BBC Television has a very popular programme called The Multicoloured Swap Shop. *Below,* MANIFOLD *presents the Multicoloured Theorem Shop, running the gamut of the integers from 1 to 5...*

the 1-colour Theorem

JOZEF PLOJHAR

The little-known one-colour theorem is due to the persistence of a long-since forgotten cartographer of about 3 years of age, who like all children of such an age covered his maps in a single wash of colour. His father immediately realised the significance of this, and burst into the mathematical journals with:

THEOREM *All coloured maps are coloured with a single colour.*

Before reprinting the proof, a comment is surely due on the power of this theorem - unlike later chromatic theorems, it does not assert that the map *may* be so coloured, but rather, that it IS!

Proof (By Mathematical Induction). Let P(n) be the proposition "all maps with n regions are 1-coloured".

P(1) is trivially true. We show that P(n) implies P(n+1). Consider a map with n+1 regions. Remove one region: by induction the resulting map is 1-coloured, with colour C, say.

Again consider the map with n+1 regions, but now remove a different single region. The remaining n are 1-coloured, with colour K, say. But K = C as there are some regions which have been 1-coloured twice (if you see what we mean!).

Hence all n regions are coloured with colour C, and P(n+1) is true. Hence, by induction, P(n) is true for all n - and all maps are one-coloured!

the 2-colour Theorem

VIVIENNE HATHAWAY

The infamous 4-colour problem asks, as you are no doubt aware, whether any map on the plane can be coloured using 4 colours so that no two adjacent regions have the same colour. It is not our purpose to go into this question here: we set our sights lower and aim at a theorem about *two* colours. Until MANIFOLD produces a colour supplement, this is more appropriate!

The theorem occurred to me some years ago, but subsequent delving into the literature revealed that it is well known. The result is this:

Suppose a finite number of circles is drawn on the plane. Then the resulting map can be coloured with two colours so that adjacent regions have distinct colours.

The theorem generalizes to closed curves rather than circles:

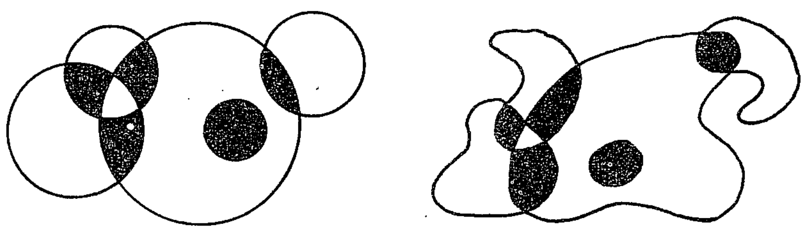

How do we go about *proving* such a theorem? If we *try* colouring a given circle-map, it becomes clear that as soon as one region is coloured the rest follow automatically. A proof based on this would have to show that all regions are reached, and that there are no contradictory choices of colour. This boils down to considering circuits of regions (regions each touching the next along an edge, starting and finishing at a given region). Only if all such circuits contain an even number of regions will the method work. They do, and it does, but that's not the best way to prove the theorem!

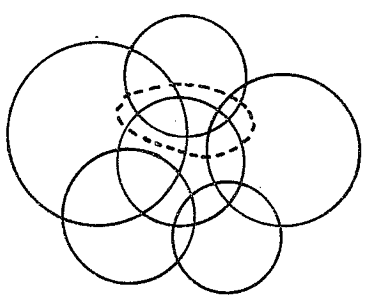

a circuit of 8 regions. Is the number always even?

It occurred to me that the theorem is accessible by Mathematical Induction on the number n of circles. If n = 1 it is easy to colour the map:

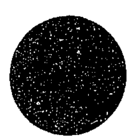

So now we assume we can two-colour any map with n circles, and try to prove that we can two-colour a map with n+1. Now any n+1 circle map comes by adding a circle to an n-circle map. The diagram on the next page is typical. If we can work out how to perform ? in general, we can prove the theorem. Inspection of the diagram reveals that:

(a) Each region outside the new circle retains its colour,

(b) Each region inside the new circle changes colour.

To see that this works in general, note that the colours

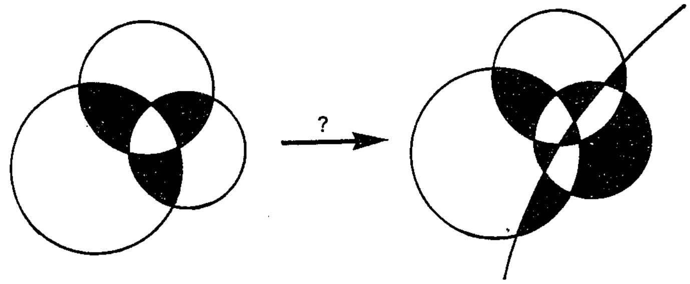

obviously change across boundary-lines outside or inside the new circle (as they did on the old map: though inside, the colours are reversed). And across the new circle's boundary, what was once a single-coloured region divides into two, of opposite colours. QED!

This is all very well - though not a very practical way to perform the colouring - try it - but by analysing the proof, we can find something better. Every time we add a new circle, *points inside it change colour*. So points inside an odd number of circles end up black; points inside an even number (or none) white. This gives us a rule: *assign to each region an integer, equal to the number of circles that contain it. If this number is even, colour the region white; if odd, colour it black.* It is obvious that this number changes by 1 from a region to the next: this gives an independent proof. Here's an example of the rule in operation:

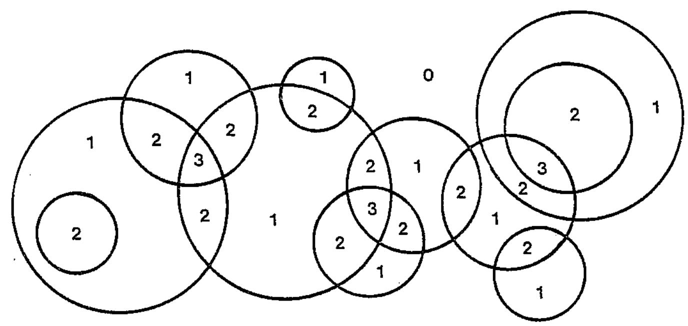

The proof obviously generalizes - e.g. to convex curves rather than circles. If a moral to the tale is needed, it is presumably that our first ideas of how to solve a problem, based on direct solutions of special cases, may not be the best way to proceed in general; and that analysis of a successful method can lead to improvements.

M-9

the $[(7+\sqrt{1+48p})/2]$-colour Theorem

Sorry about that. That's what the *Heawood Conjecture* suggests as the *precise* bound on the number of colours needed for maps on a surface of genus $p \geq 1$ (a torus with p holes, or its non-orientable analogue). It was proved by Ringel and Youngs in 1968.

the 3-colour Theorem

Now, that's a puzzle! MANIFOLD-12 set it as a competition: *Find the 3-colour theorem*. Although there *are* standard 3-colour theorems in graph theory, they all have rather artificial hypotheses. We're still waiting...

the 5-colour Theorem

was until very recently the best that was known towards that doyen of mathematical intractability, the 4-colour *problem*. Which could have caused us headaches with headlines (in contrast to our customary neckaches with necklines...). Fortunately, two Illinois mathematicians, assisted by a whacking great computer, arrived in the nick of time with:

the 4-colour Theorem

M-19

DOUGLAS WOODALL

In July of 1976, K.Appel and W.Haken, two mathematicians at the University of Illinois in America, announced the solution of what was probably the best-known unsolved problem in the whole of mathematics: the four-colour map problem. This asks whether the regions of a map can always be coloured with four colours in such a way that no two neighbouring regions have the same colour. (Neighbouring here means 'having a length of common border'. We do not insist on giving two regions different colours if they meet only at a finite number of points, like regions D and F in Fig.1.)

Fig.1

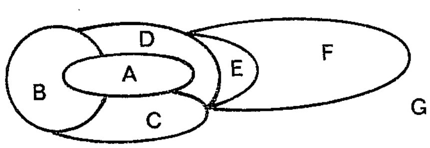

This problem was first proposed in 1852 by a London student, Francis Guthrie, who is reported to have thought of it while colouring a map of the counties of England. He noticed that four colours are sometimes needed (e.g. for regions A,B,C, and D in Fig.1) and conjectured that four colours always suffice, but was unable to prove this. The first serious attempt at a proof seems to have been made in 1879 by A.B.Kempe, a barrister and keen amateur mathematician who later became President of the London Mathematical Society. In that year he published a 'proof' in the American Journal of Mathematics which seems to have been generally accepted. But in 1890 P.J.Heawood, Professor of Mathematics at Durham, pointed out that the 'proof' contained a flaw. For some years after that the flaw seems not to have been regarded as serious, and the theorem was thought to be 'essentially proved'. However, as the years went by and nobody found a satisfactory way round the difficulty, it gradually became realised that the problem was much deeper than had been supposed. Since then, almost every mathematician of repute has probably dabbled with the problem at some time or other, so Appel and Haken's achievement in solving it (in the affirmative) is a very fine one.

As might be expected of such a refractory problem, the proof is long. It runs to 100 pages of summary, 100 pages of detail, and a further 700 pages of back-up work, plus about 1500 hours of computer time. (For comparison, the average proof presented in first year lectures probably does not last more than one page. In the published literature I would regard a 20-page proof as quite long.)

Preparatory Moves

In common with most recent workers, Appel and Haken tackled the problem in the form 'show that the vertices of every planar graph can be coloured with four colours so that no two adjacent vertices have the same colour'. A *planar graph* is a graph (= network) drawn in the plane without edges crossing: see Fig.2. It is easy to show that this version is equivalent to the original map problem (stick a vertex in the middle of each region of the map, and join vertices whose corresponding regions are adjacent). It is also easy to show that it suffices to consider *plane triangulations*, i.e. graphs that divide the plane into regions bordered by exactly three edges (can you see why?). Fig.2 shows the graph corresponding to the map in Fig.1, and the same graph made into a triangulation.

Fig.2.

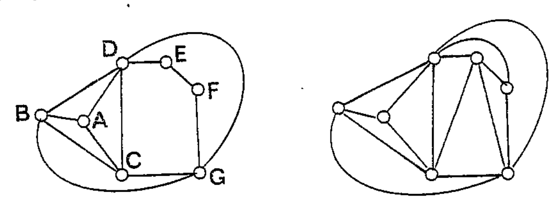

Kempe's "Proof"

In order to understand Appel and Haken's proof, it will be helpful to start by translating Kempe's attempted proof into the language of plane triangulations. Kempe started with Euler's polyhedron formula, which states that a plane triangulation T satisfies the relation $V-E+F = 2$, where V,E,F are the number of vertices, edges, and faces (regions) of T. (Can you prove this?) Since every face of a triangulation is bordered by three edges, and every edge borders two faces (the "outside" is thought of as one huge face), we must have $2E = 3F$ (why?). If V_i denotes the number of vertices of valency i (the *valency* of a vertex is the number of edges incident with it) then clearly $\Sigma V_i = V$ and $\Sigma iV_i = 2E$ (every edge has 2 ends). Substituting these in Euler's formula now gives

$$\sum(6-i)V_i = 12 \qquad (1)$$

or, more longwindedly,

$$4V_2+3V_3+2V_4+V_5-V_7-2V_8-3V_9-\ldots = 12.$$

It follows immediately that at least one of V_2, V_3, V_4, and V_5 is positive; so T must contain at least one of the four configurations in Fig. 3.

Fig.3

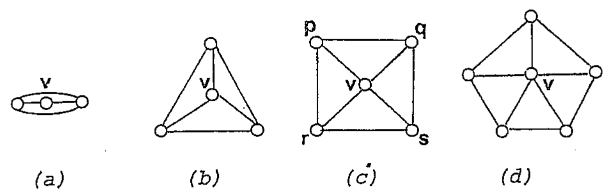

Now suppose there exists a counterexample to the 4-colour conjecture, and let T be a triangulation that is a minimal counterexample, so that every graph with fewer vertices than T *is* 4-colourable, but T itself is *not*. We naturally hope to prove this is impossible by obtaining a contradiction.

If T contains Fig.3(a) or 3(b), we need only remove v from T (together with the incident edges), 4-colour what is left, and restore v: since v is adjacent to at most 3 vertices, we can find a colour for it. Thus we have 4-coloured T, a contradiction. So T cannot, in fact, contain 3(a) or 3(b).

For 3(c) we try the same thing, but this time we are in trouble if p,q,r, and s all have different colours; in this case we cannot colour v. However, Kempe ingeniously showed, using what is now called a *Kempe-chain argument*, that here we can modify the colouring scheme so that either p and r, or q and s, have the same colour. Then we can find a colour for v, and again obtain a contradiction. (You can probably see how this can be done. If p,q,r,s are blue, green, red, and yellow respectively, then the graph T with v removed cannot contain *both* a chain of connected vertices from p to r, all blue or red, *and* a chain from q to s, all green or yellow; for these chains have to cross somewhere, and they can't.) Thus Kempe showed that T cannot contain 3(c) either.

If he could have shown that 3(d) was also ruled out, he would have completed his proof. Unfortunately, he tried to use the same trick for 3(d) as he had for 3(c), and thereby made his mistake, because the argument breaks down.

Nevertheless he made a very fine contribution towards the solution of the problem, often underestimated by later writers. Although his "proof" was fallacious, and hence technically worthless, the slightest modification of his argument yields a valid demonstration that *five* colours suffice; and his arguments have formed the foundation for most subsequent work on the problem.

The two main steps

To summarize Kempe's argument in modern terminology, he attempted to exhibit a set U of configurations (3(a)-(d)) such that:

(i) U is *unavoidable*: every plane triangulation contains one of the configurations in U;

(ii) Every configuration in U is *reducible*: it cannot be contained in a minimum counterexample to the 4-colour conjecture (i.e. any counterexample containing it also implies the existence of a smaller counterexample).

If his attempt had succeeded, it would certainly have provided a proof. It failed, because he did not show satisfactorily that (d) is reducible. Appel and Haken have been successful with exactly the same approach. But while Kempe's unavoidable set contained 4 configurations, theirs contains about 1930. (I say 'about' because they keep managing the reduce the number by 1 or 2.) The proof that these are all reducible involves massive reliance on the computer. One of their configurations is shown in Fig.4, and it is bordered by a circuit of 12 edges. All of their configurations are bordered by circuits of 14 or fewer edges. If they had used configurations larger than this, they would probably not have been able to prove them reducible with the present generation of computers.

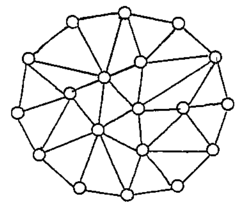

Fig.4

Appel and Haken's proof thus involves the above two steps: the construction of U, and the proof that everything in U is reducible. Each step is comparatively straightforward on its own: it is the interplay between them that is sophisticated, and in which Appel and Haken's work goes qualitatively, and not just quantitatively, way beyond anything that had been done before.

Contruction of an Unavoidable Set

To illustrate the first step, we show how Appel and Haken's method proves the set of configurations in Fig.5 unavoidable. The idea is due to Heesch.

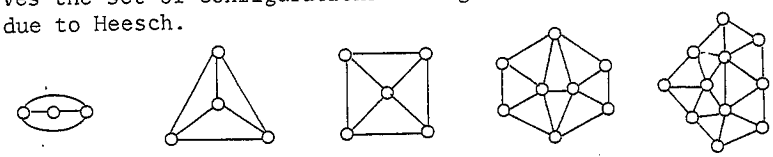

(a) *(b)* *(c)* *(d)* *(e)*

Fig.5

Suppose there exists a triangulation T not containing any of these. Assign to each vertex of T of valency i the number (6-i). Appel and Haken like to think of this as (6-i) units of electrical charge; so a 5-valent vertex receives charge +1, a 7-valent vertex charge -1, an 8-valent vertex charge -2, and so on. By (1), the total charge is positive (12 units).

We now redistribute the charge round T, without creating or destroying any, according to the following simple *discharging algorithm:* move 1/3 unit of charge for each vertex of valency 5 to each adjacent vertex of valency 7 or more. T still has positive total charge. But it is easy to check, using the fact that T contains none of 5(a)-(e), that no vertex of T can have positive charge! For T has no vertex of valency ≤ 4; any vertex of valency 5 is adjacent to at least three of valency 7 or more, so loses all its unit of positive charge; vertices of valency 6 are unaffected, ending up with charge 0, where they began; a vertex of valency 7 can have at most three neighbours of valency 5 (or two of them would be adjacent) and so recieves at most 1 unit of charge, remaining negative; and so on. This is a contradiction, so T must contain one of 5(a)-(e).

(Strictly speaking, this does not prove that one of 5(a)-(e) occurs in T *with all of its vertices distinct*. It is easy to get round this for small configurations, but for larger ones it is a serious technical problem, the *immersion problem*, and Appel and Haken had to deal with it.)

Appel and Haken proved their much larger set U unavoidable in this way, but using a more complicated discharging algorithm.

Reducibility

To illustrate this step we again take an example, showing that 6(a) is reducible.

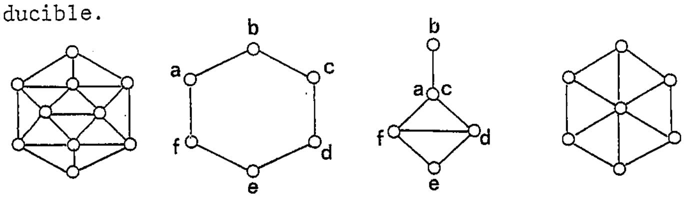

(a) *(b)* *(c)* *(d)*

Fig.6

Let T be a triangulation that is a minimal counterexample to the conjecture, and suppose T contains 6(a). Let T' be the graph obtained from T by removing the four vertices inside the hexagon in 6(a); that is, replace 6(a) by 6(b). By minimality of T, T' is 4-colourable. List the possible colour schemes for the vertices abcdef. There are 31 of them:

121212	121213G	121232	121234G	121312G	121314	121323G
121324G	121342G	121343G	123123	123124	123132G	121313
123134	123142	123143	123212G	123213G	123214G	123232G
123234	123242	123243	123412	123413	123414G	123423
123424G	123432G	123434G.				

(Here the numbers 1234 are the colours, listed in order on vertices abcdef. The G will be explained below. Note that 121211 and 121231 are not listed, since they give adjacent vertices the same colour (1); and 121214 is not listed since it comes from 121213 by permuting colours. Possibly not all of these can actually occur in T', but we don't know which do, so we have to consider them all.)

Some of these colour schemes can be extended to colourings of 6(a), so giving rise to 4-colourings of T. Call these *good* (which is what the G stands for). If all colour schemes are good, then 6(a) is clearly reducible (because we can 4-colour T, a contradiction). However, this never happens in practice.

The next step is to try to use Kempe-chain arguments to convert bad schemes into good ones. For example, 121232, which is bad, can always be converted by [13][24] interchanges into one of 121434, 121234, 121432, or 123232 - all good. If every bad colour scheme converts to a good one like this, then again 6(a) is reducible: we say a configuration that can be proved reducible this way is *D-reducible*.

The first thing the computer checks for is D-reducibility. (You should now see why the size of the outer ring is crucial!) If not, the next step is to note that we don't actually have to consider all 31 schemes on the list. By minimality, we can replace 6(a) by any configuration with fewer vertices, such as 6(c): the result T" must be 4-colourable. The effect of this substitution, here, is that we need consider only colour schemes where a and c have the same colour, and d and f are different: this rules out all but 6 of the schemes listed. If (as in this case) all the remaining schemes are good, or can be made so by Kempe-chain modifications, then again we get reducibility. There are many choices in place of 6(c) - another example is 6(d), which shows we need consider only schemes using 3 or fewer colours on abcdef. If any such substitution works, we call the original configuration *C-reducible*.

The program used by Appel and Haken, largely written by a postgraduate student John Koch and using algorithms of H.Heesch, first checked for D-reducibility; if this failed, it tried a few ways of proving C-reducibility. If these didn't work it was abandoned and the unavoidable set U modified appropriately. This may seem a

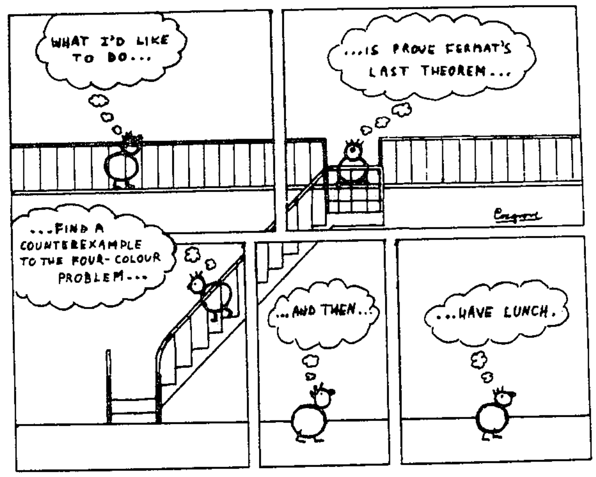

M-3

very cumbersome approach - especially since circuits like abcdef but with up to 14 vertices were involved. (Appel estimates that the amount of work goes up by a factor of 4 for each extra vertex in the circuit.) It might seem that it is better to test for C-reducibility first. But in practice this involves a lot of duplication of effort if the first substitute configuration doesn't work; and it is quicker to start by listing all the colour schemes to see which can be made good.

Conclusion

The main point I have not explained is the method by which the discharging algorithm and the unavoidable set were modified every time a configuration could not quickly be proved reducible. These modifications relied on a large number of empirical rules which have still not been given adequate theoretical justification, discovered in the course of a lengthy process of trial and error lasting over a year. By then Appel and Haken had developed such a good feeling for what was likely to work (even though they couldn't always explain why) that they were able to construct the final unavoidable set without using the computer at all. This is the crux of their achievement. Unavoidable sets had been constructed before, and configurations proved reducible; but no one could complete the monumental task of constructing an unavoidable set consisting entirely of reducible configurations.

The length of the proof is unfortunate, for two reasons. First, it makes it hard to verify. A long proof may take a long time to check, and be intellectually accessible to only a few people. This is particularly true if a computer is involved. Before the introduction of computers into mathematics, every proof could be checked by anyone possessing the necessary mental apparatus. Now an expensive computer may be needed too. Appel estimates that it would take 300 hours on a big machine to check all the details. Few mathematicians in Britain have access to this much machine time.

The other big disadvantage of a long proof is that it tends not to give much understanding of why the theorem is true. This is exacerbated if the proof involves numerous separate cases, whether it needs a computer or not. Lecturers may tend to give students the impression that proving theorems is the objective of pure mathematics; but I am sure that many of us agree that proofs are only a means to an end - understanding what is going on. Sometimes a proof is so illuminating that one feels immediately that it explains the 'real reason' for the result being true. It may be unreasonable to expect every theorem to have a proof of this sort, but it seems nonetheless to be a goal worth aiming for. So undoubtedly much work will be done in the next few years to shorten Appel and Haken's proof, and possibly find a more illuminating one. (It is doubtful that their method can be shortened enough to avoid massive use of the computer.)

In fact, there remain a number of conjectures that would imply the truth of the 4-colour theorem, but do not follow from it. One of these in particular (Hadwiger's Conjecture) is (in my opinion) most unlikely to be provable by the sort of technique that Appel and Haken have used: possibly a shorter proof of the 4-colour theorem may be found from an attack on Hadwiger's Conjecture. None of this, of course, detracts in any way from Appel and Haken's magnificent achievement.

BIBLIOGRAPHY

A slightly expanded version of this article, with references, appeared in the Bulletin of the IMA 14 (1978) 245-299. It also formed the basis for the article *The Appel-Haken proof of the Four-Colour Theorem* which formed chapter 4 of *Selected Topics in Graph Theory*, ed. L.W.Beineke and R.J.Wilson, Academic Press 1978.

* * * * * * * * * *

Early news of Appel and Haken's achievement was greeted by the mathematical community with less than unrestrained enthusiasm. One reason for caution is of course the computer involvement: it is extremely easy to make slips in long and involved programs. Now, 5 years later, no such slips have been found; and the program has been checked by a great many people. The expert view is that, if there are any errors, they do not occur in the computer part of the proof. As for the lack of elegance, MANIFOLD-19 *remarked: "maybe most theorems are true for rather arbitrary and complicated reasons. Why not?"*

Seven years of manifold

1968 – 1980

Edited by IAN STEWART and JOHN JAWORSKI