

Symmetrical Embeddings of Regular Maps R5.13 and R5.6

Inspired Guesses followed by Tangible Visualizations

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

Abstract

This report is a documentation of my trial-and-error design process to find a symmetrical embedding of the regular map R5.13 on a genus-5 2-manifold. It documents the non-linear way in which my mind homed-in on a valid solution and then refined that solution to obtain a satisfactory geometrical model. This *design-thinking log* may serve as a case study for a design approach that switches back and forth between doodling with physical materials, computer-aided template and model construction, and verification of the results on tangible visualization models. Lessons learned on R5.13 were subsequently applied to solve the embedding of the regular map R5.6.

1. Introduction

Regular maps are networks of vertices, edges, and faces of very high symmetry. All faces are p -gons and all vertices are of the same valence, q . In addition, there is global symmetry, so that any element of the network (vertex, edge, or face) can be moved to any other element of the same kind, with any orientation chosen, and the whole network will then be able to follow suit and come to coincidence with itself. The most familiar examples are the five Platonic solids, which represent such regular maps on surfaces of genus zero. A rigid-body motion can move any p -gon face to any other face and place it there in $2p$ different ways (p possible rotations, optional mirroring) and the whole polyhedron will then coincide with itself.

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 kind of global symmetry described above, so that every *flag* (a combination of one vertex and an associated edge and face) is topologically identical to every other such flag, then it represents a regular map.

These regular maps can most compactly be described by the group of all their symmetries. First we specify the infinite tessellation $\{p,q\}$ of which they are a subset. Next we have to define how large these networks are and how they are wrapping around and are closing back onto themselves. This is done with a small set of *relators*, which are sequence of symmetry operations that result in the identity transformation. One way to characterize such a sequence of transformations is by defining “**R**” to be a counter-clockwise rotation around the center of a face, and “**r**” or “**R**⁻¹” as its inverse. A second type of transformation, “**S**” defines a counter-clockwise rotation around a vertex, and “**s**” or “**S**⁻¹” its inverse. Thus a relator is then a string of these characters that would bring any flag in this regular map back to its starting point.

The search for regular maps is a relatively recent endeavor. In their description of the field in 1984 Coxeter and Moser [4] only mention a few tens of such maps. But in 2001 Conder and Dobcsányi [2] started to use the computer for a group-theoretical search for all regular maps on orientable surfaces of genus 2 to 15; and in 2006 Conder published a listing of all regular maps on surfaces from genus 2 to 101;

this list contains 6104 entries (also counting the duals) [3]. However, this list only states what regular maps exist and gives the vital data for each, such as the p and q values for the faces and vertices, the total number of symmetries, and the corresponding group presentations with their generators and relators [4]. These publications give no hint of what any possible embeddings of these regular maps might look like.

A serious effort into the visualization of such regular maps is an even more recent endeavor. By 2007, when I wrote a paper on locally regular tilings [14], I was aware of only a few explicit models or visualizations. The most famous one is Klein's map composed of 24 heptagons (R3.1 in Conder's list), which dates back to 1888 [8]. It was celebrated in 1993 with Helaman Ferguson's famous sculpture "The Eightfold Way" [5] located at the Mathematical Science Research Institute in Berkeley (MSRI). A book has been written about this sculpture and related mathematics [9]. Other visual models of regular maps comprise depictions of the Quaternion group of order 8 by Burnside [1] and my own model of an embedding of the complete tripartite graph $K_{4,4,4}$ in a surface of genus 3 [12]. In all cases, the people who constructed these models did not just want *any* crossing-free embedding, but they were looking for an embedding that would preserve as much geometrical symmetry as possible. Thus the model of Klein's map R3.1 and my own model of R3.2 exhibit the symmetry of an oriented tetrahedron, while Burnside's model of R2.1 has D_2 symmetry.

There is a general textbook procedure that will produce a legal embedding [11]. But it typically results in highly irregular patterns with a much higher density of vertices and edges in some parts of the embedding 2-manifold than in other parts. The procedure starts with a *fundamental net* that covers the regular map exactly once when different points on its perimeter that carry the same labels are sewn together. It then goes through a sequence of "cut and glue" operations where parts of the net are cut away and then reattached at some other part of the perimeter where there exist matching vertices and edges. In this manner all vertices, except for multiple instances of one single vertex, are moved to the inside of the modified fundamental net. In a second phase, involving more "cut and glue" operations, partial edge sequences along the perimeter of the net are brought into a special configuration. Pairs of edges will always appear in sequences of the form a, b, a^{-1}, b^{-1} , where a 2-edge patterns a, b is followed by the sequence of the inverses of those two edges. After proper application of the "cut and glue" operations, a regular map of genus g will exhibit g groups of such 4-edge sequences. Each such group can then be closed into a tube by joining a with a^{-1} , and this tube can then be formed into a loopy handle by gluing b onto b^{-1} . Thus a handle-body of genus g will result in the shape of g loops emerging from a single junction of valence $2g$. However, the originally regular net based on a $\{p,q\}$ tessellation will now appear in a highly distorted and unevenly distributed manner on this handle-body.

In 2005, inspired by Ferguson's *The Eightfold Way* [5], I set out to find a symmetrical embedding for R7.1, composed of 72 heptagons joined in valence-3 vertices, known as the Macbeath surface [10]. This is the next more complex regular map after Klein's map in which Hurwitz's upper limit of $84 \cdot (\text{genus}-1)$ symmetries can be reached [6]. On and off during the next three years I continued my efforts to find some symmetrical 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 symmetrically onto a cube frame. These cuboid models were exhibited at the Art exhibit in Banff [15]. It was not until January 2010 that I finally did find a satisfactory solution for the embedding of R7.1!

In spring of 2009, J. J. van Wijk sent me a draft of his 2009 Siggraph paper [17] describing a computer program designed to find symmetrical embeddings. It did not attempt to find a good solution for any particular map. Instead it used an iterative, constructive method to generate 3D tubular handle-bodies of higher genus derived from the edge patterns of simpler regular maps. Every edge in such a map was turned into a tube segment of a new handle-body. Each of these tube segments was then tessellated using some locally regular tiling in such a way that a computer search could determine whether the combined result actually was a regular map. Successful map embeddings 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 increased the number of known embeddings of regular maps 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. Some of these maps had a genus as high as 29. 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 was very inspirational and stimulating. It encouraged me to search for some of the embeddings that van Wijk's program had not yet found, and also to look for more 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 $R7.1$. In my 2010 paper [16], I describe the different approaches I took to find such solutions and the generally useful patterns that I found, which may help to construct more complicated solutions on higher-genus surfaces – by hand or by computer.

The summers of 2009 and 2010 were quite productive. By Spring 2010 I had found good solutions for all remaining cases of genus 2, 3 and 4; and at the time of this writing, only two maps of genus 5 are still awaiting the discovery of a nice embedding. A couple of weeks ago I tackled the two maps $R5.12$ and $R5.13$. These are particularly intriguing since their descriptions look almost identical. For $R5.12$ van Wijk's program had already found a pleasing simple-to-understand space model; but it did not exhibit maximal possible symmetry. It was easy to modify this model to enhance its symmetry from C_2 to D_2 , which, I believe, is the maximum that can be achieved.

$R5.13$ was much more challenging. On about the fourth day of my struggle to find a solution, I realized that this was a rather difficult problem, which might keep me occupied for quite a few more days. Even though I felt that I had had several good ideas and gained some valuable insights, I was not feeling any closer to a solution than I was when I started a few days earlier. That is when I decided to keep a log of my activities and of my reasoning. I reconstructed the log for the first three days to the best of my recollection and began to make entries once or twice a day as I found time to work on this problem, and whenever I felt that I had taken another relevant step. In this paper I am showing most of the original paper models and virtual renderings, but only about half of all the paper sketches and doodles I made.

This write-up is intended to serve as a case-study of the rather non-linear path by which solutions to difficult puzzles or design tasks come about. It caters to the belief that reflecting about the design-process will improve one's design-thinking. Already in the process of compiling these notes and trying to explain my reasoning that accompanied the creation of the various models, I have gained new insights!

2. Log of my Thinking and Model-Building Activities

Saturday 9/4/10 (am):

On this day I finally understood the connectivity of the regular map $R5.13$ and its difference to $R5.12$! This was a task that had kept me preoccupied for several days. It required careful application of more than just one of Conder's relators, e.g., "SRRRSr" to establish the proper labeling of all the edges of the network of the map displayed on the Poincaré disk – which I obtained from a very handy on-line applet [7]. Once I had it all sorted out, the two Poincaré disks look almost the same (Fig. 1a, b), except that at every vertex the labels of two opposite edges have been exchanged. As will become clear below, this has far-reaching consequences for the connectivity of the various octagonal facets!


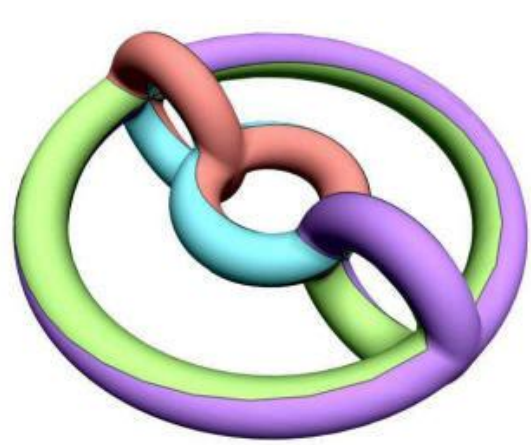


Figure 1: Representation of two regular maps on the Poincaré disk: (a) R5.12 and (b) R5.13.

I already had some nice embedding of R5.12 based on intersecting tubular rings – one from the paper of Jack vanWijk [17] (Fig.2a) and one based on my own modeling efforts in August 2010 (Fig.2b). Thus it was plausible to start with this tubular arrangement and to try to twist and bend the coloring on the tubes to accommodate R5.13. But soon it became clear that no simple twisty edge patterns on these tubular handles could solve the problem.



R5.12 {8, 8}: 4 octagons, 16 edges, 4 vertices




Figure 2: Regular embeddings of R5.12: (a) C_2 -symmetrical solution by Jack vanWijk [17]; (b) my own D_2 -symmetrical solution.

Saturday 9/4/10 (pm):

I was aiming for D_2 symmetry; thus a good structure for the handle-body is a ring (or toroidal ribbon) with 4 holes in it. I sketched such a structure and placed the vertices and their neighborhoods (cut out from the map on the Poincaré disk) at the junctions between the holes (Fig.3a). Then I tried to link up all the edges. It was tedious and confusing work and did not lead to a successful result.




Figure 3: Genus-5 handle-bodies with D_2 symmetry: (a) toroidal ring with 4 holes; (b) a paper-strip torus with four slanted handles attached.

I also tried to connect all matching edge-stubs by starting with a toroidal ribbon to which I had added four slanted handles, each going about half way around the torus (Fig.3b). I was not able to make all the connections either.

Sunday 9/5/10:

Since I could not find a solution in a top-down manner, based on a promising handle-body with four plausible vertex positions, I started to investigate how the vertices of a single facet connect back to themselves. I hoped to gain some insight as to what handle-body shapes might naturally emerge on which such a facet could be embedded crossing-free. I formed an 8-armed paper “spider” and joined opposite arms into loops. An easy-to-grasp structure resulted; it consisted of four dangling loops, all passing through one common junction (Fig.4a). This network could be embedded nicely in a genus-2 torus realized as two nested rings of different sizes, fused together in one junction (Fig.4b). I later made another model with a more customized octagon (Fig.4c). When its opposite vertices were forced to join, a surface resulted that was clearly forming a 2-ring toroidal structure (Fig.4d).

{Comment with hindsight: This structure should have been clear to me from the beginning, since it is the same topology for an individual octagon as in the map R5.12 shown in Figure 2.}




Figure 4: Connectivity of a single octagonal facet: (a) as an 8-arm spider; (b) a genus-2 surface in which it could be embedded; (c) a customized octagon, and (d) its curled-up state.

However it was not clear how one could possibly weave four of those structures together seamlessly, observing the edge-labeling of R5.13.

Monday 9/6/10:

I finally understood a key difference between the two maps! In R5.12 pairs of connected polygons have their shared edges appear in reversed sequences; so that two facets can readily be brought together back to back – forming a two-colored inflated cross-shaped cushion (Fig.5a). But in R5.13 pairs of polygons have their shared edges appear all in the same sequence. Thus they cannot be put together back to back. Stubs with corresponding edges on them need to go through some twisted contortions so that they can be joined up with proper orientation. If the face normals at the centers of the two octagons point in the same direction, we obtain the configuration shown in Figure 5b; but it is difficult to see how this embeds in a toroidal structure. If the two face normals point in opposite directions, e.g., outward (Fig.5c), a configuration results that would fit readily on a torus – except for the fact that vertices at opposite positions on this torus carry the same labels, and thus need to be merged!




Figure 5: Two octagons joined by their shared edges: (a) the mirror-symmetric solution for R5.12; (b) – (d) some twisted solutions for R5.13.

Tuesday 9/7/10 (am):

Since I was trying to embed R5.13 in a toroidal structure with either four holes or four handles, I wanted to see what R5.12 would look like in this form. I drew a projected view of Figure 2, stretching and twisting the edges at the two types of tubular junctions so that they would fit on a flat template. I tried to maintain D_2 symmetry while also keeping the front and back of the templates the same (Fig.6a). These four paper cut-outs could then be joined into a toroidal paper strip with 4 holes, preserving D_2 symmetry.




Figure 6: Embedding of R5.12: (a) D_2 -symmetric template; (b) resulting toroidal structure.

Tuesday 9/7/10 (pm):

The model shown in Figure 6b then served as an inspiration while I was looking for embeddings for R5.13. I tried to use a similar structure with overall D_2 symmetry. However, I placed pairs of vertices

symmetrically between two holes at the front and back of the template (Fig.7). It was then easy to place a set of four edges that connect such a pair of vertices on opposite sides of the toroidal strip. In order to have the two sets of four edges appear in the same order around both of the vertices, pairs of edges need to reverse order as they go from one side of the strip to the other. This can be achieved by routing one edge through the edge of the hole closest to the vertex pair, and the other edge through a point opposite on the perimeter of the same hole. In this manner I could quickly establish the proper sequences of eight colors around all vertices and join up half of all edges locally.

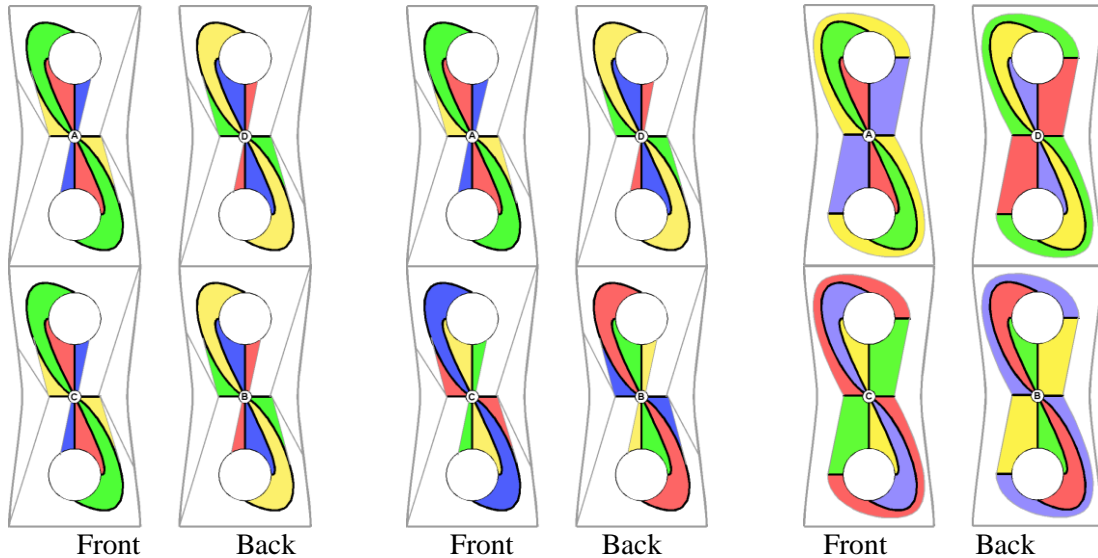


Figure 7: *Partial embeddings of R5.13: (a)–(c) different legal color sequences.*

The big problem is now to connect these two 2-hole clusters to one another!

Wednesday 9/8/10 (am):

I tried about four different combinations of colorings – equivalent of cyclical shifts of the edge sequences around the two vertices. I was never able to connect all remaining edges.




Figure 8: *More partial embeddings of R5.13: (a)–(d) trying to find proper global connectivity.*

{On 9/10/2010, after considering twisting the texture representing the fundamental net on the torus, it occurred to me that I might also try to twist the whole toroidal band. – This may be something to follow up on.}

Wednesday 9/8/10 (pm):

Each hole-pair cluster can be seen as an entity with eight connectors of four different colors sticking out in different directions. How can these dangling arms from two such clusters be connected in the right sequence? I made a conceptual model of this arrangement using paper strips of four different colors and tried to connect the loose arms in the right order (Fig.9a). Since the two clusters together already need four holes to accommodate their local connectivity, we only have available one more handle to solve our puzzle. Thus this abstracted connectivity problem has to be solved on a single torus. I could not see how this can possibly be done.




Figure 9: (a) *Abstract paper strip model characterizing the needed connectivity between clusters.*
 (b) *Face-centered fundamental net of R5.13.*

Wednesday 9/8/10 (evening):

The time seemed to have come to try to attack this problem from a completely different direction. For another difficult embedding problem, concerning map R3.3, I was able to fold up the fundamental net of the map and thereby obtain some insight into the routing of the various edges through the different handles. With the same goal, I extracted the fundamental net of R5.13 on the Poincaré disk (Fig.9b) constructed symmetrically around the center of the blue octagon shown in Figure 1. I then deformed this domain so as to obtain four 90° corners that all corresponded to the center of the red octagon and carrying equal edge labels on opposite parallel edges (Fig.10a). Thus the bounding rectangle of this domain could readily be rolled up into a torus (Fig.10b), thereby bringing together the four red corners and completing the center of the red octagon.

I also formed similar 90° corners corresponding to the center of the yellow octagon (Fig.10a). By suitably narrowing and elongating the flaps that carried these four yellow corners, I was able to also join them and thereby complete the center of the yellow octagon. However, at that stage the model became very contorted and crumpled, and it was not possible to further join some of the 8 instances of the remaining vertices which all corresponded to the center of the green octagon. I was also not able to read off the needed edge configurations on all the bridges that would have to be inserted between the edges of the fundamental net that were still left open. Perhaps a looser model with narrower, more spidery flaps carrying the red and yellow vertices might bring more clarity.



Figure 10: The fundamental net of R5.13: (a) a first distorted cutout that allows easy joining of some matching vertices; (b) the beginning of the fold-up process into a torus.

Thursday 9/9/10:

I drew a better template for the fundamental net, making sure that the red and yellow corners were all on suitably long and flexible flaps so that they could be properly joined (Fig.11a). I also colored all the regions with the proper colors of the facets and I labeled all the edges so I could keep track of them more easily. Then the model was folded up. I joined the red corners first, forming a (flattened) torus. Joining the yellow corners seemed to produce a twisted convoluted bridge. I was wondering: Is this still embeddable in a clean genus-2 surface? – I still needed to join the eight instances of green vertices. Once again this model became too cramped and too twisted to actually do this.




Figure 11: A different distortion of the fundamental net of R5.13: (a) a template for half the net; (b) the fold-up process at the stage where all red and yellow vertices have been merged.

Friday 9/10/10 (am):

Since the physical model did not help, I was trying to gain some insights from a virtual model. I first took a slightly enhanced template of the fundamental net (Fig.12a) and used it as a texture around a torus, thereby readily joining the red corners. The advantage of this model was that I could easily shift and twist the texture on the torus and thereby move the two instances of the yellow vertices into relative positions where it was easier to see what kind of bridges would be needed to actually merge these vertices and the associated edges.




Figure 12: *Partial embedding of R5.13 on a torus: (a) a suitably distorted fundamental net; (b) the net used as a texture, mapped onto a torus with 360° twist.*

This model made it clear that a single bridge is not sufficient to make a complete merger of all yellow vertices; at least two bridges are required. I quickly built a physical model of this, by printing the texture of Figure 12a on a strip of paper and folding up a thin toroidal ring (Fig.13a). This ring was enhanced with two handles carrying two yellow/green edges each from one side of the torus to the opposite one (Fig.13b). This then completed the closing of the red and of the yellow facets. What remained to be done was to also close the green facet. – Two more handles may be used to accomplish this task.




Figure 13: *The texture of Figure 12a applied to a toroidal paper strip: (a) first model; (b) narrower strip, enhanced with two handles to close off the yellow octagonal facet.*

Friday 9/10/10 (pm):

I also thought more about the shape of the octagonal facets themselves. They form basically a central cross connected to four “claws” that join pair-wise with their tips in pairs of vertices (Fig.5c,d). Since they join in a twisted, reverse-order manner, the claw tips have to be routed through a pair of holes or wrapped around a pair of tubular handles in order to appear in the right sequence. The challenge is to do this simultaneously for four separate octagons with only four holes inserted into the main torus.

I tried many ways of warping and shifting the four 2-hole clusters (where two claws meet) against one another to avoid illegal edge crossings (Fig.8a-d). Nothing seemed to work. There is a fundamental division of the toroid surface by the pairs of claws of the same color (e.g., Fig 8b: red, green) so that areas of different color (blue, yellow) that need to be connected are always separated by a barrier. Perhaps the whole 2-hole cluster approach is misguided!

For I while I also tried to stretch and interleave the claw-clusters, so that they perform their edge-switching operation with two non-adjacent holes that lie at opposite locations on the torus, and which are separated by another hole in between (Fig.14). But I still could not find a way to connect all edges.

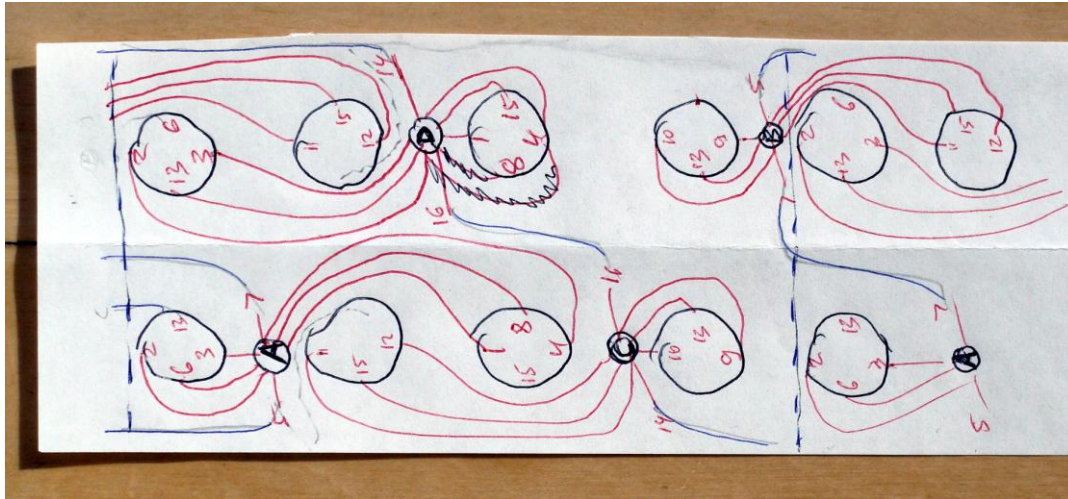


Figure 14: *Claw-clusters using two non-adjacent holes.*

I also thought some more about folding up the fundamental net. Perhaps I should start with a domain that was symmetrically centered on a vertex. This would then involve all four octagons in an identical manner. Another 4-way junction point may then also fall upon vertices and in forming the first 2-hole torus, it might become clear how D_2 -symmetry can be preserved. – This is something to try in the future!

Friday 9/10/10 (late):

I continued trying to make the “interleaved claw-cluster” idea work, but could never match up the proper edge terminations in all four holes (Fig.14).

Saturday 9/11/10:

I made a cleaner toroidal paper-strip model with more labels added. This made it easier to add the needed handles to close the yellow facets (Fig.13b). It made it also easier to see what handles have to be added so that the two instances of vertices A,B,C,D and the eight instances of the green octagon center can be brought into coincidence. I started to question whether just TWO additional handles would be enough to allow me to accomplish all this.

For a while I contemplated whether I should look at the general way of transforming the fundamental net to make a genus-5 handle body to get some idea how all the edges might be routed without intersections. The problem is, I cannot see what net is required to make such a body that also has D_2 symmetry. The default construction would lead to five loops joining in a single 10-way junction.

It is somewhat depressing that every idea seems to run into a dead end.

Sunday 9/12/10:

I decided it was time to try a different approach again. Perhaps the 2-hole clusters are too confining, since all four edges between a given pair of vertices run through the same pair of holes. Perhaps some of these edges have to run around the large toroidal ring. So I started with that premise and ran pairs of all four edge types around the big toroid. I actually put these edges at a diagonal slant within the rectangular net that represented the torus domain, so that I could introduce some global twist into the toroidal ring

(Fig.15). This also allowed for the possibility to treat the major and minor loops in the torus identically. For all of these various starts, the remaining two instances of each edge type will then be routed through four additional handles to be judiciously inserted in the right places.

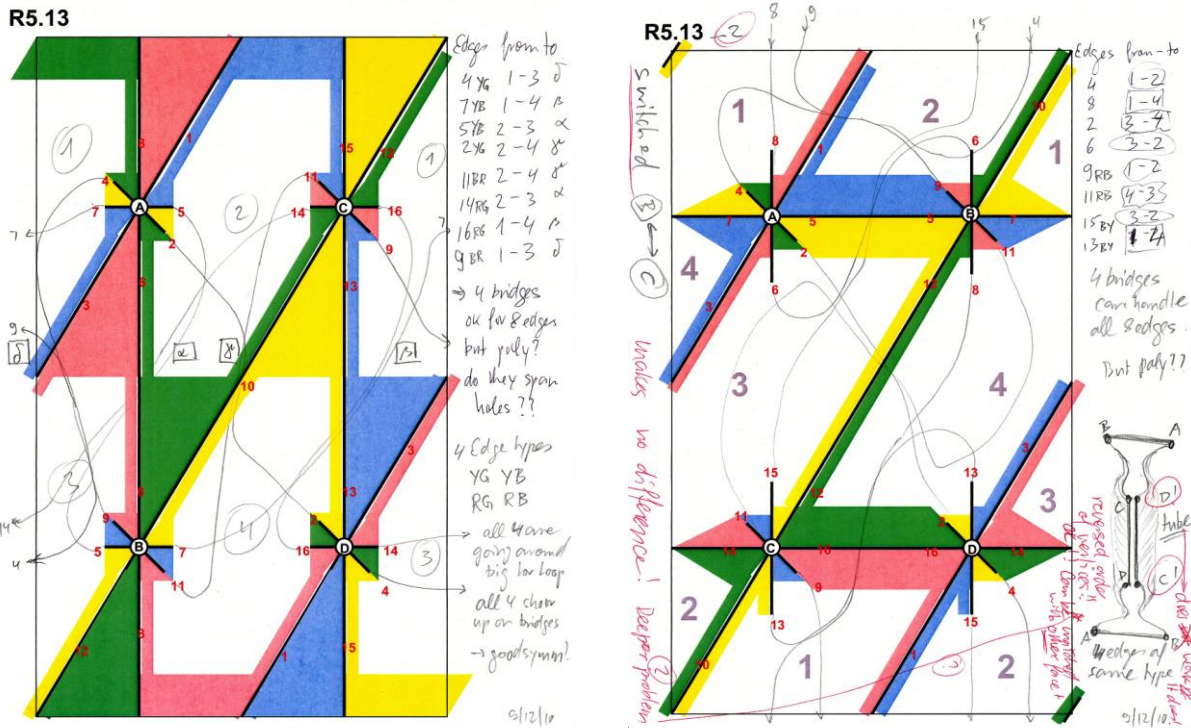


Figure 15: Routing half of the edges around the main torus: (a) all four edge types appear as major (vertical) toroidal loops; (b) two of the edge types are forming minor (horizontal) loops.

I drew out the diagonal edges on the rectangular torus domain and then noticed that another four edges can readily be connected without any crossings. There is a choice whether this second pair of circular loops also runs around the major loop of the torus (Fig.15a) or around the minor loop (Fig.15b). It seems to make no difference. In either case, we need four bridges connecting different pairs of quadrants. These four bridges can accommodate all eight edges that remain to be connected. Unfortunately, on each bridge, connecting a pair of quadrants, there are two edges that have different color pairs. So all four colors are present on each bridge, but there are only two edges. This means that the facets would have holes in them. This is not allowed!

If we examine the relationship between a pair of facets, say blue/yellow, we can glue together a pair of opposite edges (e.g., 5 and 7) and form a thin tube. This tube will be used to form a 2-colored handle. The other two edges of the same color combination (e.g., #13 and #15) will then be represented by some of the diagonal loops already placed on the big torus. This seemed to be tantalizingly close to a solution. I wondered whether by shuffling around the loops of opposite edges to form circles around the main torus in the appropriate way, I might find a solution with bridges that carry only **two** colors.

To figure out how these loops may have to be arranged, I followed up on the **tube idea** above. I made a blue and a yellow octagon in the shape of an extreme I-beam (Fig.16a) and connected pairs of opposite edges along the “bridge” part to form a tube (or double-sided, 2-colored paper strip). Then I twisted the hammer-head endings with the other two shared edge pairs, so that they could be merged in the proper orientation (Fig.16b). The resulting structure seemed to be embeddable on a toroidal ring with a twisted cross-bridge. This looked quite promising. But there are still two instances each of vertices A

and B. When these vertices were merged forcefully in the center of the toroid, the “bridge” turned into a twisted loop, the main function of which seemed to consist in carrying the edges #5 and #7 from A to B and reversing the order of their arrival. The first model became very twisty and messy, and I had to insert an extra strip of paper to make a longer, more flexible, twistable bridge. (Figure 16b is a clean subsequent reconstruction.)



Figure 16: Embedding two facets on a torus with a handle: (a) template for one octagonal facet; (b) a genus-2 structure formed by the two facets; (c) structure after merging the two instances of the vertices A and B and moving the “bridge” to the outside of the toroidal ring.

Sunday 9/12/10 (11pm):

I untangled the messy paper-strip model by cutting open the twisted bridge and by re-routing it around the outside of the main toroid ring (Fig.16c). This made the model less twisted and easier to understand. It has now genus 4. – But I have accommodated only **two** facets so far!

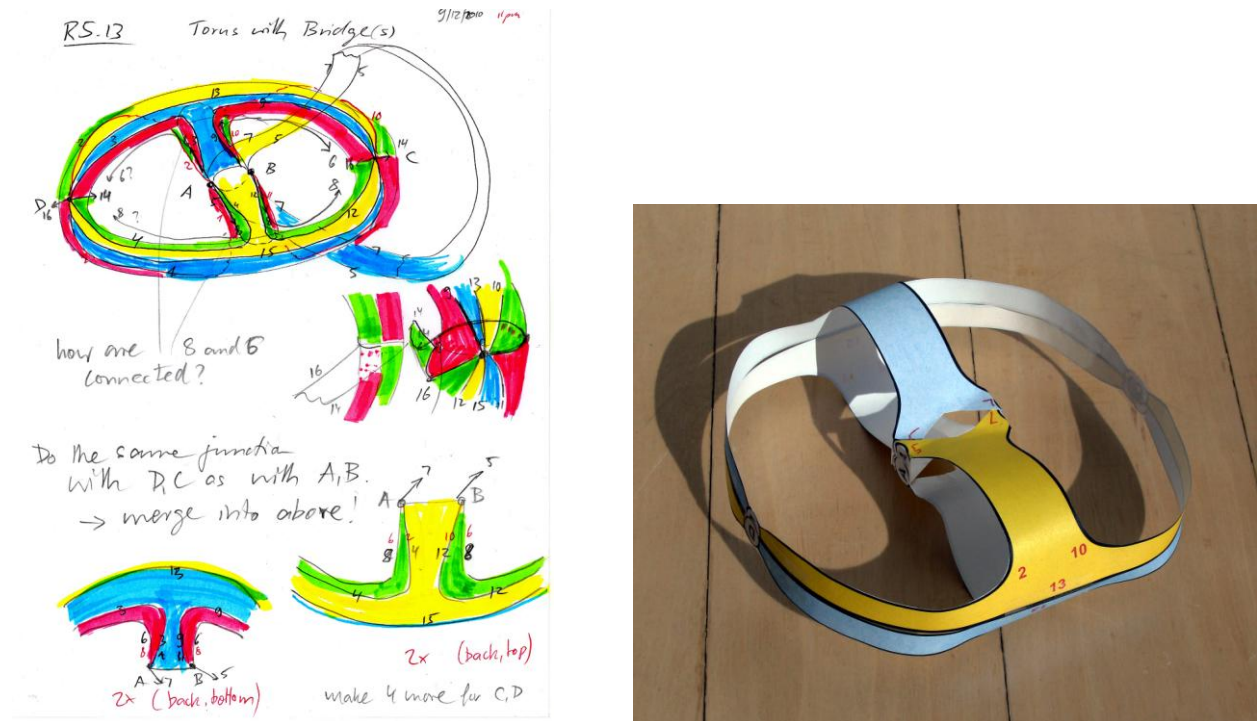


Figure 17: Breakthrough to a solution: (a) a sketch of the bridge structure, showing that a hole between A and B can replace the loop originating there; (b) a later clean reconstruction of this structure.

Based on information read off this model, I made a drawing of this embedding and sketched some stencils for making a paper-strip model of this new, more organized structure (Fig.17a). Four “T” shaped pieces with only two colors each were needed to cover the front and backsides of a paper-strip model of this configuration. In this drawing vertices C and D were located at opposite ends of the torus, but I had the hope that the same bridge construction used to accommodate vertices A and B could also be used with C and D, probably at right angle to the first bridge, so that overall D_2 symmetry could be maintained.

Monday 9/13/10 (6am):

The twisted double-sided 2-colored ribbon loop originating between vertices A and B breaks the D_2 symmetry. However, it occurred to me that the function of this ribbon, to exchange the order of two edges (#5 and #7) going from A to B, could also be performed with a **small tunnel through the bridge between A and B**. This tunnel could run along one of the D_2 -symmetry axes, and could carry the two edges whose positions need to be exchanged in the form of two helical spirals from A to B. (Figure 17b shows a model constructed later to demonstrate this concept more cleanly.)

Monday, 9/13/10 (10am):

I drew clean templates for this new configuration (Fig.18a) and assembled four T-units. I first joined two identical copies of the 2-sided T-units with their short arms, which carried two of the vertices forming a “bridge” with a small tunnel through its mid-point. When I then connected the (A,B) bridge with the (C,D) bridge, joining pairs of fat arms with matching color pairs, I did not obtain the expected large toroidal ring with two bridge structures. Rather, the four fat arms of the four T-units formed two separate loops. The overall structure thus was a loose ring with four interspersed smaller rings (Fig.18b). This could also be seen as a toroidal strip with four evenly spaced holes! Interestingly enough, the four vertices A,B,C,D were not located between the holes on the toroidal ring (as in my earlier attempts), but were sitting in opposite positions on the two small rings inserted in the major ring (Fig.18b).




Figure 18: An embedding for $R5.13$: (a) the templates for the four T-units; (b) the assembled toroid with four tunnels.

Monday 9/13/10 (7pm):

Knowing the basic structure of the solution, I was then able to draw it cleanly onto a torus strip with four holes. I constructed a nice symmetrical template (Fig.19a), which could then be looped into a toroidal model (Fig.19b,c). With this paper strip model I also tried to capture one more degree of symmetry: the inside-outside symmetry of the toroidal strip. (This would be a front-to-back symmetry if the strip is cut open and laid out flat. The rotational symmetry of the main torus ring would then be seen as a

translational symmetry in this strip). It was only a few days later that I realized that this model had only C_2 -symmetry – not the D_2 -symmetry I had been shooting for!



Figure 19: *Cleaned-up solution: (a) the template; (b), (c) two views of the C_2 -symmetrical model.*

Tuesday 9/14/10 (9am):

I made a new template with a somewhat simpler edge structure and all four holes of equal size (Fig.20a). I also enlarged the toroidal ring to make the final solution (Fig.20b) look more like that for R5.12.




Figure 20a: *Refined solution for R5.13 – the template (front and back).*



Figure 20b: *The refined model made in the same style as the final model for R5.12 (Fig.6b).*

I then pondered the essence of this solution. Why did I not find it earlier when I was looking at the plans shown in Figure 15 ? In the final solution, different combinations of edges were used to make the “horizontal” and “longitudinal” edge loops! Based on this insight, would it be possible to make an even more symmetrical (or easier to understand) solution on the surface of a pair of intersecting, horizontal and vertical toroidal rings? (Corresponding to the edge structure of a 2by2 square tile pattern on a torus). Since this is one of the regular, generated handle-bodies that van Wijk would have tried, the question arises: What is breaking the symmetry so that the full 4-fold symmetric solution with all identical patterns on each tube segment could not lead to a solution?

Monday 9/20/10:

It was only during the process of cleaning up my notes and writing this report that I realized that the models in Figures 19 and 20 did not have proper D_2 symmetry, since the geometry around the vertices A,D and around B,C is different. Also the shapes of the red and blue facets are different from the yellow and green ones. To obtain full D_2 -symmetry, these four vertices and all four face centers have to be placed on one of the symmetry axes. This can be achieved by flattening the major toroidal ring into its equatorial plane. The resulting embedding model is shown in Figure 21.

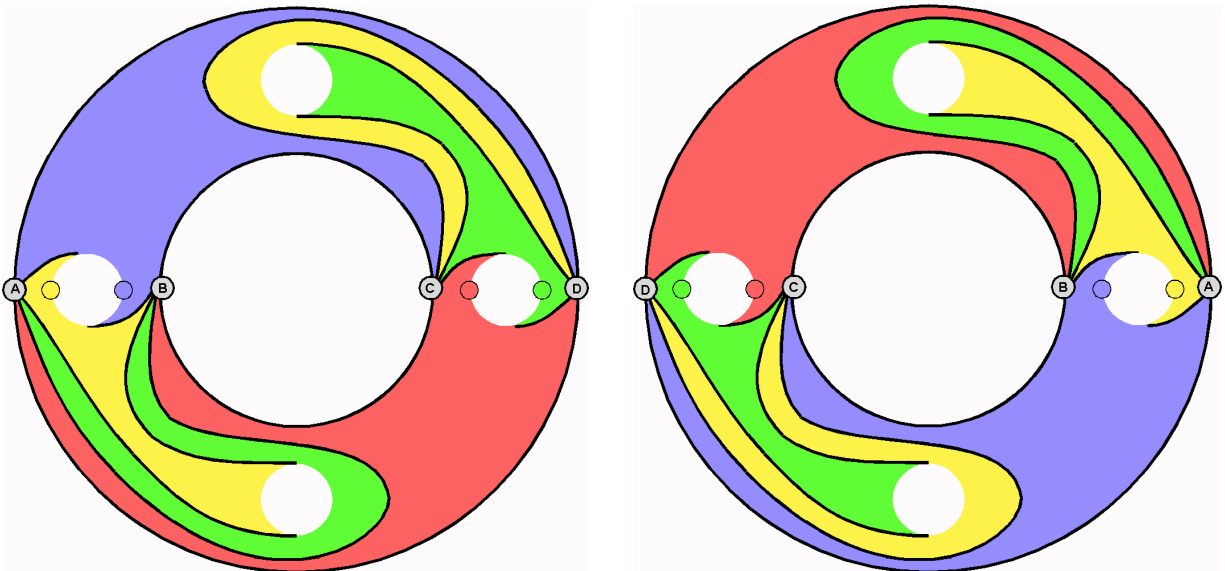


Figure 21: Refined solution for R5.13 (front and back) with full D_2 -symmetry.

3. Review of the Design Process

Looking back over this whole design process, I have to ask myself: Was this contorted, non-linear path really necessary? Was there possibly a more logical way to arrive at the solution more directly?

Clearly, some of the simpler initial steps had to be tried first, since there is a chance that they work and then lead to a simple, easy-to-understand solution. Thus, testing whether a different way of placing the vertices (as well as the centers of the octagons) onto the tubular handle-body that worked for R5.12, and then trying out various amounts of twisting of the edges on the various tubes, had to be tried first.

Of course, with the solution known, we can now map the found embedding back onto such a tubular structure. It then turns out that for R5.13 the four vertices, as well as the four octagon centers lie in the middle of four of the eight tubular segments that make up this handle body; the other four tubular segments each carry four longitudinal edges and four slivers of all four octagonal facets (Fig.22).




Figure 22: *The solution mapped onto an assembly of tubular segments (back and front).*

My subsequent attempts aiming for a D_2 -symmetrical solution based on a torus with four holes or four handles was definitely aiming in the right direction. It just turned out that connecting all edges by hand was too challenging. The final template (Fig.19a) then showed an interesting combination of connectivity patterns. Some portions were spirally with C_2 -symmetry, but others had mirror symmetry. If the handle-body were to be cut up into its individual tubular elements, homotopic to cylinders, the decoration on these handles would not be as symmetrical as the patterns used by van Wijk [17]; and that is the reason why his program could not find a solution for the map R5.13. However, using this D_2 -symmetric “canvas” with some judicious placement of the four vertices, and empowered by a brute-force path-searching program to connect matching edge stubs, with backtracking whenever it painted itself into a corner, might have found a solution. I may thus try to sign up some student(s) to develop such a program.

To see what my solution would look like on a tubular handle-body and to see how the patterns differed on the different tube segments, I constructed the virtual model shown in Figure 22. It turns out that three different texture patterns (not considering colors) are required to depict this solution; – two different edge patterns for the tubes that carry the vertices, and one more pattern for the connecting tubes with essentially a spiral edge pattern.

Perhaps the most important question is: After the successful design of the embedding of R5.13, am I now better prepared to tackle other difficult embedding problems? Did I learn some new technique(s) that will help me find other solutions more quickly?

The key steps that eventually led to the solution for R5.13 were not too different from an approach that I had taken earlier for solving the embedding of R3.3: Folding up the fundamental net to get a large fraction of the vertex mergers and edge connections taken care of. However, in R5.13 this first step led to a lesser degree of completion; more edges still remained to be connected, and more handles to be inserted to make this possible. Thus the paper model was not good enough to allow me to see the solution. At this point it was then really helpful to map the partial solution as a texture onto a virtual computer-graphics torus, where it could be shifted and twisted interactively by manipulating a couple of sliders. This experimentation allowed me to bring other elements that still needed to be connected into a proper

configuration where I could see what bridges, handles, or holes would be most helpful to make the necessary connections. This then allowed me to draw out some appropriate templates, which led to an improved paper model. By manipulating this physical model, trying to merge vertices with identical labels, I then discovered what topological form of connectivity would result. This “ping-pong” approach between virtual computer graphics models, and tangible, flexible paper models that could be easily manipulated and deformed in unexpected ways was crucial. I will certainly use this approach again as I try to solve the remaining two embedding problems of genus 5 and then start to tackle some regular maps of genus 6.

Acknowledgements

This work was supported in part by the National Science Foundation (NSF award #CMMI-1029662 (EDI)).

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] H. and C. Ferguson, *Eightfold Way: The Sculpture*. In: Levy [9], pp 133–173 (1999).
- [6] A. Hurwitz, *Ueber algebraische Gebilde mit eindeutigen Transformationen unter sich*. Math. Ann. Vol 41 (1893) pp 403–442.
- [7] D. E. Joyce, *Hyperbolic Tessellations, an applet to display your own*. – <http://aleph0.clarku.edu/~djoyce/poincare/PoincareApplet.html>
- [8] F. Klein, *Ueber die Transformationen siebenter Ordnung der elliptischen Funktionen*. Math. Ann. Vol 14, (1879). (Translation into English by S. Levy [9]).
- [9] S. Levy, *The Eightfold Way: The Beauty of Klein’s Quartic Curve*. Cambridge Univ. Press, 1999.
- [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] C. H. Séquin and L. Xiao, *K12 and the Genus-6 Tiffany Lamp*. ISAMA, Chicago, June. 17-19, 2004.
- [13] C. H. Séquin, *Patterns on the Klein Quartic*. Proc. Bridges 2006, Conf., London, pp 245–254.
- [14] C. H. Séquin, *Symmetric Embedding of Locally Regular Hyperbolic Tilings*. Proc. Bridges 2007, Conf., Donostia, pp 379–388.
- [15] C. H. Séquin, *Regular Maps on Cube Frames*. Bridges 2009, Banff, Art Show Proceedings.
- [16] C. H. Séquin, *My Search for Symmetrical Embeddings of Regular Maps*. Proc. Bridges 2010, Conf., Pécs, Hungary, pp 85–94.
- [17] J. J. van Wijk, *Symmetric Tiling of Closed Surfaces: Visualization of Regular Maps*. Proc. Siggraph 2009 Conf., New Orleans, pp 49:1-12.

APPENDIX

R5.6 Embedding - Design Log

This Appendix is an extension of the above report. It was added about two weeks after the above report was mostly finished, when I realized that the symmetrical embedding of R5.6 presented another formidable challenge. It thus constitutes an interesting test to see whether the insights gained in the R5.13 design effort and documented in the above report can help in finding another embedding.

Friday 10/1/10:

On this date I completed the labeling on the Poincaré disk for the regular map R5.6 (Fig.A4a). Making some educated guesses, using information about the number of vertices, faces, and the multiplicities of their connectivity, and exploiting symmetry to the fullest, let me complete this first phase in a few hours.

Studying the completed Poincaré disk (Fig.A4a), I noticed some patterns that may become useful when looking for an embedding with maximal symmetry:

- As expected from “ $mV = 2$ ”, opposite edges at all vertices form rings comprised of two edges. In my final labeling I labeled those edges with consecutive numbers: “ n ” and “ $n+1$ ”.
- There are also face rings consisting of four quads (always stepping straight across a quad to the opposite side). Alternate (opposite) faces in these rings have no vertices in common.
- Pairs of faces form checkerboard patterns, i.e., two such quads connect with all four corners to each other. At each vertex two such faces come together diagonally across from each other (faces “ o ” and “ m ” in Fig.A1a).

Because of the expected D_4 -symmetry, using a torus with four holes (as in R5.12/13) seems like a good first guess for a suitable handle-body. This handle-body could then be used in the following way:

- Two vertices each could be placed between adjacent holes.
- Since each vertex connects twice to four others, each vertex might connect to the two pairs of vertices beyond the nearest two holes (Fig.A1b).
- Half the edges might then be routed fairly directly (s = “short”: 90° ; orange) and the other half by going three quarters around the main torus ring (L = “long”: 270° ; green).
- The edge pattern around any vertex might then be: “ $ssLLssLL$ ”.

On second thought, this does not really work out! Opposite edges at a vertex must connect to the same neighbor vertex. The above “ $ssLLssLL$ ” pattern does not have that property, and I could find no other symmetrically repeating pattern where opposite edges are labeled “ s ” and “ L ”. Does this make this whole wiring plan obsolete?

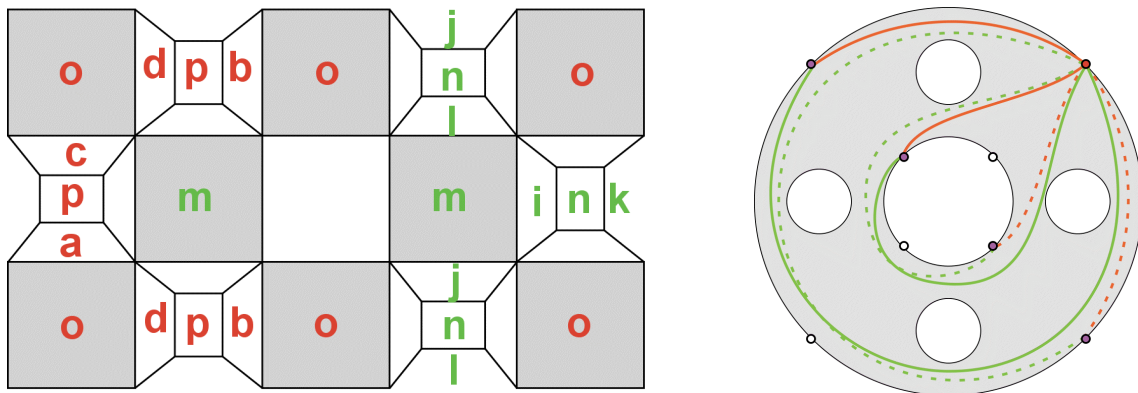


Figure A1: (a) Checkerboard pattern among face-pairs. (b) A possible vertex connectivity pattern.

Saturday 10/2/10:

I tried to accommodate one of those checkerboard face pairs, where two face diagonals form a closed ring, by putting that ring around the minor circle of a torus, e.g., between vertices **E** and **G** in Figure A2a. But I then realized that this would be a bad choice, since it would cut off all pathways running around the main torus ring for more than 90° .




Figure A2: (a) Possible placement of a checkerboard face-pair; (b) accommodating 2-edge rings.

I realized that it is crucial to accommodate these 2-edge loops (and also the 2-face-diagonal loops) in a systematic manner. One possibility is to place such loops like the pair of intersecting circles of radius R that go through the hole of a torus, passing once around the minor and the major loop. Figure A2b is a simple planar projection of four such circle pairs, where the drawn edge pattern would occur identically on the front and on the backside of this annulus. Now the four small holes have to be introduced into the torus and used in such a way as to avoid all edge intersections. However, this may not be directly useful for R5.6, since each vertex here connects with $mV=4$ to its partner. Moreover if the circles drawn are seen as the face diagonals of checker-boarding face pairs, then the dots would represent face centers – but there are only eight of them showing, and we do need 16.

Sunday 10/3/10:

I started to think about a **simultaneous** D_4 -symmetrical allocation of **vertices** and **face centers**, since the dual map R5.6d would also have to have this same symmetry:

The vertices could be placed as suggested above: two each between adjacent holes. They could be located at the inner and outer edges of the main torus ring, half-way between the holes, as shown in Figure A2, or at the edges of the small holes, closest to the adjacent hole (Fig.A3a, red disks). Since I am not shooting for any mirror symmetry, they could be located anywhere between these two mirror-symmetrical locations (Fig.A3a, orange ellipses), thus creating a C_4 -symmetrical pattern. The face centers in turn could be located at highly symmetrical locations marked by green squares in Figure A3a; or they could also be shifted away from these mirror symmetrical positions in the same way as the vertex locations. If this annulus is converted into a toroidal paper ring as in Figure 20b, then we still achieve D_4 -symmetry; the blue dots would mark the penetration points of the four C_2 -symmetry axes.

To make drawing easier and to show explicitly the symmetries between the inner and outer ring of the toroidal annulus, I could just draw a rectangular fundamental domain with a single hole in it, which is then repeated four times (Fig.A3b). The two sides of this rectangle (yellow) wrap around left to right; and the front and back of this domain would represent the inner and outer surfaces of a toroidal paper strip (as in my 3D models of R5.12/13).




Figure A3: (a) Placements of vertices and face centers on a 4-fold symmetrical annulus; (b) one fundamental domain of a toroidal ring, with some tentative edge placements.

Monday 10/4/10:

This is the day on which I actually started to write down the log captured in this Appendix, since I realized that finding a symmetrical embedding for R5.6 might be another challenging design exercise.

Today I started by using the above rectangular fundamental domain (Fig.A3b) to try out some edge placements. I placed two faces forming a checkerboard pair (Fig.A1) roughly in opposite positions on the main toroidal ring. I then tried to place all four face diagonals from the two quad centers to their four vertices without any crossings. However, I did not succeed.

Tuesday 10/5/10:

I tried to establish a nicely-bordered fundamental net on the Poincaré disk (Fig.A4a) for future fold-up experiments. When starting with face “o” in the middle, face “n” seems to resist a nice integration into a smooth boundary for the fundamental net. Thus I constructed a new fundamental net, centered around the vertex **F** (Fig.A4b). Eight complete quads were placed around it; and an additional half-quad was attached to each of the 16 edges of the perimeter of this first ring of quads. All of these half-quads added one extra edge and one face diagonal to the boundary of the fundamental net, which all connect to the vertex **A**. Vertex **A** can thus be viewed as being located “opposite to vertex **F**.”

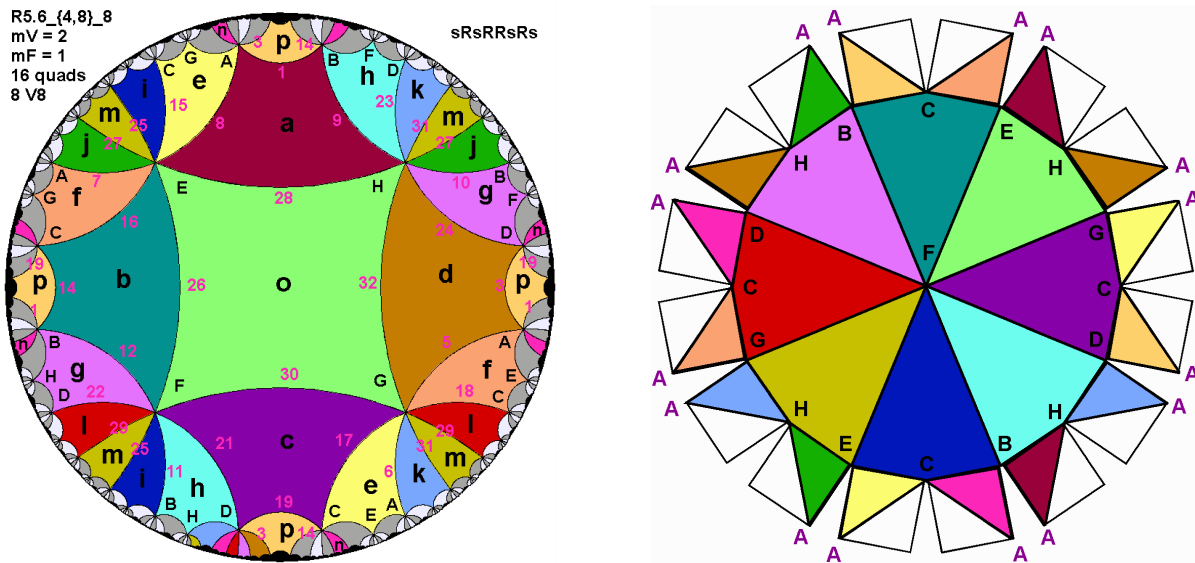


Figure A4: (a) R5.6 on the Poincaré disk; (b) vertex-centered fundamental net.

The complete fundamental net then showed four instances of vertices **C** and **H** on its boundary. Either one of which could serve as the corners of a rectangle (Fig.A5a) that could readily be folded up into a toroid (Fig.A5b). Thus the situation is very similar to what I encountered when I folded up the fundamental net of R5.13. I decided to review carefully what I did in the case of R5.13, so that I might then follow a similar path to find a solution.

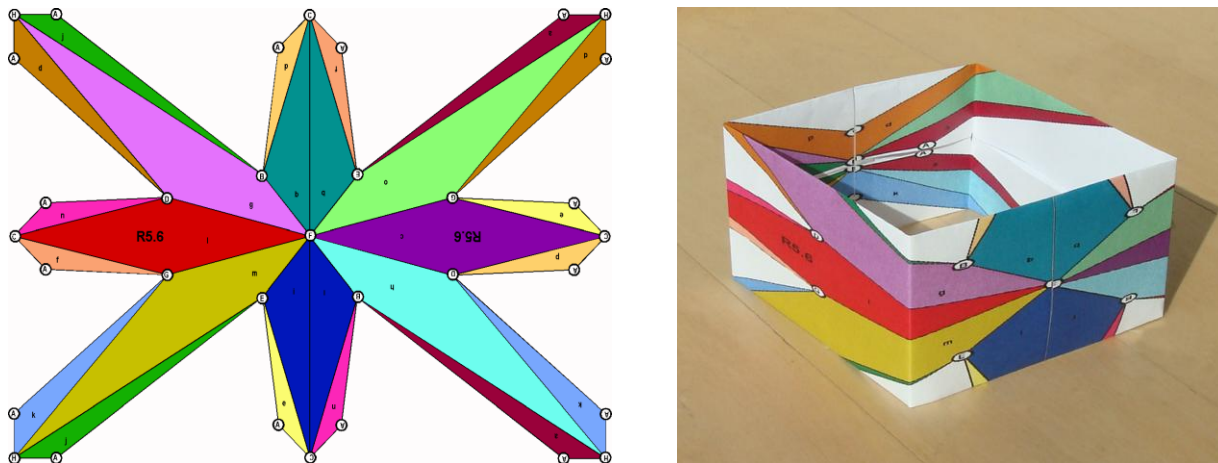


Figure A5: (a) The restructured fundamental net ready for fold-up; (b) the resulting toroid.

Wednesday 10/6/10:

I spent considerable time re-reading my report on R5.13, and in the process also cleaned up the wording in some sections. I was surprised to notice that after only two weeks I already had some difficulties understanding exactly what I did and what my thinking was behind some of my model constructions. I also re-built the models for Figures 16b and 17b.

Since the breakthrough for R5.13 came with the various experiments folding-up its fundamental net, I also focused on a similar approach for R5.6. To my surprise the outline of the fundamental nets for the two maps had the same structure. There was a direct correspondence between the 32 vertices that defined the polygonal outline of these domains. In the case of R5.13 these points corresponded to instances of all four vertices and of three of the four face centers (Fig.11a). In the case of R5.6 they were instances of seven of the eight vertices (Fig.A5a). Thus it was very tempting to just see where those points ended up in the final solution (Fig.21) and then place the vertices for R5.6 at the same locations (Fig.A6a).




Figure A6: (a) Vertex placement in correspondence with Figure 21; (b) conversion to a D_4 -symmetrical handle-body; (c) further conversion to a D_4 -symmetrical disk.

I also spent some time looking for a direct correspondence between the two maps as they appear on the Poincaré disk (Fig.1b and Fig.A4a), where some vertices might be identified with some face centers or edge mid-points – but found nothing useful.

Thursday 10/7/10:

I have become concerned that there may be no gain in pushing this correspondence too far: The embedding of R5.13 has D_2 -symmetry – but now I am pushing for D_4 -symmetry! This is based on Matthias Goerner’s programming efforts, which indicate that there should be at least an **immersion** of R5.6 with D_4 -symmetry. Up to now, in almost all cases, the symmetries of my models seem to agree with the maximal symmetry predicted for any immersion of the corresponding regular map.

I spent most of the day reviewing the steps I had taken so far and drawing clean figures for this Appendix.

Friday 10/8/10 (6am):

It occurred to me that the handle-body shown in Figure A5a could readily be converted into shapes with D_4 -symmetry – either a 3D handle-body (Fig.A6b) or a 2D disk (Fig.A6c). Thus I decided to focus my energy on pursuing this correspondence with the final embedding of R5.13.

Friday 10/8/10 (11am):

I took the disk layout from Figure A6c and turned it into the cylindrical map shown in Figure A7. On this map I started to connect the shortest edges between adjacent pairs of vertices by comparing Figure 11a with and Figure A5a to find corresponding lines, and then checking Figure 21 to see where these lines ended up. Some easy-to-place lines in Figure A7 were not actually representing edges, but lines that stayed completely inside a single face, such as the lines marked with lower-case letters and shown in the color of the corresponding facet in Figure A4a. With these key lines in place, I then started to draw other edges in the upper half of the template strip (Fig.A7) – complete ones for shorter edges (such as 26, 21, 3 or 8) or partial edges for the longer ones (such as 28, 30, 19 or 1). I then rotated the emerging edge pattern through 180° and placed it symmetrically on the lower half of the strip. For once, I actually did all this drawing directly on the computer (using Canvas from Deneba), since the repetitiveness of the pattern could so easily be exploited with copy and paste commands. The exact same edge pattern also had to be placed on the back side of this strip. Thus the set of all black edges was flipped vertically and drawn once more as dashed orange lines (meaning that these lines lie on the back side).




Figure A7: The completely wired-up D_4 -symmetrical embedding of R5.6.

Friday 10/8/10 (4pm):

After the opening ceremony of the Blum Center, I finished the drawing A7 and then made partial copies of it in order to draw colored templates for an actual 3D paper-strip model (Fig.A8).




Figure A8: Front and back templates for a D_4 -symmetrical toroidal embedding of R5.6.

Saturday 10/9/10 (am):

I refined the above templates and created a 3D paper-strip model for the embedding of R5.6 (Fig.A9a).




Figure A9: D_4 -symmetrical embedding of R5.6: (a) paper-strip toroid; (b,c) disk model.

Sunday 10/10/10 (am):

I created the disk model shown in Figure A9b,c and did some final editing for this Appendix. It is interesting to note that I might have found this pattern directly, if I had tried in the template A3 to connect a particular vertex to the four vertices associated with the two adjacent holes, rather than trying to follow a connection diagram according to Figure A1b.

Conclusions

Thus, the learning experience on R5.13 was indeed beneficial and resulted in direct pay-offs for the task of finding an embedding for R5.6. Not only did the previous experience guide me sooner rather than later towards trying to fold up the fundamental net, but it even saved me the trouble of making many new intermediate paper models, since I could use a direct correspondence to just sketch the topology of the final result directly onto the disk with appropriate vertex placements.

At this point I have only one more embedding problem in the genus-5 group that needs solving. It will be interesting to see how much the above experiences will help me in that task.