

Visualizing and 3D Printing Colorful Maps on Surfaces

Timothy Sun

Department of Computer Science, San Francisco State University; timothysun@sfsu.edu

Abstract

Map coloring problems comprise a classic area of study in topological graph theory, and three-dimensional visualizations of such maps have been a common theme in mathematical art. However, there is a gulf between the mathematical and artistic work because the former is often purely combinatorial. The present work describes how, using known algorithmic techniques, one can visualize these combinatorial maps as so-called “pillowcase models,” a 2D representation first described by Séquin. These models are further adaptable for single-material 3D printing, where the different colors are printed separately and then assembled. The technique is illustrated on examples from the Map Color Theorem, an analogue of the Four Color Theorem for higher-genus surfaces.

Introduction

Given a “map” drawn on an orientable surface, what is the fewest number of colors needed to color its regions so that bordering regions are given different colors? For surfaces besides the sphere (whose solution is given by the Four Color Theorem), the *Map Color Theorem* of Ringel, Youngs, and others [5] solves this problem by exhibiting, for each orientable surface of genus $g \geq 1$, a *complete map* with $\lfloor (7 + \sqrt{1 + 48g})/2 \rfloor$ regions, where every region borders every other region. In such maps, each region must receive a distinct color. Their constructions are purely combinatorial and require background in topological graph theory to understand. Thus, visualizations of these maps, preferably as three-dimensional tactile objects, are crucial for a non-expert’s understanding of the problem. Séquin [6] introduces the *pillowcase model* as one possible framework for describing maps on orientable surfaces. In this model, the holes of the surface are oriented in the same direction, and subsequently any map on that surface can be described using two planar pictures: the viewpoints from above and below that look through the holes. One can imagine a plane passing through the surface that separates it into two halves, corresponding to those two viewpoints.

In this work, I describe a method for converting combinatorial maps into what I call “non-equatorial” pillowcase models. Using this technique, I found models based off of old and new combinatorial maps on $n = 10, \dots, 17$ and $n = 19$ regions (the cases $n = 7, 8, 9$ are well-trodden in Bridges and elsewhere). As seen in Figure 1, the smallest three cases were taken one step further and 3D printed. Each map in this work, in concordance with a goal in Séquin [6], exhibits at least one nontrivial rotational symmetry. These symmetries come in two forms: a “circular” symmetry, where the surface is rotated in the drawing plane, and a “flip” symmetry, where the top half is rotated to the bottom half, and vice versa. Note that for $n = 12$, the model has both types of symmetry, and hence has *dihedral* symmetry.

Combinatorial Embeddings

A graph can be treated as a topological space where vertices are points, and edges are line segments connecting those points. The orientable surface of genus g , i.e., the sphere with g holes (or handles), is denoted by S_g . Let $\phi: G \rightarrow S_g$ be an embedding of a graph G . A *region* is a connected component of $S_g \setminus \phi(G)$, and two regions are said to *border* each other if they are incident with the same edge (which is sometimes also referred to as a border). In this work, regions are not required to be simply connected, i.e., homeomorphic to a disk, and any two regions are allowed to have more than one border between them. The original proof of the Map



Figure 1: 3D printed maps on surfaces involving 10, 11 (with missing adjacencies), and 12 regions.

Color Theorem [5] contains these unusual artifacts, and if one desires, they can be removed. Regions can be made simply connected by pinching the multiple boundary components together, and duplicate borders can be removed via edge contractions in the graph. If each region is simply connected, the embedding is said to be *cellular*, and if each region is bounded by a 3-cycle in the graph, then the embedding is *triangular*.

The *dual graph* of a cellular embedding replaces each region with a vertex, and two vertices are connected by an edge for each border that their corresponding regions share. The dual graph has a natural embedding in the same surface: each dual vertex is placed in the interior of its region, and dual edges cross their corresponding primal edges. In Figure 2(a), the primal graph, represented by black vertices and edges, has been embedded in the torus, and its dual graph is drawn in red.

The *complete graph* K_n is the simple graph on n vertices where every vertex is adjacent to every other vertex. The proof of the Map Color Theorem [5] is almost entirely about finding embeddings of such graphs in the surface of genus

$$\gamma(K_n) = \left\lceil \frac{(n-3)(n-4)}{12} \right\rceil.$$

Such embeddings are triangular or almost triangular, and their duals are complete maps. Unless otherwise stated, when I refer to maps on n regions, they are always complete and in the surface of genus $\gamma(K_n)$.

The drawings in this work are derived from rotation systems and the theory of current graphs (see Gross and Tucker [3] or Ringel [5]), combinatorial tools for describing graph embeddings. At present, current graphs are the only known proof method for the entire Map Color Theorem. While understanding these theories requires background in topological graph theory that the *viewer* is not expected to have, the *artist* can utilize them for visualizations.

Given a simple graph, a *rotation* at a vertex is a cyclic permutation of its neighbors. A *rotation system* is




Figure 2: A toroidal graph and its dual (a), and slicing a handle along a dual cycle (b).

an assignment of a rotation to each vertex in the graph. The Heffter-Edmonds principle states that orientable cellular embeddings and rotation systems are in one-to-one correspondence, up to orientation-preserving equivalence of embeddings. Cellular embeddings induce a rotation system from the clockwise ordering of edges leaving a vertex, and one can find a cellular embedding matching a rotation system via a procedure called face-tracing. For example, the primal graph in Figure 2(a) has the following rotation system:

0.	(1	3	4	2)
1.	(2	4	0	3)
2.	(3	0	1	4)
3.	(4	1	2	0)
4.	(0	2	3	1)

The rotation system has cyclic symmetry because if we interpret the vertices as elements of the cyclic group \mathbb{Z}_5 , the permutation $i \mapsto i + 1$ is an automorphism, i.e., it results in the same rotation system. Alternatively, we can think of the entire rotation system as being generated from a single cyclic permutation: the rotation at each vertex i is equal to the rotation at vertex 0 after adding i to each element. Heffter [4] is the first to consider such rotation systems, and the theory of current graphs generalizes his technique and presents these generating cyclic permutations in a more convenient graphical form. The interested reader can find the current graphs I used in Sun [8] (for $n = 13$ and 14) and Ringel [5] (for all others).

The backtracking algorithm of Altshuler and Brehm [1] for finding triangular rotation systems successively adds triangular regions to the surface until it becomes closed. Symmetry constraints are straightforward to incorporate: given a target symmetry group on the vertices, any time a region is added, every region in its orbit is also added. This modified algorithm found a minimum genus embedding of K_{10} with twofold symmetry, but it did not find one for K_{11} with any nontrivial symmetries. I settled for a triangular embedding with twofold symmetry of K_{11} with the edges of a 4-cycle deleted. The missing edges can easily be added using a handle, as in Figure 3.

The Automated Part: Surface Slicing

My method for finding pillowcase models starts by taking a rotation system of an embedded graph (e.g. the dual of an embedding of a complete graph) and repeatedly simplifying the surface. On a connected surface, a simple closed curve is said to be *non-separating* if cutting along the curve does not disconnect it. One can think of non-separating cycles as wrapping around handles, so cutting along such a cycle decreases the genus. A standard operation in computational topology [2] is to identify and cut along non-separating cycles, where



Figure 3: Adding four missing edges to the 11-region map with a handle.

both the input and output are represented combinatorially.

We call the intersection of the surface of a pillowcase model and its cutting plane the *equator*. The method begins by determining the components of the equator. The “internal” components, indicated by solid black circles in our figures, are non-separating cycles. For the 3D printing method (see below) to work, pillowcase models are further required to be *non-equatorial*, that no border runs along the equator. As a side note, in a non-equatorial model, it is perhaps easier to verify that all of the desired incidences are present—determining the regions demarcated by a border along the equator requires looking at both halves. The non-equatorial condition causes each of the aforementioned non-separating cycles to intersect the embedded graph at only a finite number of points. Topologically, such cycles are equivalent to cycles in the dual graph, as seen in Figure 2(b). There are algorithms for finding the shortest non-separating cycle [2], but my program simply tries all of the 3-cycles and 4-cycles in the dual graph.

One can always reduce a closed orientable surface to a sphere (with boundary) by repeatedly slicing along non-separating cycles. However, if one wants to preserve the symmetry of the original rotation system, then the symmetry group must also act on the chosen non-separating cycles. For maps expected to have a circular symmetry, the genus of the surface might require a hole to be placed in the center of the model. For example, in Figure 5, the genus is 10, so if one wants a threefold symmetry, there is at least one fixed hole. I found that Hamiltonian cycles, ones that include every vertex, are useful choices for such holes.

My backtracking program chooses a cycle in the dual graph and finds the other cycles in its orbit. Such cycles may share vertices or edges of the dual graph, but the program can only slice along them if they can be perturbed so that they are pairwise disjoint. If this process is successful, then it is repeated recursively until the surface is a sphere. Otherwise, the original cycle is discarded and the program tries another one.

The Manual Part: Drawing the Final Model

Once the embedding has been sliced by the above procedure, the combinatorial pillowcase model is generated by cutting the sphere in half so that the two boundary components of each sliced cycle lie in different halves. I manually chose where to halve the sphere, preserving any flip symmetry while aiming to (a) balance the sizes of the two halves, and (b) minimize the number of edges crossing between the two halves. Then, I drew each half as a planar drawing. For those with circular symmetry, I used polar coordinates to draw sets of borders simultaneously, which helped to avoid crossing lines.

The final maps are shown in Figures 1 and 4-9. Except for $n = 11$ and 13, these maps are complete. As mentioned earlier, the map on 11 regions is missing four edges. The drawing for $n = 13$ actually shows the dual of a “split-complete” graph [7, 8], one where a vertex of a complete graph has been split into two nonadjacent vertices, and each of the other vertices are adjacent to exactly one of those two vertices. If the two white regions denoted by asterisks are connected by a handle, then we would obtain a complete map on 13 regions. Such a handle can easily be incorporated into this model, but at the cost of the flip symmetry.




Figure 4: A “split-complete” map on $n = 13$ regions with flip symmetry.




Figure 5: A complete map on 14 regions with circular symmetry.

A 3D Printing Method for Non-Equatorial Pillowcase Models

The most popular type of hobby 3D printer uses *fused deposition modeling* (FDM), where a nozzle extrudes molten plastic to create the desired object layer-by-layer. While one can simply paint the surface of the object, I opted for printing regions and borders individually using the wide range of colorful plastic filament available. The approach is specifically designed for single-material printers, which are more affordable and less wasteful than multi-material printers.

Perhaps the greatest obstacle in FDM 3D printing is that of *overhang*, where a 3D model is placed on the printing bed, but part of the model has nothing underneath it. The nozzle cannot print in midair, so one would often need to place “support material” underneath the overhanging parts. Removing the support material leaves behind a noticeable rough patch that is not only aesthetically displeasing, but also potentially detrimental to any assembly process where the pieces fit together snugly. I transformed the pillowcase models for $n = 10, 11, 12$ into 3D objects by first giving them thickness, and then rounding the objects near their equators. These objects can then be sliced in half along the equator, and each half can be placed cut-side down to produce two solids without overhang.

Qualitatively, I observed that visible black borders help to separate non-complementary colors. These borders also divide the surface into regions and provide a frame for other pieces to attach to. Since the colors inside the surface are not visible, one can choose how each region or border extends into the surface. Each




Figure 6: A complete map on 15 regions with circular symmetry.




Figure 7: A complete map on 16 regions with flip symmetry.

piece, including border objects, can be extended vertically down to the cutting plane to obtain pieces that are all free of overhang.

The maps are constructed by first specifying the borders as curves on the surface. After giving the borders thickness (Figure 10(a)), they are projected down to the plane and vertically extruded beyond the surface to obtain a “fence” (Figure 10(b)). The final shapes are found through *constructive solid geometry*, algorithms that create new shapes by applying Boolean operators to existing shapes. The border and country objects are found by taking the intersection (Figure 10(c)) and difference (Figure 10(d)), respectively, of the surface with the fence. This method was applied to the maps in Figures 1.

Conclusion

I presented a method for visualizing maps on surfaces, applying them to complete (or near-complete) maps on $n = 10, \dots, 17$ and $n = 19$ regions. The cases $n = 10, 11, 12$ were physically realized through 3D printing, and in principle, with a large enough 3D printer and enough colors, one could 3D print any of the maps generated by this method. However, there are still a few limitations and future directions to explore.

The resulting drawings do not always preserve all of the combinatorial symmetries in the original rotation




Figure 8: A complete map on 17 regions with circular symmetry.

systems, though sometimes this is not possible due to the number of handles. For those that seem to be improvable, the current graphs for $n = 15$ and 17 produce rotation systems with dihedral D_5 symmetry (i.e. reflections and rotations of a regular pentagon), while the drawings here only have cyclic \mathbb{Z}_5 symmetry. The rotation system for the complete graph K_{16} has $\mathbb{Z}_2 \times \mathbb{Z}_8$ symmetry, but my map scrapes by with only the order 2 subgroup of the second component. Are there pillowcase models that exhibit more of those combinatorial symmetries? In particular, can one obtain at least $\mathbb{Z}_2 \times \mathbb{Z}_2 \cong D_2$ symmetry for the map of K_{16} ?

The $n = 18$ case is notably absent from the above list. Historically, while K_{17} and K_{19} were among the first few cases of the Map Color Theorem to be solved, K_{18} was actually one of the last (see Ringel [5, p.6]). A current graph solution is now known [8], but the resulting embedding does not have any nontrivial symmetries. Does there exist a pillowcase model for K_{18} with at least a flip symmetry?

The manual drawing process was by far the most time-consuming part, especially when it came to the large map on 16 regions. To what extent can this part of the process be automated? Part of the effort is trying to minimize the size of the drawing, given some minimum width for the regions, but a program that can produce any drawing that preserves the symmetries, without too much optimization on the area used, would already be desirable.

On the 3D printing side, I encountered an issue with borders near or on the boundary between the two halves. If each half of the surface can be printed without overhang and the surface is smooth, the tangent plane at the equator is vertical. Then, for example, a border that runs along the equator would project down to a thin line, hence why the non-equatorial property was introduced. Is there a way to provide more volume to thin features with more complex piece shapes? In my 3D printed examples, I circumvented this issue by choosing surfaces that are mostly flat and deliberately drawing borders that cross the equator orthogonally. These shapes perhaps conflict with our usual notion of a torus in 3D as a solid of revolution of a circle.

Finally, there is an analogous *nonorientable* Map Color Theorem for surfaces like the projective plane and the Klein bottle. Can the techniques here be extended to those surfaces? Ringel [5] exhibits drawings of small maps on nonorientable surfaces with boundary (e.g. the projective plane as a Möbius strip), but I am not aware of any visualizations of larger cases.




Figure 9: A complete map on 19 regions with circular symmetry.




Figure 10: Given a set of borders (a), they are extruded vertically (b), segmenting the solid into overhang-free pieces (c-d).

References

- [1] A. Altshuler and U. Brehm. “Neighborly maps with few vertices.” *Discrete & Computational Geometry*, vol. 8, 1992, pp. 93–104.
- [2] É. C. de Verdière. “Computational topology of graphs on surfaces.” *Handbook of Discrete and Computational Geometry*, 2017, pp. 605–636.
- [3] J. L. Gross and T. W. Tucker. *Topological Graph Theory*. Wiley & Sons, 1987.
- [4] L. Heffter. “Über das problem der Nachbargebiete.” *Mathematische Annalen*, vol. 38, no. 4, 1891, pp. 477–508.
- [5] G. Ringel. *Map Color Theorem*. Springer Science & Business Media, 1974. vol. 209.
- [6] C. Séquin. “Easy-to-Understand Visualization Models of Complete Maps.” *Bridges Conference Proceedings*, 2023, pp. 187–194.
- [7] T. Sun. “Face distributions of embeddings of complete graphs.” *Journal of Graph Theory*, vol. 97, no. 2, 2021, pp. 281–304.
- [8] T. Sun. “Jungerman ladders and index 2 constructions for genus embeddings of dense regular graphs.” *European Journal of Combinatorics*, vol. 120, 2024, p. 103974.