# A Machine Resolution of a Four-Color Hoax

Stan Wagon

Macalester College, St. Paul, MN 55105; wagon@macalester.edu

## 1. The Hoax

In April 1975, Martin Gardner, in his famed *Scientific American* column, published the four-color hoax: a planar map that required five colors (Fig. 1). It was apparently not completely trivial to 4-color the map, since many readers believed the hoax. Of course, we now know that no such map exists.

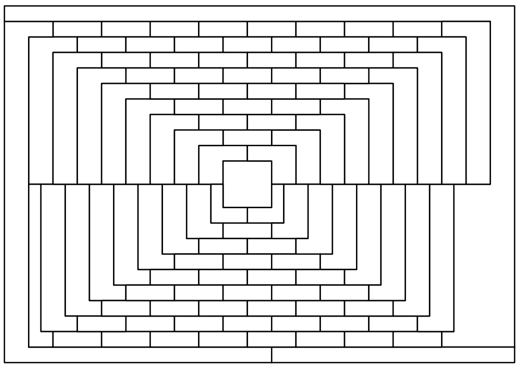

Figure 1. The April Fool's hoax map from 1975.

Developing an algorithm to 4-color planar maps leads to two interesting computational issues:

1. Devising an algorithm that will 4-color planar graphs. Since the proof that such an algorithm works would yield the 4-color theorem, one will have to deal with the fact that the algorithm is likely to be not proved to work. One hopes that it will simply work in practice.

2. Devising an algorithm to turn planar maps into planar graphs. This leads to the question of designing a good data structure for maps, and for planar drawings of graphs whose edges might not be lines, but will have to be broken lines.

The most notable work on (1) is by Morgenstern and Shapiro [9]. Unaware of their work, Joan Hutchinson and I [5, 6] developed a randomized algorithm based on Kempe's false proof of the four-color theorem. Their paper then pointed us to a 1935 idea of Kittell that leads to a much better algorithm. The Kempe–Kittell algorithm seems to work very well in practice. Point (2) required the development of a method of going from a map to its adjacency graph, and a planar drawing of it. It is important from several viewpoints that the drawing be such that the edges stay inside the countries being connected. Strictly speaking, simply having the cyclic neighbor structure for each country is enough to implement a coloring algorithm for maps. But it is reasonable to take the approach of first turning the map into a graph, and then use a coloring algorithm for graphs.

## 2. The Kempe Four-Color Algorithm

In 1879 Kempe gave an explicit method of 4-coloring planar maps, which I summarize here from the point of view of planar graphs. Assume a planar graph $G$ is given, with vertices labeled 1 through $n$; $R$, $G$, $B$, and $Y$ denote the four colors red, green, blue, and yellow. Choose the first vertex in $G$ —call it $v$ — having degree 5 or less (it follows quickly from Euler's formulas that every planar graph has such a vertex). Remove it. Color the remaining graph by induction on the vertex set. Then color $v$ as follows:

1. If the neighbors of $v$ have only 3 colors appearing among them, then there is a color left free for $v$. 2. If there are 4 neighbors of $v$ and all 4 colors appear among them, use a Kempe chain to eliminate one of the colors from the set of neighbors. 3. If there are 5 neighbors of $v$ and all 4 colors appear among them, use a Kempe chain to eliminate one of the colors from the set of neighbors.

A *Kempe chain* is defined as follows: Suppose the neighbors of $v$ have colors $R$, $G$, $B$, $Y$, in that order. Look at the subgraph (the "chain") containing the $R$-vertex and consisting of all $R$ and $B$ vertices that are reachable from the $R$-vertex by an $R$–$B$ path. If this subgraph does not contain the $B$-neighbor of $v$, simply switch all $R$s and $B$s in the chain. This turns the red vertex blue, and so red can be used on $v$. If the chain does reach all the way to the blue neighbor, then look instead at a green-yellow chain starting from the $G$-vertex. This cannot reach the yellow vertex because of the protective red-blue fence that separates $G$ from $Y$. So switching colors on that chain will free up green. This takes care of case (2) above.

For case (3) suppose the colors on the neighbors, in circular order, are $R$, $G$, $B$, $Y$, $G$ (if the $G$'s are adjacent, the reasoning is similar). Start with red, letting $u$ be the vertex so colored. Let $w$ be the neighbor of $v$ colored blue. Look at the connected component of $u$ among all vertices colored $R$ or $B$. If $w$ is not in this component, switch all $R$s and $B$s on the component, thus freeing up $R$ for use on $v$. If, on the other hand, the target $w$ does show up on the chain, then try to make a color-switch using the pair $R$–$Y$. If that fails, try to eliminate green by chaining from the first green to yellow and then from the other green to blue. Kempe's quite-believable topological argument shows that the

Kempe chain method will always succeed in eliminating a color.

## The Bad Example

It took a dozen years, but Heawood (also de la Vallée Poussin) saw the flaw in Kempe's argument. When dealing with the *R*, *G*, *B*, *Y*, *G* case just described, it can happen that one color switch can demolish a protective fence. Heawood's example is often quoted, but in fact there is a better refutation, due to A. Errera [3], that provides an explicit graph on which Kempe's algorithm, the one just outlined, fails. That graph, which I call the Errera graph *E*, has 17 vertices and is nicely realized as a map on the sphere.

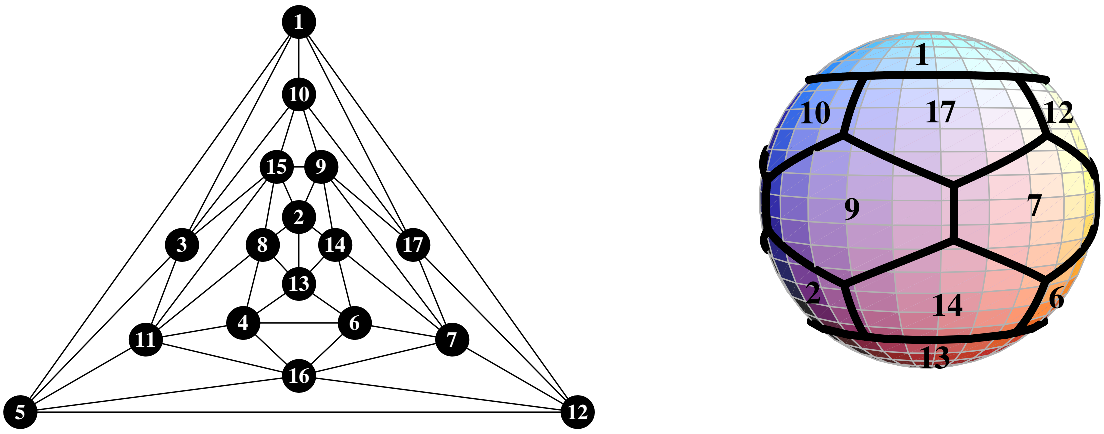


Figure 2. The Errera graph, drawn in the plane, and also as a map on a sphere. In the latter form it is the fullerene *C30*.

When all the vertices except 1 are colored, Kempe's method runs into an impasse (see Fig. 3). Vertex 1 has two green neighbors, 10 and 12. The chain (red edge in figure) from 12 to yellow (3) works and eliminates one of the greens, but causes a switch in 7's color from yellow to green, which breaks the protective red–yellow fence around 10 and causes the other green chain (purple in figure), from 10 to 5, to fail.

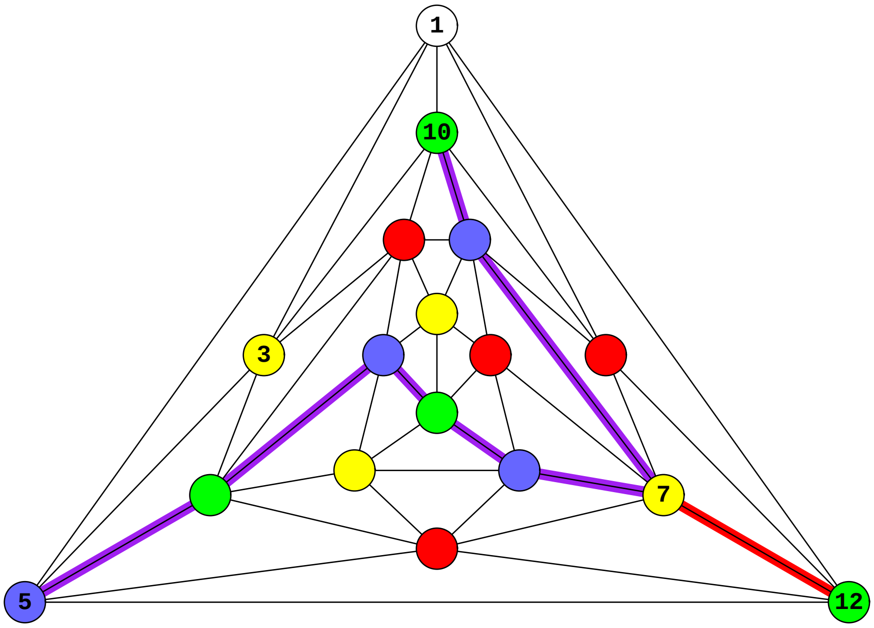


Figure 3. Kempe's algorithm on the Errera graph leads to an impasse, since the chain method fails to color the remaining vertex.

But the Kempe chain idea is so pretty, and counterexamples so rare, that it is not surprising that it leads to a good algorithm. The Errera graph is a counterexample only if it is labeled in a certain way. Empirical trials shows that roughly 10% of the labelings cause a problem; for 90% of the labelings, the Kempe order is different and the impasse does not arise. So one attack on the 4-color problem is: give the graph a random labeling and try Kempe's algorithm; if it fails, relabel and retry; continue. That will work for $E$. But it will fail badly on, say, 30 disjoint copies of $E$. For if 90% of the $E$-labellings are good, then only 4% of the labelings of $30E$ are good! As we will see, a simple modification turns this algorithm into one that, apparently, works much better. But we are left with an intriguing question regarding the random Kempe algorithm just presented, maely: Does it always halt?

QUESTION. Is it true that any planar graph admits a labeling for which Kempe's coloring method works? That is, does the algorithm just described — the ”relabel and retry” strategy — always halt?

## 3. Kittell's Improvement

Irving Kittell in 1935 had some ideas that he hoped would lead to a solution of the four color problem. From a computational viewpoint, his ideas can be viewed as leading to the following method: To eliminate the impasse that Kempe's method might run into, make color-switches on Kempe chains connecting two randomly chosen neighbors, even when those neighbors are consecutive and adjacent to each other. Note that Kempe would never have looked at such a color-switch, since it clearly will not free up a color. Kittell's point is that, nevertheless, it might cause a change that leads to resolution later. For the Errera graph, if one follows Kempe's method until the impasse and then tries random Kittell chain switches, then the impasse gets resolved in a small number of steps. In a typical case, it took 13 random switches to get to a position where only three colors appear among the five neighbors of vertex 1, leaving the fourth free to use on 1. The chosen pairs, after the switch on (12,7), were (10,5), (3,14), (10,12), (17,8), (12,10), (3,5), (5,12), (3,10), (10,5), (10,3), (5,3), and (12,6). In the last case the chain from 12 did not reach the target, 3, which is what freed up a color. In the (3,14) and (17,8) cases the chain did not reach the target either, but this did not free up a color, since the color at the start of the chain appeared twice.

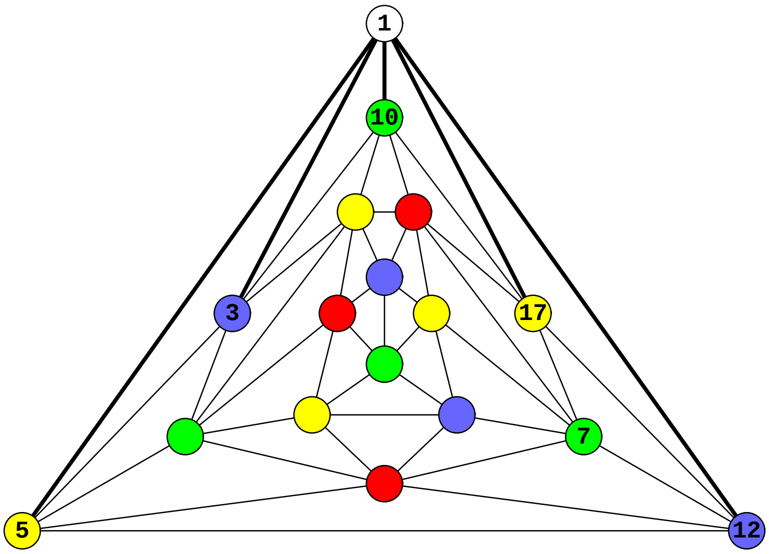


Figure 4. Eleven random Kempe switches resolve the impasse and free up red for the remaining vertex.

Thus we can formulate the following coloring algorithm. A variation on this idea is studied in [9]; they looked at many instances and made conjectures about running times. They also considered some other algorithms.

**The Kempe–Kittell Algorithm.**

1. Given a planar graph, label the vertices randomly. 2. Get the Kempe order of the vertices, by removing a vertex of degree 5 or less, and repeating until only a single vertex is left. 3. Color the single vertex red. 4. Add vertices back in the reverse order in which they were removed, coloring them as follows: i. Follow Kempe's method to the letter, switching colors on chains defined by nonconsecutive vertices in an attempt

to free up a color. ii. If Kempe's method fails, use Kittell's random Kempe-chain-switches until the impasse is resolved.

As before, a proof that this will always work is lacking. But it seems to work very well in practice, and has no problem getting a 4-coloring of the graph consisting of multiple copies of the Errera graph. The following conjecture was stated in [6].

CONJECTURE. The Kempe–Kittell algorithm always halts.

## 4. From a Map to a Graph

Given a planar map, with countries that are polygons, the simplest way to color it is to apply an algorithm to a planar representation of the adjacency graph (where each country is a vertex, and two vertices are adjacent if the countries share a part of a linear border). While simply having the abstract information about the cyclic ordering of neighbors around each country is adequate, it is more satisfying to turn the map into a planar graph, and then apply an algorithm that is set to work on graphs. Here is one way to do that (Wagon [13]), where the map is assumed to be given as a finite set of points, along with index sets delineating the points forming the border of each polygonal country.

1. For each country, locate a "capital" by triangulating the polygon and choosing the centroid of the largest triangle.

2. For each country, use the triangulation to generate a tree of triangles, and generate paths from the capital to the center of each border segment by uniformly dividing the diagonals that separate the triangles.

3. For each pair of adjacent countries, choose one of the border segments that are shared and form the piecewise linear curve from the capital of one country to the center of the chosen border edge, and then on to the capital of the second country. This construction yields a planar graph with edges that stay entirely within the two countries that they represent; see Figs. 5 and 6. In Fig. 6, Michigan is taken to be one state, and Lake Michigan is added as a region to keep the map simply-connected (so that Euler's formula applies).

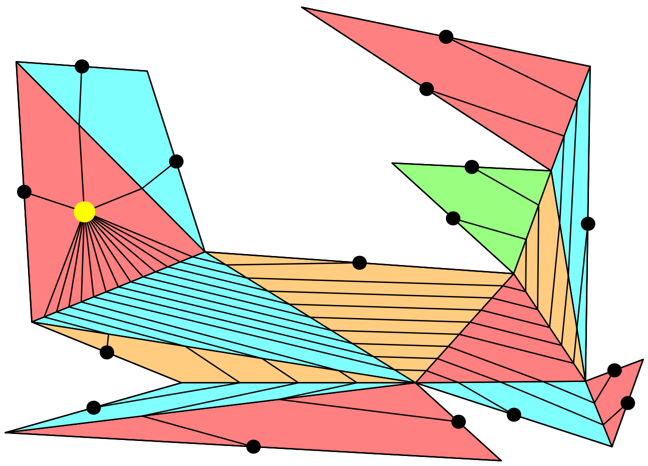

Figure 5. A triangulation yields disjoint paths from the capital of the polygon — the (yellow) centroid of the largest triangle — to the midpoint of each edge.

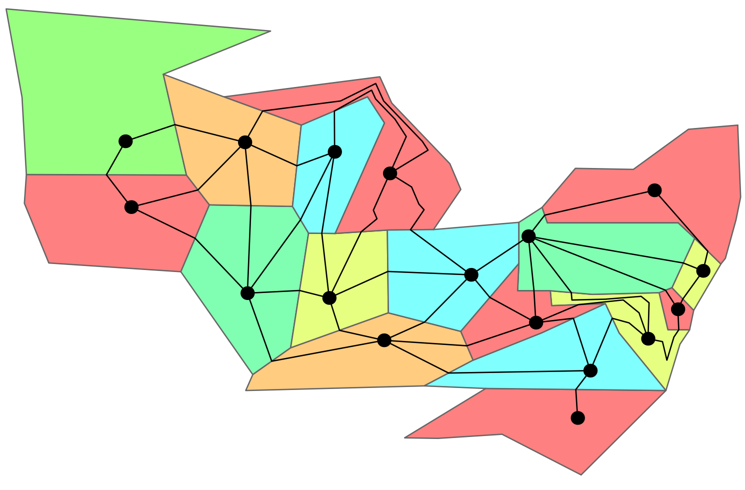

Figure 6. Choosing paths for each adjacency yields a system of broken lines that defines the adjacency graph of a map, here a part of north-central U.S.

It is not hard to further refine the edges in a couple of ways: The edges can be straightened by eliminating corners and checking, at each step, that the resulting drawing is still planar; and this can be done repeatedly. Or one can replace the broken lines with Bézier curves. I have programmed this to some extent (Fig. 7).

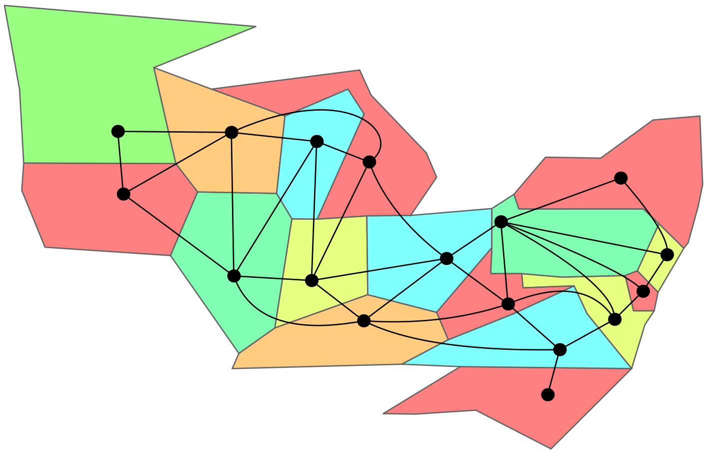

Figure 7. The adjacency graph of the map in Figure 6, using Bézier curves to smooth the edges.

And now we can put it all together to solve the hoax by computer. Figure 8 shows the adjacency graph for the April Fool's map, and Figure 9 shows a machine generated 4-coloring of the map using Kempe's method. There are no impasses, so Kempe's method, in the exact form he envisioned, has no problem coloring it.

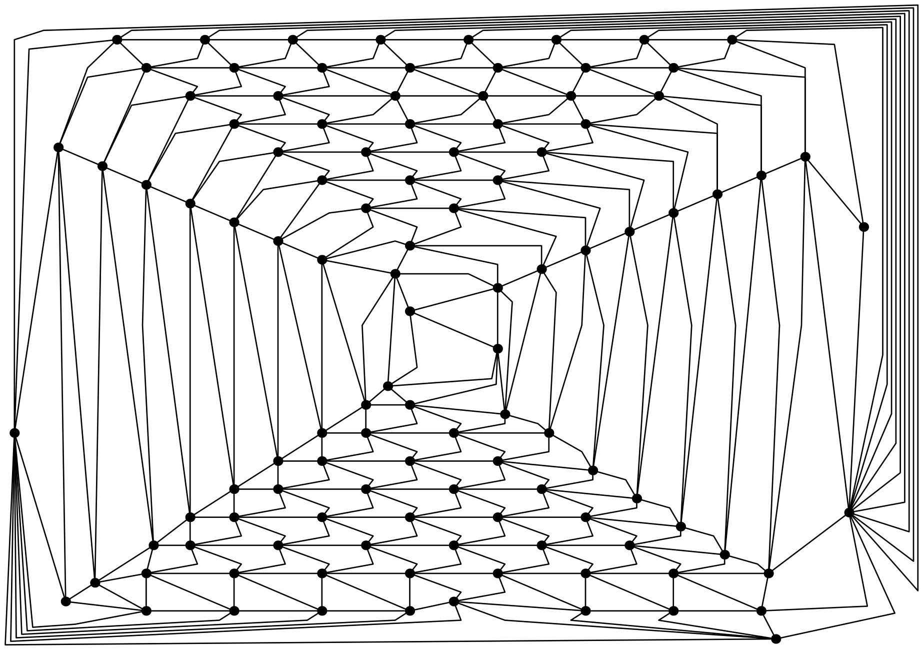

Figure 8. The adjacency graph of the April Fool's map.

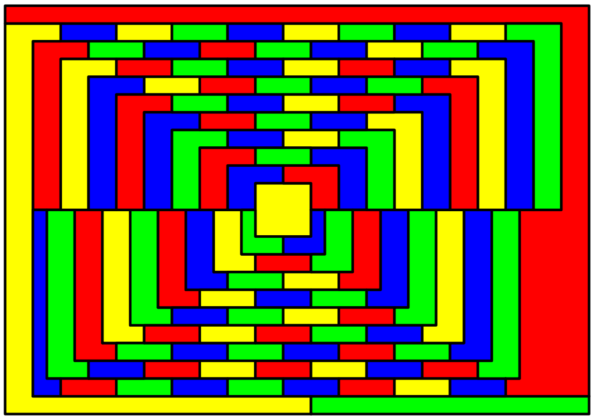

Figure 9. A 4-coloring of the April Fool's map obtained by Kempe's algorithm.

## 5. A Penrose Application

When applying the preceding map-coloring ideas to Penrose tilings, Wagon noticed that they seemed to use only three colors. The question as to whether any Penrose tiling could be 3-colored turned out to have been posed by John H. C. Conway many years ago. Sibley and Wagon [12] found an amazingly short proof that Penrose rhombus tilings are 3-colorable: it turns out to be quite easy to show that any such finite tiling has a tile with 2 or fewer neighbors. Indeed, this is true for any tiling by parallelograms that meet edge-to-edge or vertex-to-vertex (meaning: two tiles intersect in a full edge of each, a vertex of each, or not at all). It follows by induction that every such finite tiling is 3-colorable, and then from a general result about infinite graphs that any infinite tiling is 3-colorable. The case of Penrose kites and darts is harder, and was done later by W. Paulsen and R. Babilon, independently [2]. A first question (Sibley [11]) is whether any finite collection of parallelepipeds in 3-space that meet either face-to-face, edge-to-edge, vertex-to-vertex, or not at all, has one that has three neighbors at most. If this were true, then any such 3-dimensional map — the countries being the interiors of the boxes — would be 4-colorable. But the hoped-for result is false. A counterexample appears in [10], reproduced in Figure 10. The question whether such maps are 4-colorable remains open.

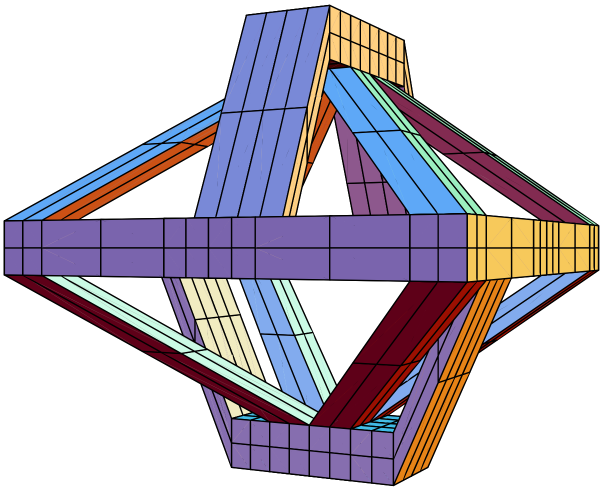

Figure 10. A 3-dimensional map made up of boxes that meet neatly and such that every box has four or more neighbors.

Returning to the traditional planar situation, we can take a spherical view of maps using parallelograms. Namely, consider such a map having the topology of a sphere; that is, imagine a convex polyhedron whose faces are all parallelograms. Is this planar map 3-colorable? Such polyhedra, in the more general situation where the faces can be *2n*-gons with opposite sides parallel, are called *zonohedra*, and these objects exhibit a lot of structure. Much has been written about them (see the treatise by Bjorner et al [1]). Yet apparently the coloring of the faces has not been considered. But under another guise, the question has been considered. It turns out that the map of the zonohedral surface is equivalent to the arrangement graph of great circles on a sphere. By this is meant any collection of great circles on the sphere such that no three have a point in common; the vertices of the graph are the intersection points of the circles, and the edges are the arcs connecting neighboring points. Felsner et al [4] raise the interesting question of whether such arrangement graphs are 3-colorable, and that question remains open. Steven Tedford (Franklin & Marshall College) has recently shown that if there are $m$ great circles and they somewhere form an $m$-gon, then the entire arrangement is 3-colorable. The great circles are easily projected onto the plane, where they become circles. But circle graphs in the plane can require four colors, as proved by Koester [8]. His example is shown in Figure 11.

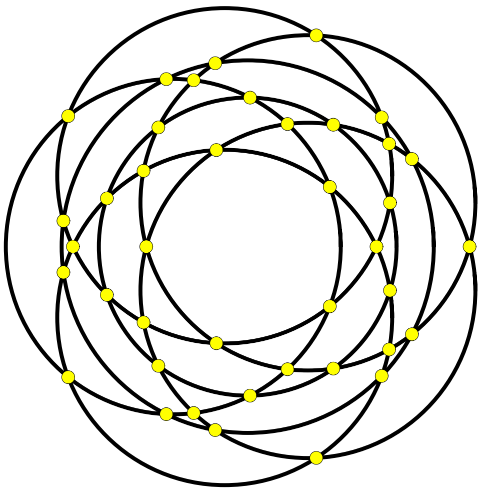

Figure 11. An arrangement of circles in the plane that leads to a graph requiring four colors.

**Acknowledgement:** I am grateful to Joan P. Hutchinson for many valuable contributions at all stages of the preparation of this paper.

## References

1. A. Bjorner, M. Las Vergnas, B. Sturmfels, N. White, and G. Ziegler, *Oriented Matroids*, sec. ed., Cambridge Univ. Pr., Cambridge, 1999.

2. R. Babilon, 3-colourability of Penrose kite-and-dart tilings, *Disc. Math.* 235 (2001) 137–143.

3. A. Errera, *Du Colorage des Cartes et de Quelques Questions d'Anaysis Situs*, Thesis, Gauthier-Villars, Paris, 1921, 66 pp.

4. S. Felsner, F. Hurtado, M. Noy, and I. Streinu, Hamiltonicity and coloring properties of arrangement graphs, in *Proc. Symp. on Disc. Algorithms*, ACM Press (2000) 655–664.

5. J. P. Hutchinson and S. Wagon, The four-color theorem, *Mathematica in Education and Research* 6:1 (1997) 42–51.

6. J. P. Hutchinson and S. Wagon, Kempe revisited, *Amer. Math. Monthly* 105 (1998) 170-174.

7. I. Kittell, A group of operations on a partially colored map, *Bull. Amer. Math. Soc.* 41 (1935) 407–413.

8. G. Koester, 4-critical, 4-valent planar graphs constructed with crowns, *Math. Scand.* 67 (1990) 15–22.

9. C. A. Morgenstern and H. D. Shapiro, Heuristics for rapidly four-coloring large planar graphs, *Algorithmica* 6 (1991) 869–891.

10. L. Robertson, M. Schweitzer, and S. Wagon, A buttressed octahedron, *Geombinatorics* 9:4, 2002.

11. T. Sibley, Coloring space tilings of generalized parallelograms, *Geombinatorics* 10:2 (2000) 75–76.

12. T. Sibley and S. Wagon, Rhombic Penrose tilings can be 3-colored, *Amer. Math. Monthly* 106 (2000) 251–253.

13. S. Wagon, An April Fool's hoax, *Mathematica in Education and Research* 7:1 (1998) 46–52.