

My Search for Symmetrical Embeddings of Regular Maps

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

Abstract

Various approaches are discussed for obtaining highly symmetrical and aesthetically pleasing space models of regular maps embedded in surfaces of genus 2 to 5. For many cases, geometrical intuition and preliminary visualization models made from paper strips or plastic pipes are quite competitive with exhaustive computer searches. A couple of particularly challenging problems are presented as detailed case studies. The symmetrical patterns discovered could be further modified to create Escher-like tilings on low-genus handle bodies.

1. Introduction

Regular maps are networks of edges and vertices embedded in closed 2-manifolds of arbitrary genus [4]. The most familiar examples are the five Platonic solids, which represent such maps on surfaces of genus zero. These five polyhedra have high symmetry: Any p -gon face can be moved to any other face and placed there in $2p$ different ways (p possible rotations, plus optional mirroring) and the whole polyhedron will then coincide with itself, i.e., every vertex, edge, and face will fall again on a corresponding vertex, edge, and face. Thus we say that the map of all vertices, edges, and faces exhibits *flag-transitive* symmetry. On the sphere there are two more types of regular maps: the infinite series of *hosohedra* (Fig.1a), a sphere partitioned by q equally spaced longitudinals, resulting in two vertices of valence q at the North and South poles, – as well as their duals, the *dihedra* with q vertices placed regularly around the equator.

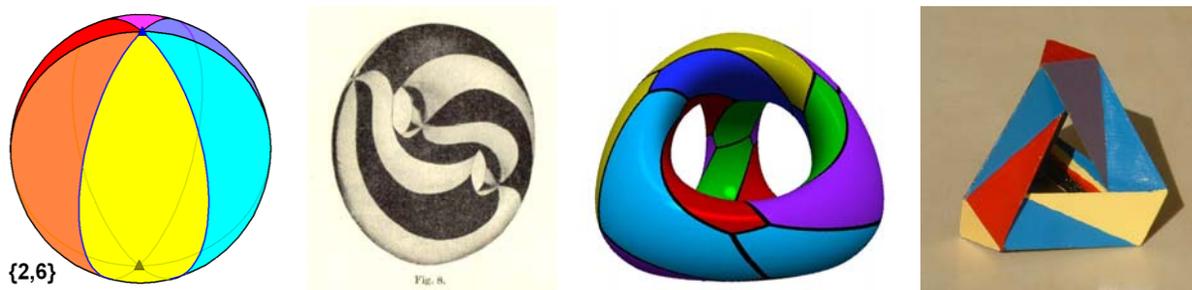


Figure 1: Examples of regular maps: (a) Hosohedron, (b) R2.1, (c) R3.1 (d) R3.1 dual.

Such regular maps also exist on surfaces of higher genus, such as tori or donuts with multiple holes, spheres with multiple handles, or 3-dimensional tubular constructions. Now we assume that the surfaces of these solids are made of a highly elastic fabric, so that portions of the surface that are on the inside of small handles can readily be stretched to fit also around the outside of large bulky loops. If the network drawn on this fabric has the same flag-transitive symmetry described above, so that every *flag* (vertex, edge, face) is topologically identical to every other such flag, then it represents a *regular map*. Note that it is not sufficient that all nodes (vertices) have the same branching valence and all meshes of the network (faces) have the same number of sides. Such *regular graphs* embedded in 2-manifolds can form *equivelar* $\{p,q\}$ tilings, where q p -gons always meet at each vertex. But these tilings are only *locally regular*. They may exhibit quite attractive patterns [15], but regular maps exhibit additional, more *global* symmetries.

For instance, the length of all their *Petrie polygons* is always the same. A Petrie polygon is the closed zig-zag path that is obtained if one moves along the edges of the network and at every vertex one makes alternately the sharpest right-turn and sharpest left-turn possible, thereby hugging every face touched for exactly two consecutive edges. Finding Petrie polygons of different lengths in such a network is the quickest way to show that a given equivelar $\{p,q\}$ tiling is not a *regular map*.

All the different transformations (automorphisms) that bring such a regular map into coincidence with itself form a group, and this group explicitly describes all the topological symmetries inherent in the given regular map. Only recently, due to brute-force computer searches by Conder and Dobcsányi [2], it has become known how many different regular maps can exist on surfaces of various genus. Conder's table of 2006 [3], lists all regular maps on surfaces from genus 2 to 101; it contains 6104 entries (including duals). This list gives the vital data for each map, such as the p and q values for the faces and vertices, the lengths of the Petrie polygons, the number of symmetries, and the corresponding group presentations with their *generators* and *relators* [4]. For each genus g the number of existing maps are consecutively numbered "Rg.#_{p,q}_PL" by increasing p and q , – shown here in curly brackets and followed by the Petrie length. Thus the first regular map of genus 2 we label as R2.1_{3,8}_12. However, Conder's listings give no hint of what possible embeddings of these regular maps might look like.

A serious effort into the visualization of such regular maps is a very recent endeavor. By 2007 when I wrote my paper on locally regular tilings [15], I was aware of only a few explicit models or visualizations. The most famous one is Klein's map (Fig.1c) composed of 24 heptagons (R3.1_{3,7}_8 in Conder's list) which dates back to 1888 [9]. It was celebrated in 1993 with Helaman Ferguson's famous sculpture *The Eightfold Way* [6] located at the Mathematical Science Research Institute in Berkeley (MSRI). A book has been written about this sculpture and related mathematics [12]. Another visual model of a regular map is a depiction of R2.1_{3,8}_1, the Quaternion group, by Burnside [1](p396).

In 2005, inspired by Ferguson's *The Eightfold Way*, I set out to find an embedding for R7.1_{3,7}_18, composed of 72 heptagons, known as the Macbeath surface [10]. This is the next regular map after Klein's map in which Hurwitz's upper limit of $84*(\text{genus}-1)$ symmetries can be reached [7]. However, after struggling with this problem for several months, I finally had to suspend my search and focus on a simpler problem for my 2007 Bridges presentation. In my scaled-back quest, looking primarily for locally regular tilings, I nevertheless found a solution for the regular map R4.2_{4,5}_6 embedded in an octahedron with four tunnels connecting opposite faces [14]. On and off during the next three years I continued my efforts to find some embeddings for regular maps. I started with the simpler cases and looked for patterns that might become useful in my quest to solve the R7.1 problem. In the fall of 2008 I created models for some simple genus-2 and genus-3 problems and also studied the genus-5 cases that would map onto a cube frame. These cuboid models were exhibited at the Art exhibit in Banff [16].

In spring of 2009, Jack van Wijk sent me a draft of his 2009 Siggraph paper [17]. This work changed everything! Suddenly the number of known embeddings of regular maps went from just a handful to more than 50 by the time his final paper was sent to SIGGRAPH. And 15 more solutions were found in the months after that. Van Wijk used an iterative constructive method to generate 3D tubular surfaces of higher genus from the edge patterns of simpler regular maps. Every edge in the map was turned into a tube segment of a new handle-body. He then tessellated its surface into various locally regular tilings in such a way that a reasonably efficient computer search could determine whether some of them actually were regular maps. These maps were then used in turn to generate more complex tubular surfaces of yet higher genus; and the process could be repeated.

Van Wijk's first approach had some limitations. The patterns drawn on each tube segment always had 4-fold symmetry (front-to-back and mid-point symmetry). Thus the only regular maps this algorithm could produce were maps with $4E$ -fold symmetry. This is too limiting in many cases. On the other hand, his program had produced embeddings with a genus as high as 29, and other generated models of lower genus were still of a complexity that could not be handled by a human without the help of a computer search. This paper encouraged me to look for some of the regular embeddings that van Wijk's program

had not yet found, and also to look for more regular, symmetrical embeddings, where the computer solutions looked more twisted than necessary. This program also prompted me to work out a more methodical approach to finding such embeddings by hand, in preparation for the ultimate quest for a nice regular embedding of $R7.1$. This paper describes the different approaches I took to find such solutions and the generally useful patterns that I found, which in turn may help finding more complicated solutions on higher-genus surfaces – by hand or by computer. Thus the focus of this paper is more on the process of finding symmetrical embeddings than on the final results.

2. Methods for Finding Embeddings

It should be noted that there exists a general method for taking a cut-out domain of a given genus g and converting it with a series of cutting and reassembly operations into a special $4g$ -gon. Then there is a general method to fold up this $4g$ -gon into a genus- g handle body (see Section 5). In this way the connectivity pattern of the regular map can be transferred to the surface of a genus- g surface, where the regular map can thus be embedded crossing-free. Thus for “logic-oriented”, “formula-minded” mathematicians there is no mystery left.

Once we have found one solution, we can readily generate other solutions by cutting around any handle or tunnel in the embedding surface and applying a 360° Dehn twist [5], which reconnects all the same edges and faces after the application of a 360° helical twist. Topologically the map embedding has not changed, but the appearance may now be quite different! This process can be repeated arbitrarily many times for closed curves around any handle or tunnel or combinations thereof. The fact that there is in principle such a general way to create arbitrarily many embeddings may be another reason why “pure” mathematicians do not care to “see” any actual embeddings of these maps and are happy with the knowledge that solutions exist. But for more vision-oriented, geometry-minded folks (like the attendees of the Bridges conferences) this is not good enough. We would like to see actual models of such embeddings, and would strongly prefer them to be nice and symmetrical!

My own approach is based on symmetry, intuition, and “numerology.” The technique is very similar to what I used to find the most symmetrical embeddings [13] for the Dyck graph $K444$ on a genus-3 surface (which yields $R3.2_{\{3,8\}_6}$) and for the complete graph $K12$ with twelve vertices on a genus-6 surface (which does not yield a regular map). Since then I have learned to pay more attention to the ordering of the individual edges around each of the vertices. I typically first print a Poincaré disk with the proper $\{p,q\}$ to represent the basic pattern of the tessellation. This is easy; there are many applets on the web to do this, e.g. [8]. From the infinite hyperbolic $\{p,q\}$ tessellation on the Poincaré disk I cut out a finite domain with the right number of faces to represent the desired map. Identifying sets of identical faces on the Poincaré disk for a given map could be quite a tricky task. But fortunately, Conder’s tabulation lists the Petrie lengths for all the maps as well as some select group *relators* that define other wrap-around closures. Using any of these closed paths, one can readily identify points, edges, and faces that represent one and the same instance in the embedding of the regular map. Assigning the same color to repeated instances of a tile then yields a very informative visualization of the complete network and of all its expected wrap-around connections around tubular arms or tunnels.

Next I choose a suitable handle body. This is where the “numerology” comes in. I look at the total number of vertices and tiles and try to find a highly symmetrical surface of the right genus that also has a number of “elements” (tube junctions, connector arms, tunnels...) that have a large GCD with the number of vertices and polygons that need to be placed. To make finding such a match more likely, I have built a comprehensive catalog of usable handle-bodies. Five of these are based on inflated wire frames of the Platonic solids. Others are based on the regular hosohedra, or on other small regular graphs. The availability of suitable handle bodies up to genus 5 is not a problem!

For every vertex of the map I then create a small paper disk that describes the *star* of the vertex, i.e., the circularly ordered sequence of its nearest neighbors. A convenient way to do this is to give unique labels to all the different edges in the cut-out of the regular map. Then each vertex disk is a star of stubs

labeled with these edge labels. If possible, I try to orient these stars in such a way that as many of the stubs point in the directions of their endpoints on other vertices. Now the remaining challenge is to connect all the pairs of stubs carrying the same edge labels, while making sure that all the generated faces of the map are of genus zero and do not contain any tunnel entrances or handle terminations within them.

When the mesh is completed, I trace out several Petrie polygons to check that all of them are of proper length. If not, a different vertex placement needs to be tried. Finally, if everything still looks promising, the set of colors or labels established in the cut-out domain on the Poincaré disk can be transferred to the embedded graph, and it should now nicely wrap around and close on itself as expected.

I have found that working with a physical model is crucial. Just playing with computer displays, even in an interactive graphics environment, is not good enough. It is too hard to trace reliably a highly twisted path to determine its Petrie length, and I do not have any graphics tools at hand that allow easy sketching and annotation on the surface of a higher-genus manifold. The physical models I use may be plastic pipe structures of the right genus and desired symmetry, Styrofoam structures built from heat isolation sleeves for water pipes, paper models constructed from simple paper strips, or more refined net-models that I can assemble into polyhedral space models with a detailed pattern of labeled edges and faces representing the desired map (Fig.2).

3. A Case Study of Medium Complexity

As mentioned above, my main approach is to compare the vertex and face count and the numbers of symmetries of the group of automorphisms of a given regular map with the number of handles, junctions, and types of symmetries of the target handle-body. Then I try to place the vertices in equal numbers near all junctions or at the middle of each handle. For my tutorial exposition I chose an example that did not immediately yield the right solution, but which first misled me to some pretty models that exhibited only local regularity and did not correspond to the map I was trying to embed.

The case I want to discuss is the self-dual map $R5.10_{\{6,6\}_4}$. I started to look at some genus-5 models well before I had found solutions for all maps of genus 2 and 3, because they promised to make particularly attractive models, and because I had already created locally regular tilings on “inflated” cube frames in 2007 [15]. The map $R5.10$ seemed to be a particularly obvious case, since it had 8 vertices, which could naturally be associated with the 8 corners of the “cuborus.” Figure 2a shows my first attempt, which to my surprise however yielded some Petrie polygons of lengths 6 rather than always 4. In a second attempt, I placed half the vertices on the inside of the cuborus junctions (Fig.2b), inspired by the “vertex-flower” solutions I had found on maps of genus 2 and 3 (see Section 4); but this one now had some Petrie lengths of 2. It took quite a bit more experimentation with the three pointy ends of the star hexagons to get them wrapped around the cube-frame handles to obtain the proper Petrie lengths (Fig.2c,d). In this exploration, it helped to consider the self-duality of this map: Ideally, the center of the faces should lie in equivalent locations to the positions of the vertices.

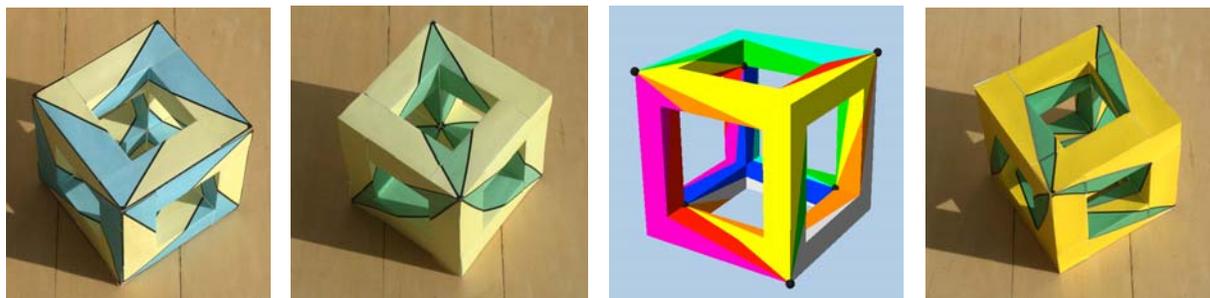


Figure 2: *Cube-frame models with different locally regular tilings of type $\{6,6\}$; only the two models on the right side represent actual embeddings of the regular map $R5.10_{\{6,6\}_4}$.*

4. Useful Recurring Patterns

After I had obtained Jack van Wijk’s paper, I also took an interest in the more exotic maps that had only one or two “countries” and some vertices that were visited more than once by the same “polygon”. Some of his solutions for these cases seemed more twisty than necessary. Focusing on the maps $R2.5_{\{6,6\}_2}$ and $R3.11_{\{8,8\}_2}$, which both are self-dual with exactly 2 polygons and 2 vertices, I came up with very simple and symmetrical solutions on the *Hosorus* structures with 3 and 4 arms, respectively. One vertex is placed on the outside of the handle body at the North pole, the other at the inside of the South pole. The shared boundary of the two intertwined $2^*(g+1)$ -gons then zig-zags back and forth along the arms of the *Hosorus* between the two vertices. The shape of the two polygons reminds me of two interleaved sets of flower petals – hence I call the resulting embeddings *Vertex-flowers*. Figure 3 shows these patterns for genus 1 through genus 5; specifically including Conder’s maps $R2.5$, $R3.11$, $R4.11$, and $R5.15$. Of course, the *Vertex-flower* pattern in this series can readily be applied to embeddings of higher genus.

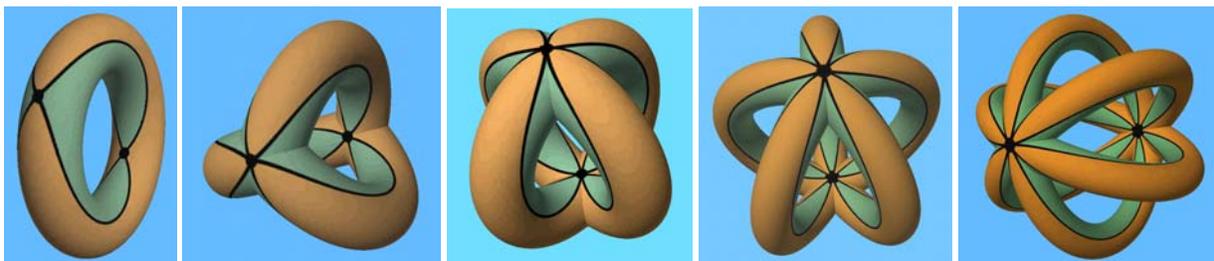


Figure 3: *Vertex-flower* solutions on the *Hosorus* from genus 1 through genus 5, corresponding to regular maps $R1.x_{\{4,4\}_2}$, $R2.5_{\{6,6\}_2}$, $R3.11_{\{8,8\}_2}$, $R4.11_{\{10,10\}_2}$, $R5.15_{\{12,12\}_2}$.

The last map in each genus group, i.e., $R2.6_{\{8,8\}_2}$, $R3.12_{\{12,12\}_2}$, $R4.12_{\{16,16\}_2}$, and $R5.16_{\{20,20\}_2}$, etc, was more challenging. The classical technique described in textbooks [11] explaining how to fold up a single $4g$ -gon into a genus- g surface (see Section 5) does not yield regular maps, since the individual edges then do not arrive at the single vertex in the right sequence. Each individual edge must approach the vertex from opposite directions. I started by drawing a suitably colored vertex star onto a handle body with a single junction and then traced non-crossing paths for the differently colored edges around the available handles. I could not find any symmetrical solution for the genus-2 case, but almost immediately found a nice symmetrical pattern with 3-fold symmetry for the genus-3 case (Fig.4a,b). From that I could eventually develop a general approach to handle a map of this type for any genus (Fig.4c). The basic idea is to provide g loops that do not intersect and which approach the junction where the single vertex is placed from opposite directions. On each of these handles, two edges use the handle to twist around each other and exchange their positions. Thus $4g$ edges join together at the vertex.

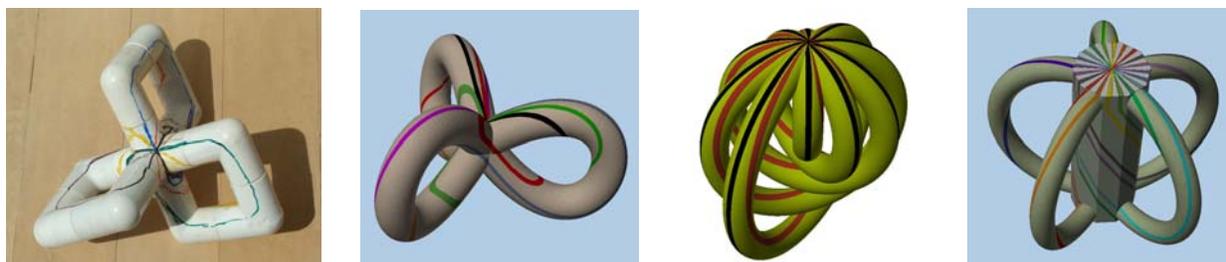


Figure 4: *Single-vertex single-polygon* models corresponding to the regular maps $Rg.last_{\{4g,4g\}_2}$: (a,b) the case of $R3.12_{\{12,12\}_2}$; (c) general solution for any genus; (d) map $R5.16_{\{20,20\}_2}$.

A different, more symmetrical way to embed this type of map is shown in Figure 4d. It uses the *stalk* structure. Attached to a central *stem* are g semicircular handles that carry two edges each and allow them to switch positions. The return path of each such pair of edges to the single vertex at the top of the stalk

occurs in a helical pattern around the stem, which transports each pair to the side opposite its handle. For odd genus, a nice pattern of g -fold rotational symmetry results; the case for genus 5 is shown in Figure 4d. For maps of even genus, this symmetry has to be broken with some extra twist in the helical stem. This is the underlying reason that I could not find a nice symmetrical model for the genus-2 case.

For each genus, there is also a map that has a single vertex, but has two polygons sharing it. The simplest case is $R2.4_{\{5,10\}_2}$, and it is one of the embeddings that van Wijk's program could not find; thus it presented an interesting challenge. Studying the basic $\{5,10\}$ tiling pattern in the hyperbolic Poincaré disk revealed that once again the individual edges had to approach the single vertex from opposite directions. But now there are two additional edge-endings using this vertex. It turns out that this additional edge is actually the one that splits the surface into two separate domains.

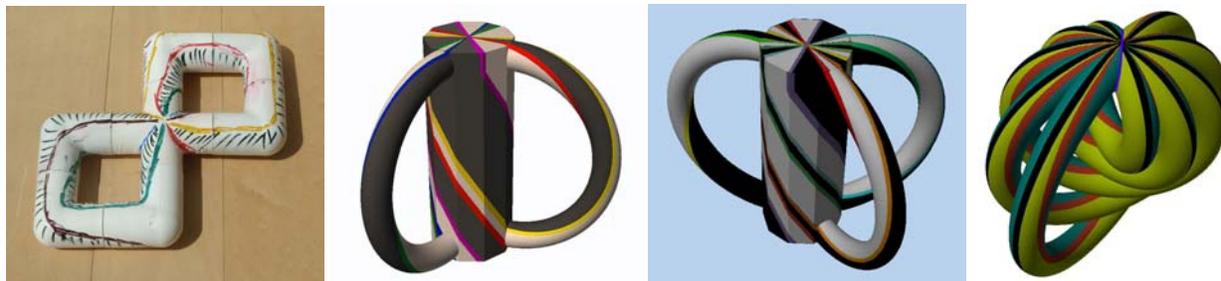


Figure 5: One-vertex two-polygon models corresponding to: (a,b) regular maps $R2.4_{\{5,10\}_2}$ and to: (c) $R3.9_{\{7,14\}_2}$; (d) shows a way to embed the general case, with $R5.14_{\{11,22\}_2}$ as an example .

For the map $R2.4$ every one of the five differently colored edges approaches the vertex twice from opposite directions. I painted this 5-color crossing with marker pens onto plastic pipes and then I tried to connect these colored edges without any crossings. I quickly found 3 different solutions, but I could not find any nice symmetrical pattern (Fig.5a). Later, after fully understanding the solutions to the 1-vertex 1-face maps discussed above, I saw how this new problem could be derived from the previous one. The same general structure shown in Figure 4c can be used; but now I bunch the handles in a way that makes room for an additional (purple) edge that runs around the top junction and which splits the single $4g$ -gon into two $(2g+1)$ -gons (Fig.5d). Again, a space model based on the *stalk* geometry can be used. But now the symmetry is broken for maps of even as well as odd genus. Examples for the maps of genus 2 and 3 of this type are shown in Figures 5b and 5c. This section demonstrates how previously established patterns can help in the search for solutions to other problems.

5. Some Difficult Maps of Genus 3

Some regular maps just don't want to yield to such methods. Examples are the maps $R3.3_{\{3,12\}_8}$ and $R3.5_{\{4,8\}_8}$, and I will use them to show the difficulties one may encounter in trying to find embeddings of regular maps. These are examples, where van Wijk's program did not find solutions.

Failed Intuitive Attempts

The Poincaré patterns of these two maps show obvious 3- and 4-fold symmetries. There is also a strong local 4-fold symmetry around every vertex: Every third edge connects to the same vertex, and there are 3 such sets interleaved. This local symmetry most strongly suggests the use of a 4-arm Hosorus surface (Fig.3c). The four vertices can then be placed on the insides and outsides of the two junctions of this structure. It then seems natural to connect the inner and outer vertices at each pole to one another by 4 edges that run in the valleys between adjacent pairs of Hosorus arms. But after placing these 8 edges, I was never able to accommodate all the other edges of the $R3.3$ or $R3.5$ maps crossing-free.

I then tried placing each of the four vertices into the middle of one of the four handles. Particularly for the case of the $R3.5$ this seems promising, because every vertex is connected to only two other

vertices, and those could thus be placed on the two adjacent handles. One edge to each neighbor vertex could then travel above the tunnel in between, and the other one below it. To keep the arrival at the neighboring vertex in the proper cyclic order, two of these edges may run entirely on the outside of the Hosorus structure, while the other two edges travel around the inside of one of the Hosorus arms to implement the required topological crossing of two edges. Yet, again, I failed to complete the edge pattern for the R3.5 map.

Procedural Approaches

Perhaps I was just not diligent enough to pursue any of the above approaches to its final successful conclusion. After all, we know that a solution must exist. Also, because of the purely topological nature of this problem, it should not matter where exactly the vertices are placed – in principle they can be moved smoothly around the whole surface while dragging the whole network along without creating any new edge crossings; of course, this may break the desired symmetry

As stated in Section 2, there is a guaranteed procedure to find an embedding of any of these maps: We start by identifying a connected region in the Poincaré disk that contains every facet of the desired map exactly once (Fig.6a). On the boundary of this cut-out domain we identify corresponding edges through which the map closes back onto itself (around one of the many arms or holes of the chosen embedding surface). We link together sequences of edges that connect to other identical sequences somewhere else on the boundary into “Boundary Edge Trains” (BET) of maximal lengths. For a map that embeds into a genus g surface, there will be n such BETs, where n is at least equal to $4g$, but is likely to be bigger. Topologically the map cut-out is thus an n -sided disk with the n BETs as its boundary. This 2-manifold patch needs now to be folded up into a genus- g surface.

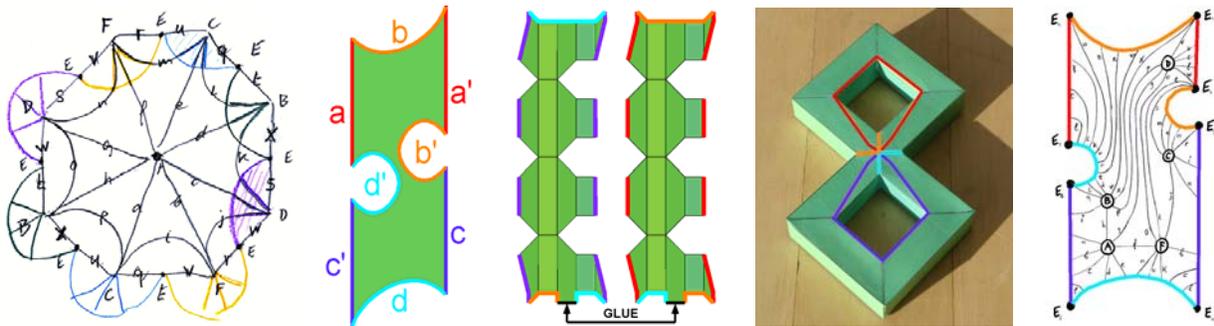


Figure 6: General map embedding procedure: (a) cut-out of map, (b) $4g$ -gon, (c) foldable net for $g=2$, (d) folded-up handle-body of genus 2, (e) resulting edge pattern for the map $R2.1_{\{3,8\}_{12}}$.

To achieve this in a generally applicable manner, the n -gon needs to be transformed into a $4g$ -gon with a very special sequence of edges. The edges must occur in g groups of four oriented edges, each with the pattern $\{a, b, a', b'\}$ (Fig.6b). Here identical characters mean that these “edges” will be identical BETs that need to be fused, and the “prime” attached to a character means that this BET has the opposite directionality along the boundary from the un-primed character. Each such set of four BETs in the proper order will then allow us to form a toroidal handle. The two BETs labeled with “a” join to form the longitudinal seam along the whole handle; the “b-BET” forms the end of the toroidal hose, and “b’-BET” forms an opening onto which the end of the hose can be fused, forming a closed, edge-free toroidal handle. Topology books describe how the original n -gon can be transformed into the properly sequenced $4g$ -gon [11]. The two major steps are: Re-form the fundamental domain to have all its boundary vertices identified to a single vertex \mathbf{V} of the given map, thus making all BETs into closed loops starting from and ending at that one vertex. This can be achieved by repeatedly cutting off a polygonal region starting at an instance of \mathbf{V} and ending at an adjacent vertex \mathbf{W} that we would like to remove from the boundary of the current polygon, and then gluing the cut-off piece back onto the remaining polygon so that the number of distinct instances of \mathbf{W} is reduced by one. Once all original vertices except \mathbf{V} have been removed from

the boundary, a similar kind of re-assembly step can bring 4 BETs that have the relationship (a, b, a', b') into adjacent positions. This step is repeated until the whole boundary consists of g sequences of BETs with the desired ordering [11](Ch.7). Figures 6c,d show an explicit fold-up net for a genus-2 handle-body.

While doing all this, we carry along the whole network of edges and faces inside the transforming boundary of the cut-out. After all the above operations, the inside of our $4g$ -gon still has exactly one copy of the whole map that we want to embed in a genus- g surface, and the connectivity of the edges has not changed. I have manually performed this procedure for the case of the map R2.1_{3,8}_12, but the results (Fig.6e) is not helpful at all to find a nice regular space model; the surface is crowded with edges in some regions and very sparsely populated in others.

Breakthrough

I finally realized that the two problems R3.3 and R3.5 are just different variants of the same problem, and that my assumption that R3.5 was perhaps easier to solve, because it had only 8 edges emanating from each vertex, was misleading. The (invisible) diagonals in each quadrilateral facet represent just as strict topological constraints as the additional edges in R3.3. The two problems can readily be transformed into one another by just turning on and off a set of $1/3$ of all edges that act as diagonals in the quadrilateral faces of R3.5. From this moment, I just concentrated on R3.3 alone, knowing that R3.5 would fall out for free once I had solved the first problem.

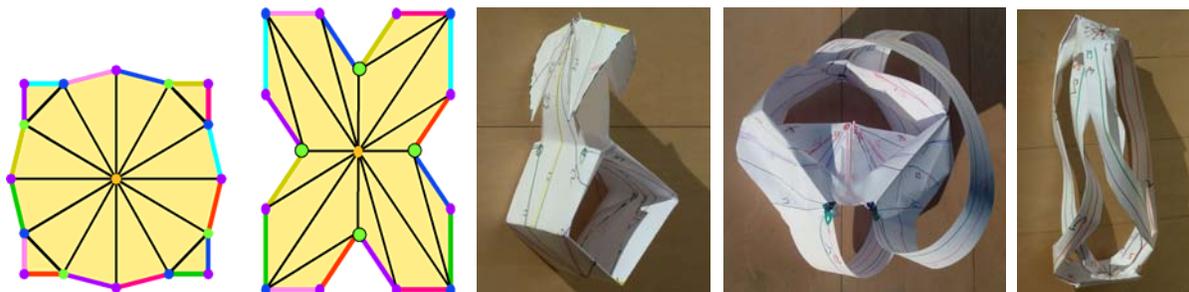


Figure 7: Folding up map R3.3: (a) cut-out; (b) deformed cut-out; (c) torus-fold; (d) adding missing connections via two paper strips; (e) redrawing all connections on a 4-arm hosorus.

The key breakthrough came when I found a way to see how the cut-out domain could be folded up into a genus-3 surface in a more direct way, that did not require all the cuts, reassemblies, and deformations described above for the general fold-up procedure. I took the cut-out consisting of 16 triangles that most naturally describes the desired map in the Poincaré disk (Fig.7a) and deformed its shape in such a way that edges that need to be joined all appear parallel to one another. This first led to an 8-pointed star, which I then deformed into a rectangular domain (Fig.7b) that could readily be folded into a toroidal configuration with two openings bounded each by 4 BETs (Fig.7c). Forcing two more pairs of identical vertices (large, green) to fall upon one another contracted these openings into two holes each, bounded by only two BETs. At this point my model corresponded to a toroid with four openings with 2-edge borders. These holes had to be connected pair-wise with two additional tubes glued upon the base torus to make it genus 3. In Figure 7d these two new handles are represented simply by narrow paper strips carrying the relevant edges to complete all connectivity required in the given R.3.3 map. In particular, the remaining two green vertices had to be moved along these handles to the other end and joined with the other already merged pair of green vertices; this stretched several more edges connected to these vertices along the two paper strips. While the resulting paper model is rather ugly and confusing, it contained the most important information: What edges run along what handle, and in what sequence are they placed around each arm? This then determines the embedding of the whole map.

With this information I could form a new paper model that consisted primarily of four tubular arms with the right edges embedded in them and two hubs that carried two vertices each and showed the circular sequence of edges emerging from each of them (Fig.7e). I lined up the tubular ends with the

proper edges near the vertices to which they had to connect and then massaged the whole structure to see what natural configurations might emerge. The two most promising candidates were a tetrahedral frame with two (opposite) short thick arms carrying two of the vertices each, or the 4-arm Hosorus structure with the 4 vertices placed on the inside and outside of the two 4-way junctions. I pursued the latter approach to see how it differed from the solutions that I had failed to find for this starting configuration.

Model Refinement

My next model consisted of a genus-3 surface made of 1" propylene PPR pipes, on which I drew and labeled all edges the way I was reading them off my crumpled paper model, carefully maintaining the topological twists and turns of every edge. And, wondrously, I was now able to accommodate all 24 edges of $R3.3_{\{3,12\}_8}$ in a crossing-free embedding and thus had found the desired solution!

To obtain a nice, colored, reproducible model, I designed a cut-out net for making paper models that I could hand to other people (Fig.8a). For ease of design and assembly, I divided the model into its horizontal surfaces and a prismatic mantle of its vertical surfaces. On these surfaces I drew the various edges and labeled them, and optionally colored the facets between them according to the color code shown on the Poincaré disk. I displaced the edges laterally to keep them nicely separated (pinning them to some near-by corner, if necessary), yet I did not let them deviate too much from geodesic lines.

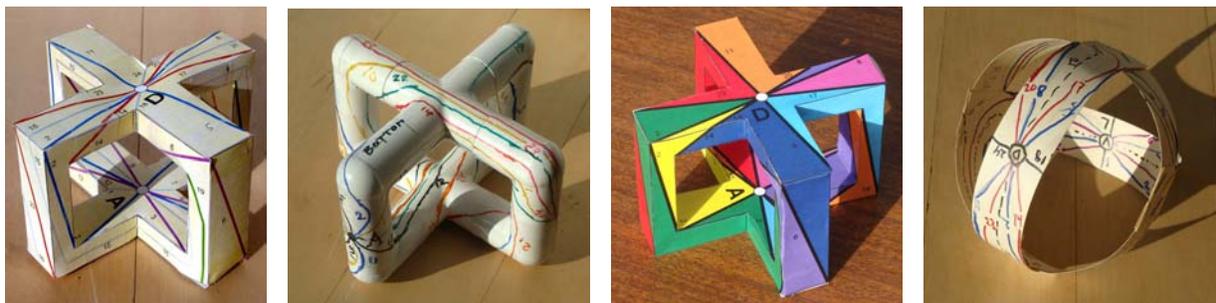


Figure 8: Steps towards a refined model: (a) paper-frame model of $R3.3$; (b) edge pattern resulting from different vertex positions; (c) paper-frame model of $R3.5$; (d) new paper-strip model of $R3.3$.

It is interesting to study how this model deviates from the initial partial solutions that I had obtained on the Hosorus with exactly the same vertex placement. Most importantly, there is no 4-fold symmetry, just $D2$ symmetry! In particular, the inner and outer vertices at each pole are not directly connected by four short edges running in the valleys between them – two of the edges take long detours around two other arms of the Hosorus. I also wanted to see what the solution might look like if I placed the vertices in the middle of each Hosorus arm. For that purpose, I used my “canvas” constructed from plastic pipes and moved the four vertices to their new locations, dragging along all attached edges. This destroyed all remaining symmetry (Fig.8b)! Starting from the solution shown in Figure 8a, by just omitting one third of all the edges and joining the two triangular faces on either side of it into quadrilaterals, I readily obtained an embedding of $R3.5$ (Fig.8c). The resulting pattern differs somewhat, depending on which set of edges is omitted. Finally, Figure 8d shows a paper-strip model of an embedding of $R3.3$ on a handle body that does not suggest more symmetry than can actually be achieved.

6. Discussion and Conclusions

I am still surprised and puzzled by the varying degrees of difficulty associated with finding nice embeddings for different regular maps. While some symmetry groups yield a nice solution almost instantly, others stubbornly refuse to yield a good solution, and often allow only solutions with lesser symmetry than one would expect initially. In addition to my adhoc, intuition-based approach, I have also started to look into some methods based on computer searches. In the fall of 2009, I teamed up with 3 undergraduate students, and we have started to explore some geometry and graphics algorithms involving

extensive search trees that could help to solve such embedding problems. But the programs have not yet matured enough to yield interesting new results. For readers with a lesser familiarity with the fields of mathematics and/or computer science this paper may be a surprising eye-opener concerning the way that some mathematical problems get solved. This often happens not by writing long equations on a blackboard or by developing complicated computer programs! Instead it may involve folding paper, bending wire, cutting styrofoam, or drawing patterns on tennis balls or on plastic tori from a toddler's toy set... The combination of these many activities may then lead to new insights and the right mental state in which one can suddenly "see" a new solution for the problem at hand.

My main goal was to find simple and symmetrical embeddings of regular maps. With some judicious coloring of the various tiles, the resulting models should already have some aesthetic merits. Their artistic properties may be further enhanced by deforming the tile boundaries into Escher-like patterns. A new difficulty that arises in the context of a "3D-canvas" is the fact that the tessellation can never be seen as a whole; some parts are always hidden for any viewing direction. Possible ways to deal with this issue are the use of transparency or structures that only show the edges of the maps, e.g. as ribbons or tubes. Other approaches use the symmetry of these objects to form repetitive grids, in which different instances of the basic cut-out net can be seen from different perspectives. Results will be shown at the conference in Pécs.

Acknowledgements

I am grateful for the stimulating e-mail correspondence with Jack van Wijk and the many educational tutorials I received from some of my mathematically gifted friends and colleagues. This work was supported in part by the Center for Hybrid and Embedded Software Systems (CHESS) at UC Berkeley, which receives support from the National Science Foundation (NSF award #CCR-0225610 (ITR)).

References

- [1] W. Burnside, *Theory of groups of finite order*. Cambridge University Press 1911.
- [2] M. Conder, and P. Dobcsányi, *Determination of all regular maps of small genus*. J. of Combinatorial Theory, Series B, 81 (2001), 224–242.
- [3] M. Conder, *Orientable regular maps of genus 2 to 101*. <http://www.math.auckland.ac.nz/~conder> (2006).
- [4] H.S.M. Coxeter and W.O.J. Moser, *Generators and Relations for Discrete Groups*. Springer, 1980.
- [5] Dehn_twist: http://en.wikipedia.org/wiki/Dehn_twist
- [6] H. and C. Ferguson, *Eightfold Way: The Sculpture*. In: Levy [12], pp 133–173 (1999).
- [7] A. Hurwitz, *Ueber algebraische Gebilde mit eindeutigen Transformationen unter sich*. Math. Ann. Vol 41 (1893) pp 403–442.
- [8] D. E. Joyce, *Hyperbolic Tessellations*. <http://aleph0.clarku.edu/~djoyce/poincare/> (2002)
- [9] F. Klein, *Ueber die Transformationen siebenter Ordnung der elliptischen Funktionen*. Math. Ann. Vol 14, (1879). (Translation into English by S. Levy [12]).
- [10] A. M. Macbeath, *On a Curve of Genus 7*. Proc. London Math. Society 15 (1965) pp 527–542.
- [11] W. S. Massey, *Algebraic Topology*, Springer, 1991.
- [12] S. Levy, *The Eightfold Way: The Beauty of Klein's Quartic Curve*. Cambridge Univ. Press, 1999.
- [13] C. H. Séquin and L. Xiao, *K12 and the Genus-6 Tiffany Lamp*. ISAMA, Chicago, June. 17-19, 2004.
- [14] C. H. Séquin, *Patterns on the Klein Quartic*. Proc. Bridges 2006, London, pp 245–254.
- [15] C. H. Séquin, *Symmetric Embedding of Locally Regular Hyperbolic Tilings*. Proc. Bridges 2007, Donostia, pp 379–388.
- [16] C. H. Séquin, *Regular Maps on Cube Frames*. Bridges 2009, Banff, Art Show Proceedings.
- [17] J. J. van Wijk, *Symmetric Tiling of Closed Surfaces: Visualization of Regular Maps*. Proc. Siggraph 2009 Conf., New Orleans, pp 49:1-12.