clever complementary decorations. In a few instances, we simply have painted the tile surface.

The second fundamental domain we studied is based on the densest sphere packing, where each cell has 12 neighbour cells. For the fully symmetrical case, the fundamental domain is a rhombic dodecahedron. Again we added several slider-controlled parameters that allowed us to drastically deform this basic shape without altering the connectivity between adjacent tiles

At “Level-1” of our deformation modules, we allow the three base vectors of the cubic lattice to be changed into a base for an arbitrary triclinic lattice. This allows us, for example, to make strongly elongated tiles, if we want to represent, say, a skinny fish.

At “Level-2” we allow the user to further deform the fundamental domain by individually moving an independent set of the dodecahedron corners (i.e., the tiling vertices) and of its face mid-points, while the program makes sure that opposite faces in the fundamental domain automatically maintain identical geometry. The level-2 controls break the rhombic dodecahedron faces into four non-coplanar triangles.

At “Level-3” we also allow the edge midpoints of the dodecahedron to be adjusted. Each quadrilateral contact surface between two adjacent 3D tiles now has a total of 9 control points. So, in one version of the tool, we allow each such face to become a quadratic Bézier patch. Since the three quadrilateral faces that join in a single edge all refer to the same three control points for that edge, the edge is defined unambiguously as a quadratic Bézier curve, which then serves as the boundary edge for the three patches connected to it. In total, eight edges of the rhombic dodecahedron can be individually modified. Figure 5c shows some fish-like forms that can be constructed with this level of editing.

Tile Design by Simulated Annealing

Since manually editing the shape of the fundamental domain is rather difficult, we also tried to automate the shape-matching process, using simulated annealing, as inspired by Kaplan’s pioneering work [5]. A manually specified fish form served as the goal shape in an iterative process that subdivided the boundary of the fundamental domain and then tried to minimize the Hausdorff distance to the goal shape. Initial results (Fig.1c, 4d) were unsatisfactory; the tile surfaces produced were too noisy, and when subjected to a smoothing post-processing step, they lost all interesting details contained in the original geometry.

4. Tiles of Genus 1

In 3D we can do something that cannot be done in 2D: We can create tiles that not only interlock geometrically – but interlink topologically. For that purpose, tiles must be of genus 1 or higher, and the “handles” on these tiles are then interlinked. We start our explorations with simple ring-shaped tiles of genus 1; subsequently we extend our study to geometries of higher genus.




Figure 6: (a) Cubic lattice of linked rings, (b) 24-facet tile, (c) split 16-facet tile, (d) two 12-facet tiles.

The Basic Ring Tiles

In 1995, I created a 3D tiling of interlinked toroidal rings, based on the body-centered cubic lattice (BCC), in which each ring-tile is interlinked with 4 other rings (Fig.6a). One way of thinking of these shapes is as a polyhedral approximation of the Voronoi zones associated with square wire frames around every possible cube face in the two lattices. To avoid singularities, we uniformly shrink each of these wire

squares by a tiny amount, before finding the locus of all points in space closer to it than to any other wire square. In 2D, the locus of equal distances between a point and a line is a parabola. In 3D, the locus between two skewed lines is a saddle surface – a hyperbolic paraboloid. In the tiles used in my original cubic assembly, these saddle regions are approximated with 4 small quadrilaterals, leading to a tile with a total of 24 planar facets (Fig.6b).

There are various approximations that can be made for these saddle regions (Fig.6b-d). The simplest one is just the mid-plane parallel to both lines. This leads to a toroidal tile with only 12 planar faces (Fig.6d). This tile was pointed out to me by John Conway in 2007. Later, at the Math-Art 2007 conference, he showed me a simple and direct procedure to find these minimal 4-segment ring-tiles, and taught me [2] that the very same procedure also leads to symmetrical polyhedral rings with 6 and 10 segments when starting from the diamond and *triamond* lattices [1], respectively, rather than from the cubic lattice. The *triamond* lattice is an intriguing, often overlooked lattice, also known as the (10,3) lattice [10], closely related to the diamond lattice. But in this lattice each “atom” is linked in a regular manner to only three equidistant nearest neighbors, 120° apart. The shortest loops of these connections comprise 10 links in a wiggly ring configuration (Fig.7c).

The construction of these toroids starts with finding the shortest edge-rings in these three lattices; they have n edges, with $n = 4, 6, 10$, respectively. Correspondingly, the toroidal tiles will have 4, 6, and 10 prismatic segments. Each of the three lattices can be interspersed with a complement lattice in which the “atoms” are placed “farthest away” from the atoms of the primary lattice. In each lattice, exactly one of the edges of the complement lattice will serve as a symmetry axis or as a “hub” or “axle” for these shortest edge-rings. In a fully tiled lattice, n complementary toroidal tiles will join along this hub axis, while simultaneously interlocking with the primary ring tile. To obtain the simplest polyhedral geometry for these interlocking ring tiles, one forms the n tetrahedra between the hub segment and each one of the edge-rim segments in turn, and then splits each tetrahedron with the mid plane between the two skew edges used to construct the tetrahedron. The volume closer to the edge-rings forms one of the prismatic rim segments of the primary ring tile; and the other half of the split tetrahedron forms a segment for one of the n complementary rings that interlock with the primary one. Figures 7a,b show the resulting 6-segment ring tiles resulting from the diamond lattice. Figures 7c,d show the results for the triamond lattice. Note that the rings of the primary lattice and those of the complement lattice are of opposite chirality. The coordinates for the ring geometry were kindly sent to me by Chaim Goodman-Strauss [3].




Figure 7: (a) Two 6-segment ring-tiles and (b) an assembled 3D array based on the diamond lattice; (c) a split 10-segment ring-tile and (d) an array based on the triamond lattice.

3-Segment Ring-Tiles

We have discussed ring-tiles with 4, 6, and 10 segments that interlink and fill 3D space in a highly symmetrical and regular manner. What about ring-tiles with different number of segments, such as 5 or 8 or 12? In particular, I wondered, whether the simplest conceivable such tile – with only three segments – could serve in a similar function. The simplest 3-segment-ring is a triangular frame with the cross section of an obtuse (30-120-30 degree) triangle. The obtuse angle points outwards, and two triangle frames interlock along their large, square inner faces (Fig.8a). Every 3-ring interlocks with three neighbors, and their three obtuse-profile legs fill its hole completely (Fig.8b). Now let’s just interlink many of those tiles

and see what happens. On the computer, we track the expanding construction with a simple 3-valent graph, in which each node represents the centroid of a 3-ring, and each edge represents a link between two adjacent tiles. The two links at each end of an edge lie in two planes that are mutually perpendicular. At generation 4 the graph looks as in Figure 8c. At generation 6 it becomes clear that this construction runs into inconsistencies and that any further tiles would partially intersect (Fig.8d).




Figure 8: *Straight 3-rings (a,b) do not tile space, as shown by the conflict in chain-loop-closure (c,d).*

However, if we make the 3-segment tile somewhat twisted (Fig.9a), so that adjacent tiles lie in planes that intersect at an angle of 70.5 degrees (the dihedral angle of the tetrahedron), then in generation 6 we find that tiles that come from different branches of our trivalent tree coincide exactly and close smoothly into chain loops of length 10 (Fig.9c). And this tiling is indeed space filling! Further investigations showed that this tiling can be obtained directly from the triamond lattice: Connect the nearest neighbors of an atom with an (equilateral) triangular “wire frame.” Shrink these triangles slightly to avoid degenerate coincidences, and then form polyhedral approximations to the Voronoi zones around them. The result is the tria-tile (Fig.9a). Note, that in this case space is filled completely with toroidal tiles that all have the same chirality!




Figure 9: *(a) Skewed tria-tile, (b) split version, (c) closed chain of 10 tria-tiles. (d) Sliced 6-ring.*

Ring Slicing

In all four cases of ring-tiles discussed, we can double the number of tiles that interlink with a single toroid by slicing the ring-tiles longitudinally around the toroidal loop (Fig.9d). This is possible, because all these tiles have multiple C_2 symmetry axes that lie in the main plane of the toroid. In this slicing operation, each segment of the ring-tile is split into two narrower wedges that equally split the inner faces that form the central hole. This immediately gives us four new tilings with ring-tiles with 6, 8, 12, and 20 toroids interlinking with an individual toroid belonging to the interspersed complement lattice.

5. Tiles of Higher Genus

Now we aim to construct isohedral tilings using handle-bodies of higher genus. We could start with various interlocking grid structures and break these down into repeated modular elements. The simplest example is based on the cubic lattice. Take a cubic lattice and its interspersed complement lattice and shrink the conceptual wire frames that outline each cubic cell by a small amount so as to avoid degenerate

coincidences between adjacent wire frames; then form the simplest polyhedral Voronoi zones around these wire-frames. The result is shown in Figures 10a,b.




Figure 10: Lattices of interlocking genus-5 tiles: (a) simplest cube frame, and (b) an array built from it; (c) a cube cage assembled from six 4-segment ring-tiles, and (d) a tetrahedral cluster built from it.

Cages as Composites from Rings

There is also an easy “bottom-up” approach that emerges from the ring-tiles discussed in the last section. When playing with these toroids, one notices quickly that several of them assemble into “cage-like” structures and that these “cages” can then readily be repeated throughout 3D space to form a dense tiling. In the case of the cubic lattice and the 4-segment tile (Fig.6b), we can assemble 6 such ring-tiles into the “cage” shown in Figure 10c. These polyhedral shapes are a rough approximation of the Voronoi zones of two sets of interlocking cube-wire-frames in which only cells at grid locations of one parity have been instantiated. Each genus-5 tile interlocks with 4 nearest neighbors in a tetrahedral configuration (Fig.10d).

If we start with the diamond lattice, four of its toroidal tiles (Fig.7a) readily assemble into a cage of genus 3 with tetrahedral symmetry (Fig.11a). Its convex hull is the rhombic dodecahedron. Each such tile touches 12 neighbors along a common face and also interlinks via each of its 6 bent “tetrahedron” edges with 6 nearest neighbors of the complement lattice (Fig.11b).




Figure 11: Lattices of cages; (a) genus-3 cage based on four 6-rings forming a diamond lattice (b); (c) genus-2 cage based on three 10-rings forming a triamond lattice (d). (Some cages are cut in half).

For the triamond lattice, three of its ring-tiles (Fig.7c) assemble into a cage of genus 2 with D_3 symmetry (Fig.11c). Each such tile then interlinks with 8 nearest neighbors (Fig.11d); but the pattern is more involved. The tile interlinks with 6 of its neighbors via single entangled edges; in addition it interlinks with 2 neighbors via mutually interspersed 3-way Y-junctions, so as to maintain 3-fold rotational symmetry for that assembly of two cages.

Others Arrays

The presented examples are only just a basic starting point in the exploration of 3D tiles of higher genus, depicting regular arrays with high symmetry. There are many ways in which we can extend our collection of such tilings. The basic toroidal tiles can be combined in ever larger numbers to form “cages” of higher genus, which then can interlink with their neighbors in a deeper and more involved manner.

Alternatively, we can use the “sliced” ring-tiles and assemble them into various cages. For the 4-segment ring-tile, using the sliced toroids would result in the generic cube frame shown in Figure 10a. For the diamond lattice, we obtain skinny genus-3 cages that touch four neighbors along several faces (those along which the 6-segment rings have been split), and they interlink with ten cages of the complementary lattice – with six via a single entangled handle, and with four via symmetrically entangled Y-junctions. In case of the triamond lattice, the skinny genus-2 cages touch three neighbors of the same net and interlink with 14 neighbors of the complement lattice.

Physical Models

If we plan to build a physical model of such a 3D array from individual tiles, at least one of every two interlocking handles has to be “opened up.” For simple ring-tiles, roughly half of the manufactured tiles must be made of the split variety (Figs.6c, 7c, 9b). In addition to such topological considerations, we also have to keep the actual assembly process in mind. With such warped tiles as the 10-segment ring (Fig.7c) or the tria-tile (Fig.9) insertion of the last half-tile into a dense cluster of tiles may be physically obstructed. Thus it is generally advantageous if the split half-tile parts terminate in straight shanks that are easy to insert into holes. The two half-tiles themselves can be joined together with small wooden pegs.

Finally, there are some manufacturing considerations. My models were mostly fabricated on a fused deposition modelling (FDM-) machine. In this layered manufacturing process, overhang areas need to be supported by some scaffolding material, which is removed after the part is finished and has cooled down. First, we want to minimize the amount of such scaffolding that has to be built. Moreover, we also want to avoid such scaffolding material to clog up the small assembly holes into which pegs are inserted to hold the two half-parts together; this has implications for the orientation in which the part should be built.

6. Interlinked Knot Tiles

In November 1995, Ian Stewart, discussed tiling 3D space with knots [9]. He depicted a rather unsatisfactory, irregular tile, four copies of which filled a cube – which then can tile space trivially. At that time I devised a more symmetrical solution of three congruent interlocking trefoil knots that fill a hexagonal prism (Fig.12a). Of course, it would be preferable to find a tiling in which knots interlink with all nearest neighbors in such a way that the array cannot be taken apart.

In 2005 I presented the connectivity for several loosely linked arrays of knots that “knit together space” [8]. Such arrays can be used as a starting point to find dense tilings with knot-tiles. Conceptually, we have to “inflate” the skeletal knot geometries until they fill 3D space without voids. However, there is a tricky issue that must be addressed. Where strands cross near each other within one and the same knot, they must not be allowed to touch, so that they could “fuse” into a “blob” with attached loopy handles; the tile body should remain strictly of genus 1. Thus enough strands of neighboring knots must pass through the inner regions of each knot, so that all of its loops stay nicely separated. This is rather difficult to figure out for all but the simplest knots. So let’s start with the trefoil knot.

Ideally, every trefoil knot should interlink with at least three other trefoils. Most naturally this leads to a 3-fold symmetric configuration. Now, how can we combine these into a periodic 3D lattice? The triamond lattice (Fig.9) comes to the rescue! Around each atom we find 3-fold symmetry and can thus place trefoil knots that equally interlink with all three neighbors. Now we have to work out the details of the geometry of the interlinked knot strands to prevent strands of equal color from being adjacent to one another. Once again, pipe cleaners and foam-rubber strips served as my “3D haptic CAD tool” (Fig.12b) to find possible ways of winding the strands of several trefoils around each other, so that they prevent crossing strands of any individual trefoil from fusing together. The next step is to use a real CAD tool to generate a set of smooth and symmetrical trefoils, placed at the proper angles required by the triamond lattice and properly interlinking with one another (Fig.12c). This connectivity skeleton must now be inflated to find the Voronoi regions associated with each knot, which then represent the sought-after tiles. Work to find the detailed geometry of these tiles is in progress.




Figure 12: (a) Isohedral tile with the topology of a trefoil knot. (b) Mock-up of interlinked trefoil knots using pipe cleaners. (c) 3D arrangement of the four trefoil knots in the context of the trianet cell.

7. Conclusions

Many new and intriguing tilings have been found. The examples shown are just that – examples! There is a vast domain of possible other tiles in each category, each of which could easily lead to a whole paper by itself. Most of these tiles would not have been found, if I had just tried to deform the basic fundamental domain of certain tilings. It was important to pursue several structured approaches, following several different paradigms – including morphing operations between different given shapes, combining ring tiles into tiles of higher genus, or entangling loose knot structures. I am convinced that there are many additional fruitful paradigms that can be pursued to find further novel and fascinating classes of 3D tiles.

In this paper we have discussed only isohedral tiles. If one steps outside this “conceptual box,” one can experience an explosion of many more possibilities. In one approach, one could start with any monohedral tiling of 3D space and then split those tiles into two or more different parts. Alternatively, one could “squeeze” some kind of “coupling tiles” between pairs of the original isohedral tiles.

Acknowledgements

I am indebted to several of my students: Matthias Goerner, Mark Howison, and Adam Megacz, who have created some of these special tiles or constructed tools to design them. I also would like to thank several colleagues, and in particular John Sullivan, for giving me feedback on various drafts of this manuscript.

References

- [1] J. H. Conway, H. Burgiel, and C. Goodman-Strauss, *Symmetries of Things*. A. K. Peters, Ltd, 2008.
- [2] J. H. Conway, – *private communication*. Math-Art Conf., Tampa 2007.
- [3] C. Goodman-Strauss, – *private communication*. November 8, 2007.
- [4] B. Grünbaum and G. C. Shephard, *Tilings and Patterns*. W. H. Freeman and Co., 1986.
- [5] C. S. Kaplan and D. H. Salesin, *Escherization*. SIGGRAPH 2000, pp 499-510.
- [6] “Sardines” Puzzle → <http://www.passionforpuzzles.com/puzzles/sardines.php>
- [7] D. Schattschneider, *M.C. Escher – Visions of Symmetry*. Freeman and Co., New York, 1990.
- [8] C. H. Séquin, *Tangled Knots*. "Art+Math=X" Intl. Conf., Boulder CO, June 2-5, 2005, pp 161-165.
- [9] I. Stewart, *Mathematical Recreations*. Scientific American, Nov. 1995.
- [10] A. F. Wells, *Three-Dimensional Nets and Polyhedra*. Wiley, New York, 1977.
- [11] J. Yen and C. H. Séquin, *Escher Sphere Construction Kit*. Interactive 3D Graphics Symposium, Research Triangle Park, NC, March 19-21, 2001, pp 95-98.