
ARTICLES

Coloring Ordinary Maps, Maps of Empires, and Maps of the Moon¹

JOAN P. HUTCHINSON

Macalester College
St. Paul, MN 55105

1. Introduction to Map Coloring

The following statement is *not* true:

“Every map can have one of four colors assigned to each country so that every pair of countries with a border arc in common receives different colors.”

But isn't this the statement of the famous Four Color Theorem that was proved about 15 years ago using lots of computer checking? Has a flaw been found in that massive piece of work?

FIGURE 1 shows a small example of a map that needs five colors if every pair of adjacent countries is to receive different colors. The important feature is that one country (#5) is a disconnected country of two regions.




FIGURE 1

A planar map needing five colors.

Maps having disconnected countries are certainly possible; an example is the map of North America. This raises the following cartography question: Can today's map of the world be colored with four colors so that all pieces of each country receive the same color and so that no two different countries with a border arc in common receive the same color?

In any case, FIGURE 1 is not a counterexample to the Four Color Theorem; here's what that famous theorem really says [1].

FOUR COLOR THEOREM. *Every map drawn in the plane (or on the surface of the sphere) can have one of four colors assigned to each **connected region** so that every pair of regions with a border arc in common receives different colors.*

¹This article is based on talks given at the Pi Mu Epsilon Conference at St. John's University, Collegeville, MN in March, 1990. This research was supported in part by an NSF Visiting Professorships for Women grant, #RII-8901458, and by the University of Washington, Seattle, WA.

This result was first conjectured in 1852 by Francis Guthrie; a proof was published in 1879 by A. B. Kempe, but, 11 years later, P. J. Heawood found a fatal flaw in that proof. Finally in 1976 the theorem was proved by K. Appel and W. Haken, although in an unusual manner. (A thorough history of the problem can be found in [5].) Briefly, Appel and Haken's proof consists of showing that every planar map contains one of a list of at most 2,000 configurations, and that each configuration admits a reduction, allowing a proof by induction. Although the numbers involved are unusually large, the most unusual (and controversial) part of the proof is that about 1,200 hours of computer time were used to generate the list of 2,000 configurations and to check that colorings on these admit the necessary reduction. Thus the proof depends upon electricity—that is, upon a physical experiment—and that dependence is most unusual in the history of mathematics. Both the theoretical and the computational aspects of the proof have been carefully checked, but mathematicians would very much like to see a simpler or more “natural” proof. Several papers ([2, 12, 28]) contain detailed and substantive analyses of the proof.

The map in FIGURE 1 might also be a representation of political boundaries in which area #5 represents an empire, two countries joined together. In the past, when countries were annexed to form empires, it was the custom of mapmakers to color all parts of an empire with the same color; thus maps of empires might require more than four colors. The mathematics of empire coloring was first studied by Heawood over a hundred years ago. In this paper we survey some of his work and more recent map-coloring variations by G. Ringel. Then one of their results will be applied to a modern problem in the testing of printed circuit boards.

2. Coloring Maps of Empires

When Heawood both found a flaw in Kempe's argument and discovered that he could not repair it, he invented some generalizations about map colorings that, to a certain extent, he could solve. His research began the field of topological graph theory as studied today. First he investigated the empire-coloring problem: If maps are made up of countries united into empires, then how many colors are needed to color such maps provided that all countries in an empire receive the same color and that empires with a border in common receive different colors? Heawood proved that if every empire consists of at most M connected regions (called an M -pire by B. Jackson and G. Ringel [16]), then the map can be colored with at most $6M$ colors.

We shall present a proof of Heawood's upper bound, but first, we ask whether $6M$ is the best possible bound. For $M = 1$, it is not: If empires consist only of single connected regions, then four colors are enough by the Four Color Theorem. But for $M = 2$, Heawood found an example of 12 empires, each consisting of two regions and so forming a 2-pire, such that every pair of empires has a border in common. Thus 12 colors are needed; however, Heawood's example was sufficiently irregular that he was unable to generalize it for $M > 2$.

After an introduction to some facts and techniques of graph theory, we'll prove Heawood's theorem and then summarize recent work of Jackson and Ringel [16] that settles the question of the best possible coloring bound for M -pires with $M > 1$.

We make a change now from maps to graphs as is customary with practitioners in this field. Given a map C drawn in the plane, we form a *planar graph* by creating a vertex for each connected region of C , and by joining two vertices by an edge if the corresponding regions have a border arc in common. This graph is called planar because it can be drawn in the plane without edge crossings. FIGURE 2 shows the

planar graph that results from the map in FIGURE 1; notice that a vertex is created also for the outer, infinite region of the map. And given a planar graph, we could construct a map with countries having the prescribed borders in common.




FIGURE 2
The planar graph derived from the map of FIGURE 1.

Previously we wanted to color regions of a map; now we color vertices of the graph. For k a positive integer, a graph is said to be k -colored (or k -colorable) if each vertex is (or can be) assigned one of k colors so that every pair of vertices that are joined by an edge receives different colors. Why bother with this change from maps to graphs and then coloring vertices rather than regions? This change certainly allows us to use less paint and, more to the point, this convention is better suited to current algorithmic approaches to graph theory problems, allowing for efficient storage and usage of graphs by computer programs.

Here then in the language of graph theory is Heawood's M -pire problem: Prove that every graph that is derived from an M -pire map can be $6M$ -colored so that (in addition) all vertices corresponding to the same empire receive the same color.

Here is the first, crucially important, tool of the trade. A *face* in a planar graph G , drawn in the plane, is a connected region in the plane minus the edges and vertices of G .

EULER'S FORMULA. *If G is a connected planar graph, drawn in the plane with n vertices, e edges, and f faces, then*

$$n - e + f = 2.$$

For a proof, see [3, 7]. An immediate consequence of this formula is that the number of edges in a planar graph is limited. Suppose we count the number of edges bordering each face: Let e_i denote the number of edges on the i th face (counting an edge twice if both sides border on the same face). Then we get

$$2e = e_1 + e_2 + \dots + e_f \geq 3f,$$

since every edge is counted exactly twice in the sum of the e_i , and since each face has at least three bordering edges (provided the graph contains neither loops nor multiple edges). Thus

$$2e/3 \geq f.$$

Applying Euler's Formula,

$$n - e + 2e/3 \geq 2,$$

or

$$e \leq 3n - 6. \quad (1)$$

Similarly, if the i th vertex has \deg_i incident edges, (\deg_i is called the *degree* of the vertex), and we sum the degrees of all vertices, we get

$$2e = \deg_1 + \deg_2 + \cdots + \deg_n$$

since each edge is counted twice, once at each end-vertex. Then, by (1)

$$\deg_1 + \deg_2 + \cdots + \deg_n \leq 6n - 12. \quad (2)$$

We are almost ready to $6M$ -color M -pires. Notice that our goal is to color certain graphs with an extra constraint imposed because the graph comes from an M -pire: Certain vertices are supposed to receive the same color. That's a bit unnatural to a graph colorer. Most graph-coloring theorems and algorithms assume that a graph just should be colored so that no two adjacent vertices receive the same color; this extra constraint is a bother. So we'll get rid of it appropriately.

The idea is a simple one. Recall that we began with a planar M -pire map, and from this formed a planar graph—call it G —with a vertex for each connected region. From G we form a new graph G^* by identifying the set of vertices of G corresponding to an empire into one vertex of G^* and removing any multiple edges. We call G^* the *M -pire graph* associated with an M -pire map. It is the graph G^* that we want to color (with as few colors as possible) since all countries of the same M -pire will necessarily receive the same color and all map adjacencies have been recorded in G^* . But G^* most likely will no longer be a planar graph.




FIGURE 3
 G^* formed from G of FIGURE 2.

We shall prove that G^* can always be $6M$ -colored. Suppose G^* has n^* vertices and e^* edges. Since each vertex of G^* comes from at most M vertices of G , a planar graph with, say, n vertices and e edges, $Mn^* \geq n$, and $e^* \leq e$. Thus in G^* we have from (2)

$$\begin{aligned} \deg_1 + \deg_2 + \cdots + \deg_{n^*} &= 2e^* \leq 2e \leq 6n - 12 \\ &\leq 6Mn^* - 12. \end{aligned}$$

Hence the average degree of vertices in G^* is bounded:

$$\begin{aligned} (\deg_1 + \deg_2 + \cdots + \deg_{n^*})/n^* &\leq 6M - 12/n^* \\ &< 6M, \end{aligned}$$

and so G^* contains a vertex of degree at most $6M - 1$.

THEOREM 1 (Heawood) [15]. *Every M -pire graph G^* can be $6M$ -colored.*

Proof. By induction on n^* . If $n^* \leq 6M$, then G^* can easily be $6M$ -colored by placing a different color on each vertex. Assume the theorem is true for every M -pire graph with fewer than n^* vertices, and let G^* be an M -pire graph with $n^* > 6M$ vertices.

Find a vertex v of G^* of degree at most $6M - 1$; delete v and all its incident edges. The resulting graph is still an M -pire graph (why exactly is this true?) and so, by induction, can be $6M$ -colored. Since v in G^* is adjacent to at most $6M - 1$ different colors, there is a color available to place on v . Thus G^* is $6M$ -colored.

COROLLARY 1. *Every (ordinary) planar graph (or map) can be 6 -colored.*

But are $6M$ colors sometimes necessary (when $M > 1$)? Since 1980, examples have been known of 3-pires and 4-pires that need 18 and 24 colors respectively; colorful pictures of these are presented in [8]. In 1983 B. Jackson and G. Ringel were able to settle the whole question as follows.

M -PIRE THEOREM [16]. *For every $M > 1$ there is an M -pire graph that requires $6M$ colors. In fact, the graph consisting of $6M$ mutually adjacent vertices is an M -pire graph.*

The graph of $6M$ vertices with every pair of vertices joined by an edge is called the *complete graph* on $6M$ vertices and is denoted by K_{6M} . K_{6M} is $6M$ -colorable, but cannot be colored with fewer colors. The M -pire theorem can be restated as follows: There is a planar map (or graph) that consists of $6M$ empires, each with at most M regions, such that every empire has a border in common with every other. Jackson and Ringel proved this theorem by constructing symmetrical maps using the theory of "current graphs" as developed by Ringel and J. W. T. Youngs for a solution of another problem invented by Heawood (this problem will be discussed in Section 5). Thus the problem of coloring M -pire maps has been completely and beautifully settled when $M > 1$. A thorough exposition of this subject is given in [14].

3. Coloring Maps of the Moon

Ringel suggested the following variation on the empire coloring problem. Suppose the Moon were colonized and we want to color a map of the Earth and the Moon so that

1. adjacent regions on Earth or on the Moon receive different colors, and
2. a country on Earth and its lunar colony receive the same color.

Suppose now that there are *no empires* on the Earth or Moon; each country and colony is a connected region.

But this is just a 2-pire problem. Let G_e and G_m be the corresponding planar graphs of the Earth and Moon maps; side-by-side they form one planar graph that comes from a 2-pire map (see FIGURE 4). When each vertex of G_m is identified with its country-vertex in G_e , the resulting graph G^* is a 2-pire graph and so can be 12-colored by Theorem 1.




FIGURE 4
 K_8 as an Earth/Moon graph.

A graph G^* is said to be an *Earth/Moon graph* (or to have *thickness 2*) if it can be divided into two planar graphs by making two copies of the vertex set of G^* , and then assigning each edge of G^* to one of the two copies so that two planar graphs result. Think of one planar graph as representing the terrestrial countries and the other the lunar colonies. More generally, a graph is said to have *thickness t* if its edges can be partitioned into t , but not fewer, planar graphs. This thickness parameter is studied by graph theorists both for its relevance to general graph theory and for applications in computer science. In the next section, we'll consider an application of thickness-2 graphs.

So what's new in Ringel's question about coloring Earth/Moon graphs if we already know they can be 12-colored? Again the important point is whether or not 12 colors are needed. Notice that the definitions of an Earth/Moon graph and of a 2-pire graph are not the same, and consequently a 2-pire graph might not be an Earth/Moon graph. If G^* is a 2-pire graph, some vertices can be split into two so that a planar graph, say, G results. But it may not be possible to divide G into two separate planar graphs, each containing one copy of each vertex, as is required for an Earth/Moon graph. For example, the complete graphs K_9 , K_{10} , K_{11} , and K_{12} are all 2-pire graphs, but it is known that none of them is an Earth/Moon graph. They all have thickness 3 because of the following result. (For a summary of the several proofs involved, see [4].)

THEOREM (Beineke, Harary, Alekseev, Gonchakov, Vasak, Mayer). *The thickness of the complete graph K_n is given by*

$$t(K_n) = \begin{cases} \lfloor (n + 7)/6 \rfloor, & \text{if } n \neq 9, 10, \\ 3, & \text{if } n = 9, 10. \end{cases}$$

FIGURE 4 shows how K_8 arises from two planar subgraphs. K_9 cannot be so created, but if the Earth had two moons, then it would be possible to find maps on the three spheres that would unite to form a K_9 . Thus Earth/Moon graphs will need at least 8 colors, but since the graph K_{12} has thickness 3, we can't conclude that 12 colors are needed for Earth/Moon graphs.

The least number of colors that we need to have on hand so that we can color every Earth/Moon graph is called the *chromatic number* of Earth/Moon graphs (and of thickness-2 graphs). As reported in [8], T. Sulanke found a map of the Earth and Moon that needs 9 colors, although it does not come from K_9 ; FIGURE 5 shows a schematic version of the corresponding Earth/Moon graph. It is a worthwhile exercise to check that this graph can be 9-colored, but not 8-colored, and that its edges can be separated into two planar subgraphs.




FIGURE 5
A 9-chromatic Earth/Moon graph.

Open Problem. What is the chromatic number of Earth/Moon graphs? 9? 10? 11? 12?

Progress would be made if any of these possibilities could be eliminated. In general the chromatic number of thickness- t graphs is not known exactly. By the preceding theorem and Theorem 1, the chromatic number is one of $6t - 2$, $6t - 1$, or $6t$ for $t > 2$.

4. An Application of Earth/Moon Coloring to the Testing of Printed Circuit Boards

We turn to an application that, perhaps surprisingly, uses the coloring results of the previous section to develop an efficient procedure for the testing for certain errors in printed circuit boards. This example comes from three researchers at AT&T Bell Laboratories in Murray Hill and Whippany, NJ [9].

It is not hard to think of potential applications of the thickness parameter to electronics. Suppose you have a system of electrical units, certain pairs of which should be connected electrically. Such a design has a graph naturally associated with it. A goal could be to develop a layout for that graph with, say, electrical wires or solder joining the units. To avoid crossings and unwanted electrical connections, the joining should be done on multiple planar layers. Or, as in the design and fabrication of VLSI chips, connections should be etched onto separate layers of silicon.

For this application we focus on printed circuit boards (PCBs), small electrical circuits that are widely used in motorized, electrical, and electronic equipment. (For examples and illustrations, see [24].) We postulate a simple, but realistic, mathematical model for these circuit boards. We assume that a PCB consists of electrical units

placed on a rectangular array with some of the units joined electrically along horizontal and vertical lines. Typically an array consists of a 100×100 grid with about 500 electrically connected components (called *nets*). And each net has a simple structure: Every pair of units in one net is joined by a unique conductor path; this structure is known to graph theorists as a *tree*. An example is shown in FIGURE 6.

We focus on one specific problem, that of finding certain extraneous and unwanted connections that may have mistakenly occurred during the manufacturing process. In practice a design is set for a PCB; thousands or even millions of these inexpensive PCBs are made; and as they roll off the production line, a quick and accurate test is needed to see whether there are errors on the board. If so, the board is discarded; no attempt is made to find and correct the error. (As an example of what happens when quality control is not effective, a front-page article in the *New York Times* [20] reported recently on errors in a \$12 circuit board that caused a delay in manufacturing 2,500 Air Force rockets, each worth about a million dollars.)

Here is a mathematical statement of the problem we consider.

Problem. Determine whether any extraneous horizontal or vertical conductor paths have been introduced in the manufacturing process, connecting two nets that should not be electrically connected (such a connection is known as a *short circuit*).

In practice, erroneous horizontal and vertical paths are the most common short circuit, and so this problem focuses on such mistakes.

There is a simple answer to the problem: Check all pairs of nets to see if any is incorrectly connected. If there are n nets, such a check requires $n(n-1)/2$ tests and is by far too slow a procedure. For example, with 500 nets, about 125,000 tests would be needed. The better solution presented in [9] necessitates only 11 checks per board, regardless of the number of nets. We also show how to reduce the checking to a mere four per board.

Not surprisingly, we make a graph from the intended PCB. The *PCB graph* is defined to have

1. a vertex for each net of the (correct) PCB, and
2. two vertices joined by an edge if there is a horizontal or vertical line going through the corresponding two nets and passing through no intermediate net.

FIGURE 7 shows the PCB graph for the example of FIGURE 6—for the moment, ignore the labels on the vertices. In this figure the vertices are positioned to correspond to the position of the nets in the PCB of FIGURE 6.

Perhaps a better name for this graph would be a “graph of possible mistakes.” Imagine that in the fabrication process two passes are made over the PCB, one pass creating all horizontal connections and the second creating all vertical. The problem of concern is that too much connecting might be done; the fabrication process might not shut off correctly and so might connect more than is required. Thus in the PCB graph two vertices are joined if it is possible for the fabrication machine to mistakenly join the corresponding nets by a direct horizontal or vertical connection. But what about the condition of passing through no intermediate net? Look at FIGURE 6: Suppose nets x and y were mistakenly connected by a horizontal line. Then the connection would necessarily connect x with z and z with y . In other words, if we check that x and z are not connected and that z and y are also not connected, then we know for sure that x and y are not mistakenly connected. Hence there is no need for an x -to- y edge in the PCB graph.

The key observation now is that the PCB graph has thickness 2. Its edges can be divided into two planar graphs, one with all edges corresponding to possible vertical connections and one with edges from horizontal connections. That observation and

the fact that a thickness-2 graph can be 12-colored leads to the following checking algorithm.

PCB-checking Algorithm. Given a plan for a PCB and the corresponding PCB graph G :

1. 12-color G , and transfer this coloring to the nets of the plan.
2. From the plan for the PCB, construct 12 external conductor tree-structures, called "supernets," so that when each supernet is attached to the PCB, all nets in the same color class become electrically connected.
3. Check all pairs of these supernets to see if any two of these are (mistakenly) electrically connected.

Here are a few more words of explanation. Consider all nets that receive one color, say, color #1. Since no two of the corresponding vertices in G are joined by an edge, no two of these nets need to be tested for a mistaken connection. A "supernet," such as those shown in FIGURE 8, is some simple, 3-dimensional electrical connection (shaped perhaps like an octopus or an n -pus) that attaches to the circuit board and electrically connects a set of nets.




FIGURE 6
A printed circuit board model.




FIGURE 7
The PCB graph from FIGURE 6 with its vertices 8-colored.

In step (3) an electrical probe can check each pair of supernets. If, say, supernets #1 and #2 are electrically connected, then some net colored 1 and some net colored 2 are mistakenly connected (and we throw the board out). But if not, then no net of color 1 is connected to a net of color 2, and we proceed to test other pairs of supernets. With 12 supernets, there are $12 \cdot 11/2 = 66$ pairs that might need to be tested. Granted, the creating of these supernets takes some time and money, but recall that we are imagining that a machine (or several) is creating thousands of PCBs following one master plan so that the checking supernets can be used repeatedly as the PCBs roll off the line.

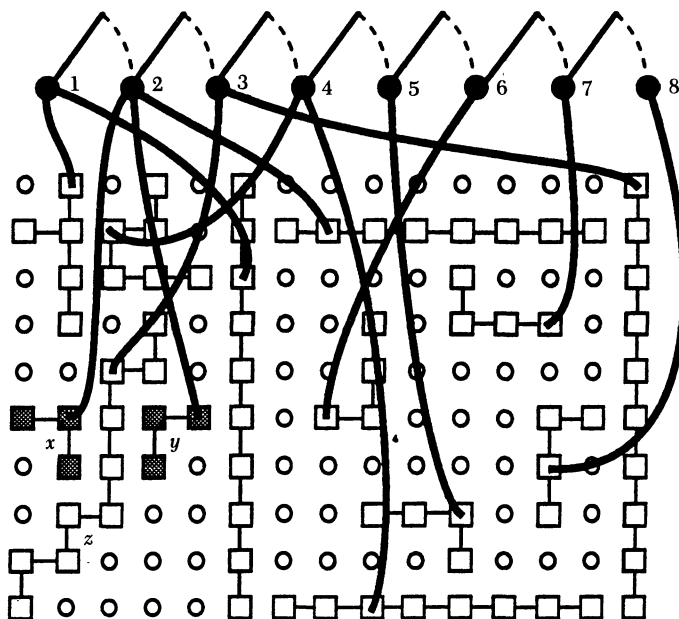


FIGURE 8

The printed circuit board of FIGURE 6 with supernets attached.

With another gadget we can reduce the number of tests even more. Suppose we test supernet #1 versus #2. If they are electrically connected, we throw the board out, but if not, we add a “gate”, some simple electrical connection to make these two supernets electrically connected. Next we test this united #1–2 supernet versus supernet #3. If there is an electrical connection, then a net in color class 3 is mistakenly connected to one of color class 1 or 2. We don’t care which; we just throw the board out. But if not, we connect supernet #3 up with #1 and #2. Continuing with this checking and connecting, we see that, with a correct board, after only 11 connections of supernets and 11 tests, we detect that we have a board free of the kinds of mistakes for which we were checking. FIGURE 7 shows the PCB graph from FIGURE 6 together with an 8-coloring of its vertices. Then FIGURE 8 shows the corresponding supernets and the connecting devices to reduce the checking to 11 steps.

Allen Schwenk (personal communication) has recently pointed out how to further reduce the number of checks to four, with additional gadgetry. Take the existing supernets, numbered 1, 2, . . . , 12, and think of these numbers, expressed in binary, each with exactly four binary digits (called bits). Make a supersupernet that connects the supernets labeled with numbers beginning with a 0 bit, and make a similar supersupernet that connects the supernets labeled with numbers beginning with a 1

bit. For the first test, check if these two supersupernets are electrically connected. If not, create two supersupernets, one connecting the supernets with a 0 in the second bit and one connecting the supernet with a 1 there. For the second test, check if these two are connected. Do the same creation and test for the third and fourth bits. If there is an erroneous connection, we will detect it: If there is a connection between a net colored i and a net colored j , the binary representation of i and j differ in some bit, and thus electricity would flow when the two supersupernets for that bit were tested. Notice that in these tests we have not detected all possible errors, such as those when too few connections are made within a net or when some extra zigzag-like connections are made between nets, but this approach does solve the problem of concern to the AT&T researchers.

5. Coloring Ordinary Maps and Empire Maps on Surfaces

Now we return to the realm of map coloring. Recall the first example of a map that needed more than four colors in FIGURE 1. There's another way in which maps might need more colors, and Heawood thought of this one too. Suppose a map were drawn on the surface of a torus (a donut) or a 4-holed torus (like a pretzel). Then the map might need as many as 7 or 10 colors, respectively. Here's some motivation for such topological map drawing and coloring.

Thinking now in terms of graphs, notice that a graph can be drawn in the plane (without edge crossings) if and only if it can be so drawn on the surface of a sphere. But what about graphs that can be drawn in neither place, such as K_5 or any K_n with $n \geq 5$? One trick would be to add bridges or handles to the plane or sphere so that edges can traverse these and so avoid edge crossings. The sphere plus a handle or plus g handles ($g > 0$) is essentially the same as the torus or the g -holed torus, respectively; topologists call this essential similarity a homeomorphism. For our purposes, we use the fact that a graph can be *embedded* (i.e., drawn without edge crossings) in the plane plus g bridges if, and only if, it can be embedded on the g -holed torus, also known as the surface of genus g . The *genus* of a graph is defined to be the least g such that the graph can be embedded on the sphere plus g handles.

For example, it is not hard to see that K_5 embeds on the torus; but with a little more effort one can also embed K_6 and K_7 there. So in particular, graphs that embed on the torus may need as many as 7 colors. FIGURE 9 shows the corresponding map situation, a map of seven mutually adjacent regions. Heawood discovered such a map and proved that every graph on the torus can be 7-colored.




FIGURE 9

Seven mutually adjacent regions on the torus.

He also proved a related result for each surface of positive genus. Let $\chi(g)$ equal the minimum number of colors needed to color every graph that embeds on the sphere plus g handles, $g > 0$.

THEOREM 2 (Heawood) [15]. For $g > 0$, $\chi(g) \leq \lfloor (1/2)(7 + \sqrt{48g + 1}) \rfloor$.

So, for example $\chi(1) \leq 7$, $\chi(2) \leq 8$, $\chi(6) \leq \chi(7) \leq 12$, and $\chi(17) \leq 17$. (Note that substituting $g = 0$ yields $\chi(0) \leq 4$, but Heawood's proof does not cover this case.) A good introduction to the mathematics and history of this problem is found in [25].

Why is Heawood's theorem true? Soon (in Theorem 3) we'll give a proof of this and a more general result about embeddings on surfaces that depends upon a generalization of Euler's Formula to surfaces of positive genus, but first, let's see where Heawood's bound comes from.

EULER-POINCARÉ FORMULA. If G is a graph, embedded on the g -holed torus with n vertices, e edges, and f faces, then

$$n - e + f \geq 2 - 2g.$$

If every face is a contractible region, then

$$n - e + f = 2 - 2g.$$

For a proof, see [11, 29]. FIGURE 10 shows three embeddings of K_4 on the torus; only the last has the nice property that every face is contractible (or homeomorphic to a planar disk). Another way to think of this property is that equality occurs in the Euler-Poincaré Formula when the embedding uses all the handles and is not really embedded on a sphere with fewer handles.




FIGURE 10
Embeddings of K_4 on the torus.

Using the Euler-Poincaré Formula and arguing just as we did for inequality (1) in section 2, we have in all cases that

$$e \leq 3n + 6(g - 1).$$

How large a complete graph K_n could we embed on the sphere plus g handles? Such a graph has $n(n - 1)/2$ edges and so it must be the case that

$$n(n - 1)/2 \leq 3n + 6(g - 1). \quad (3)$$

Thus

$$n^2 - 7n + 12(1 - g) \leq 0,$$

and by the quadratic formula

$$n \leq (1/2)(7 + \sqrt{49 + 48(g - 1)})$$

so that

$$n \leq \lfloor (1/2)(7 + \sqrt{48g + 1}) \rfloor.$$

Note that from (3) it also follows that

$$g \geq \lceil (1/12)(n - 3)(n - 4) \rceil. \tag{4}$$

But what is the genus of K_n ? Is it the lower bound of line (4)? Equivalently, when $n = \lceil (1/2)(7 + \sqrt{48g + 1}) \rceil$, can K_n be embedded on the sphere plus g handles? If so, then $\lceil (1/2)(7 + \sqrt{48g + 1}) \rceil$ colors are needed for graphs that embed there. Heawood blithely assumed that since the answers to the preceding two questions were YES for the torus, $n = 7$ and $g = 1$, they must be YES for all larger n and g . His intuition was correct, but there was no proof of these facts until 1968 when the following deep and difficult result was proved principally by Ringel and Youngs, with help from others on a few cases.

MAP-COLOR THEOREM (Ringel, Youngs, Gustin, Guy, Mayer, Terry, Welch) [23]. *The genus of K_n is $\lceil (1/12)(n - 3)(n - 4) \rceil$, and consequently for $g > 0$, $\chi(g) = \lceil (1/2)(7 + \sqrt{48g + 1}) \rceil$.*

Thus the equivalent of the Four Color Theorem was solved first for all surfaces except for the sphere (and without computer help).

But Heawood imagined even more, combining the ideas of maps on surfaces and empires. Suppose an M -pire map is drawn on the sphere plus g handles, and suppose G is the corresponding graph embedded on the same surface. Let the minimum number of colors needed for all such graphs be denoted by $\chi(g, M)$ where we require that all vertices coming from the same empire receive the same color.

THEOREM 3 [15]. *For all $g \geq 0$ and $M \geq 1$, except for the case $g = 0$ and $M = 1$,*

$$\chi(g, M) \leq \left\lceil \frac{6M + 1 + \sqrt{48g + (6M + 1)^2 - 48}}{2} \right\rceil.$$

As in the case of $M = 1$ and $g > 0$, this formula can be motivated as above by supposing that K_n is an M -pire graph that arises from vertex-identifications of a graph embedded on the sphere plus g handles.

Notice first that for $g = 0$ and $M \geq 2$,

$$6M \leq \frac{6M + 1 + \sqrt{48g + (6M + 1)^2 - 48}}{2} < 6M + 1,$$

and so the upper bound in these cases of Theorem 3 is $6M$. We have seen a proof that $\chi(0, M) \leq 6M$ in Theorem 1, and by the M -pire Theorem of Section 2, $\chi(0, M) = 6M$ for $M \geq 2$. Notice also that for $g > 0$ and $M = 1$, Theorem 3 coincides with that of Theorem 2, whose proof we haven't yet seen, and by the Map-Color Theorem $\chi(g, 1)$ actually equals this upper bound when $g > 0$. Now we'll prove Theorem 3 in general, and then summarize what is known when $\chi(g, M)$ equals this upper bound. (Notice the parallels with the proof of Theorem 1 although, as we'll point out, this argument fails for $g = 0$.)

Proof of Theorem 3. The case of $g = 0$ was proved in Theorem 1, and so we assume that $g > 0$.

If G is an n -vertex, e -edged graph embedded on the sphere plus $g > 0$ handles that arises from an M -pire map, we may identify vertices from each empire to obtain a graph G^* that should be (normally) colored and that may no longer embed on the same surface. If G^* has n^* vertices and e^* edges, then $n \leq Mn^*$, $e^* \leq e$, and by the Euler-Poincaré Formula

$$e^* \leq e \leq 3n + 6(g - 1) \leq 3Mn^* + 6(g - 1).$$

The proof now proceeds by induction on n^* . If

$$n^* < \frac{6M + 1 + \sqrt{48g + (6M + 1)^2 - 48}}{2},$$

the result is clearly true; so suppose

$$n^* \geq \frac{6M + 1 + \sqrt{48g + (6M + 1)^2 - 48}}{2}.$$

$$\text{Then } \frac{2e^*}{n^*} \leq 6M + \frac{12(g-1)}{n^*} \leq 6M + \frac{24(g-1)}{6M + 1 + \sqrt{48g + (6M + 1)^2 - 48}}.$$

(Notice that the second inequality holds only for $g \geq 1$.) Clearing the radical in the denominator of the previous fraction, one obtains:

$$\begin{aligned} \frac{2e^*}{n^*} &\leq 6M + \frac{24(g-1)\{6M + 1 - \sqrt{48g + (6M + 1)^2 - 48}\}}{-48(g-1)} \\ &= 6M + \frac{-(6M + 1) + \sqrt{48g + (6M + 1)^2 - 48}}{2} \\ &= \frac{6M - 1 + \sqrt{48g + (6M + 1)^2 - 48}}{2}. \end{aligned}$$

Thus there is a vertex v of degree at most

$$\begin{aligned} &\left\lfloor \frac{6M - 1 + \sqrt{48g + (6M + 1)^2 - 48}}{2} \right\rfloor \\ &= \left\lfloor \frac{6M + 1 + \sqrt{48g + (6M + 1)^2 - 48}}{2} \right\rfloor - 1. \end{aligned}$$

Remove v , use induction to color the remaining vertices with

$$\left\lfloor \frac{6M + 1 + \sqrt{48g + (6M + 1)^2 - 48}}{2} \right\rfloor$$

colors, and extend this coloring to v since it has one fewer adjacent vertices than the number of colors being used.

This argument with $M = 1$ (the nonimperialist case) gives a complete proof of Heawood's bound (in Theorem 2) for coloring ordinary maps on surfaces. Good expositions of this and related aspects of topological graph theory can be found in [3, 14, 25], but, surprisingly, almost no introductory text on graph theory except for the recent book [14] contains a proof of this ($M = 1$) bound (although one text contains an incorrect proof). Proofs of this nonimperialist case can be found in more specialized texts [6, 11, 22, 29], while the fully general case of Theorem 3 appears in [21].

Is the upper bound of Theorem 3 always achieved by some M -pire graph? Notice that for the torus, where $g = 1$, the upper bound is simply $6M + 1$. H. Taylor [26] has announced that for each M there are M -pire graphs on the torus, achieving the bound of $6M + 1$; his work is based on results of S. W. Golomb [10] on "graceful" labelings of graphs. And for the fully general case, Jackson and Ringel [18] have studied the situation intensively and have proved that the upper bound is achieved for at least 12.5% of the remaining cases.

THEOREM [18]. For $g > 0$,

$$\chi(g, M) = \left\lfloor \frac{6M + 1 + \sqrt{48g + (6M + 1)^2 - 48}}{2} \right\rfloor \quad (5)$$

- (1) when M is even and the right-hand side of (5) is congruent to 1 modulo 12, and
 (2) when M is odd and the right-hand side of (5) is congruent to 4 or 7 modulo 12.

It is also possible to consider empire maps on nonorientable surfaces such as the projective plane, the Klein bottle, or the sphere plus k cross-caps, $k > 0$. Then a formula analogous to that of Theorem 3 can be proved similarly, but again the hardest work is showing that equality can be achieved. Specific results are known, and in general the upper bound has been shown to be achieved in about 25% of the cases; a summary of these results is contained in [18].

Heawood thought up a variety of problems and questions that have intrigued researchers for years. Others since Heawood have embellished upon his ideas, some in equally fanciful map terms (e.g., [17]), others with more abstraction [13], and several with applications in theoretical computer science. For example, the concept of thickness has considerable applicability, beyond that of PCBs, in the area of complexity of algorithmic problems and in the theory of NP-completeness [19]. General graph-coloring questions and related algorithmic problems, not just in the context of maps, are some of the most widely studied today, because of both their difficulty and their applicability (see [27]).

Acknowledgements. I would like to thank Nora Hartsfield, Brad Jackson, Allen Schwenk, and Arthur White for helpful comments on this paper, and Stan Wagon for producing the diagrams.

REFERENCES

1. K. Appel and W. Haken, Every planar map is four colorable, *Bulletin Amer. Math. Soc.* 82 (1976), 711–712.
2. K. Appel and W. Haken, A proof of the four color theorem, *Discrete Math.* 16 (1976), 179–180.
3. L. W. Beineke, Are graphs finally surfacing?, *College Math. J.* 20 (1989), 206–225.
4. L. W. Beineke and A. T. White, Topological graph theory, *Selected Topics in Graph Theory*, L. W. Beineke and R. J. Wilson (eds.), Academic Press, London, 1978, 15–50.
5. N. L. Biggs, E. K. Lloyd, and R. J. Wilson, *Graph Theory, 1736–1936*, Clarendon Press, Oxford, 1976.
6. B. Bollobás, *Extremal Graph Theory*, Academic Press, London, 1978.
7. G. Chartrand, and L. Lesniak-Foster, *Graphs and Digraphs* (2nd edition), Wadsworth and Brooks/Cole, Publishing Co., Belmont, CA, 1986.
8. M. Gardner, Mathematical Games, *Scientific American* 242 (February 1980), 14–19.
9. M. R. Garey, D. S. Johnson, and H. C. So, An application of graph coloring to printed circuit testing, *IEEE Trans. Circuits and Systems* CAS-23 (1976), 591–599.
10. S. W. Golomb, How to number a graph, *Graph Theory and Computing*, R. C. Read (ed.), Academic Press, New York, 1972.
11. J. L. Gross and T. W. Tucker, *Topological Graph Theory*, John Wiley & Sons, Inc., New York, 1987.
12. W. Haken, An attempt to understand the four color problem, *J. Graph Theory* 1 (1977), 193–206.
13. N. Hartsfield, B. Jackson, and G. Ringel, The splitting number of the complete graph, *Graphs and Combinatorics* 1 (1985), 311–329.
14. N. Hartsfield and G. Ringel, *Pearls of Graph Theory*, Academic Press, New York, 1990.
15. P. J. Heawood, Map colour theorem, *Quart. J. Pure and Applied Math.* 24 (1890), 332–333.
16. B. Jackson and G. Ringel, Solution of Heawood's empire problem in the plane, *J. Reine Angew. Math.* 347 (1983), 146–153.
17. _____, Coloring island maps, *Bull. Austral. Math. Soc.* 29 (1984), 151–165.
18. _____, Heawood's empire problem, *J. Combin. Theory Ser. B* 38 (1985), 168–178.
19. A. Mansfield, Determining the thickness of graphs is NP-hard, *Math. Proc. Camb. Phil. Soc.* 93 (1983), 9–23.

20. *The New York Times*, Sept. 15, 1989, p. 1.
 21. G. Ringel, *Färbungsprobleme auf Flächen und Graphen*, VEB Deutscher Verlag der Wissenschaften, Berlin, 1959.
 22. G. Ringel, *Map Color Theorem*, Springer-Verlag, New York, 1974.
 23. G. Ringel and J. W. T. Youngs, Solution of the Heawood map-color problem, *Proc. Nat. Acad. Sci. USA* 60 (1968), 438–445.
 24. *Scientific American* 237 (Sept. 1977).
 25. S. Stahl, The other map coloring theorem, this *MAGAZINE*, 58 (1985), 131–145.
 26. H. Taylor, Synchronization patterns and related problems in combinatorial analysis and graph theory, *University of Southern California Electronic Science Laboratory Tech. Report #509*, June, 1981.
 27. B. Toft, 75 Graph Colouring Problems, *Graph Colourings*, R. Nelson and R. J. Wilson, eds., Longman Scientific & Technical, Harlow, UK, 1990.
 28. T. Tymoczko, The four-color problem and its philosophical significance, *J. Philosophy* 76 (1979), 57–83.
 29. A. T. White, *Graphs, Groups and Surfaces* (revised edition), North-Holland, Amsterdam, 1984.
-

Fractal Basin Street Blues

MAURICE MACHOVER

St. John's University
Jamaica, NY 11439

There was a young fractal named Fracta,
Who certainly knew how to factor;
She tried out ceramics but switched to dynamics,
And married a strange attractor.

