G R E S H A M

C O L L E G E



# BEYOND THE FOUR COLOUR THEOREM

A Lecture by

**PROFESSOR IAN STEWART MA PhD FIMA CMath**
**Gresham Professor of Geometry**

19 November 1997

# *GRESHAM COLLEGE*

## Policy & Objectives

An independently funded educational institution, Gresham College exists

- to continue the free public lectures which have been given for 400 years, and to reinterpret the 'new learning' of Sir Thomas Gresham's day in contemporary terms;

- to engage in study, teaching and research, particularly in those disciplines represented by the Gresham Professors;

- to foster academic consideration of contemporary problems;

- to challenge those who live or work in the City of London to engage in intellectual debate on those subjects in which the City has a proper concern; and to provide a window on the City for learned societies, both national and international.

***Founded 1597***

Gresham College, Barnard's Inn Hall, Holborn, London EC1N 2HH
Tel: 020 7831 0575 Fax: 020 7831 5208
e-mail: enquiries@gresham.ac.uk
Web site : www.gresham.ac.uk

# Gresham Lecture
# Beyond the Four Colour Theorem

Ian Stewart 19 November 1997

## Introductory Concepts

A *map* is an arrangement of regions, either in the plane or on a surface such as a sphere. Each region is a single connected portion of the plane or surface, and the regions make contact along common boundaries, which are curves. Often we make additional assumptions — for example that no region completely contains another region.

A *graph* is a diagram formed from a number of blobs, called *nodes* or *vertices*, which are joined together by a number of lines, known as *edges*. Graphs are simpler and more abstract than maps.

However, any map can be represented by assigning a node to each region and joining two such nodes by an edge if and only if the corresponding regions share a common stretch of border (**Fig.1**). Imagine the nodes as capital cities, and the edges as highways that join cities in adjacent countries, crossing at their common border. This is the *map graph*. It represents which regions share a common boundary with others, but removes from consideration various distracting complications, such that the shapes of the regions. For many questions, the shapes don't matter, and it's often easier to get rid of them altogether — hence the map graph.

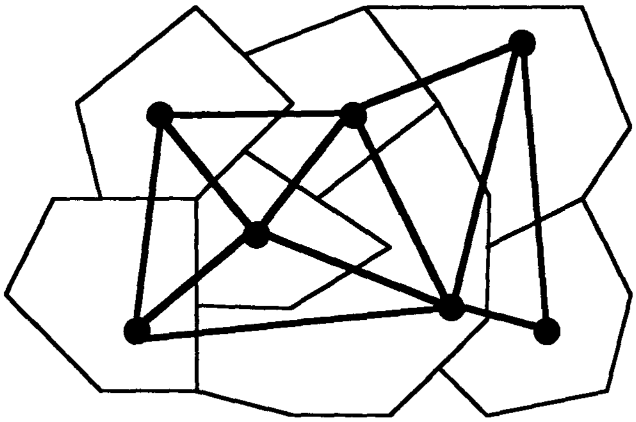

**Fig.1** A map and the corresponding graph.

A grap is said to be *planar* if it can be drawn in the plane without any edges crossing. If we start with a map in the plane, then its map graph is obviously planar. More surprisingly, if map is drawn on the the sphere, or on several disconnected planes and spheres — as is the case for 'Earth/Moon' maps, introduced below — then the resulting graph is *still* always planar. To see why, imagine a map drawn on a sphere. Put a node in each region and whenever two regions have a common boundary, connect the corresponding nodes with edges. The result is a graph that can be drawn on a *sphere* without any edges crossing. However, any such graph can be opened up and spread out on a plane. To do this, imagine cutting a small hole in the sphere, which does not meet any of the nodes or edges of the graph. Now imagine that the sphere is made from elastic sheeting. You can pull that tiny hole, making it bigger and bigger. The rest of the sphere stretches and deforms, carrying the graph with it. By pulling it far enough you can flatten it out into a disk. Lay the disk on the plane, and you've now drawn the map graph on a plane without any edges crossing.

If the map is drawn on several spheres, we just do the same for each of them, and lay all the resulting disks out in the same plane without overlaps. The resulting graph will be disconnected — it will fall into several separate pieces, one for each sphere — but that's quite a common feature of graphs, and is allowed by their definition, so it doesn't matter.

An important graph for this column is the *complete graph* $K_n$, which has $n$ nodes, an an edge joining every pair of distinct nodes. **Fig.2** shows $K_5$. If $n \geq 5$ then $K_n$ is *not* planar.

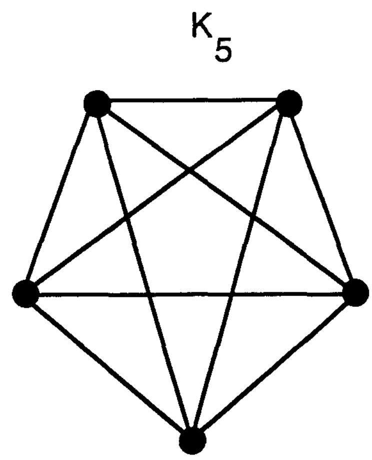


**Fig.2** The complete graph on five nodes.

A map (on a plane, sphere, several spheres, whatever) is said to be *k-colourable* if its regions can be coloured, using no more than $k$ colours, so that regions that share a common boundary curve receive different colours. (Regions that meet only at a point, or finitely many points, can if necessary receive the same colour.) The analogous property for a graph runs along very similar lines. A graph is $k$-colourable if its nodes can be coloured, using no more than $k$ colours, so that nodes joined by an edge receive different colours. It is easy to see that a map is $k$-colourable if and only if its map graph is $k$-colourable. Just colour each capital city, each node of the graph, with the colour of the corresponding country.

The smallest such $k$ is called the *chromatic number* of the graph: it tells us the minimum number of different colours needed for that graph — hence also for the corresponding map, if it is a map graph. Evidently $K_n$ has chromatic number $n$, because each node is joined to every other node, so no two nodes can be coloured the same.

**The Four Colour Theorem**

Colouring problems have been the object of mathematical study for about a century. The best known result is the famous Four colour Theorem, which says that every map in the plane can be 4-coloured. Percy Heawood proved long ago that every plane map can be 5-coloured: the number was reduced to four in 1976 by Kenneth Appel and Wolfgang Haken in a tour de force that combined mathematical analysis with extensive computer searches and calculations. To this day, no proof that avoids heavy use of computers is known. Many generalisations have been studied too, among them *Earth-Moon maps.* Each Earthly country has annexed a region of the Moon, to create an empire that consists of two connected regions: one on Earth, the other on its satellite. Between them, these regions cover both worlds completely. What is the smallest number of colours that will colour a map of *any* such disposition of territory, in such a manner that both countries in any particular empire receive the same colour, but no two adjacent regions receive the same colour — either on the Moon or the Earth? The answer remains unknown: it is either 9, 10, 11, or 12.

**m-Pires**

A problem closely related to Earth/Moon maps was introduced by Percy Heawood in 1890. The problem is set on the Earth only, but now each country is part of an empire containing a maximum of $m$ countries, and the same colour must be used for every country in a given empire, again with adjacent regions having different colours. (Countries in a given empire are assumed not to touch each other.) Such a map is punningly known as an *m-pire.* Heawood proved that an $m$-pire can always be coloured with 6m colours, for all $m \geq 2$.

Since an $m$-pire is a particular type of map, it has an associated map graph with one node per country. However, it is no longer true that every legal colouring of the map graph corresponds to a colouring of the empire. The reason is that the standard colouring rules for a graph fail to fulfil the requirement that nodes from the same *empire* receive the same colour. It is difficult to handle this condition using the map graph. Instead, the construction of the graph is modified so that the colouring rules are automatically correct.

Here's how.

The *m-pire graph* associated with a given $m$-pire map has one node for each empire (not one for each region). If you find this confusing, think of the node as representing the emperor. Two nodes are joined by an edge if and only if the corresponding empires include at least one pair of adjacent countries. You might think of the $m$-pire graph as the 'invasion graph' of emperors whose empires can go to war across a common border. One node per emperor, one edge for every possible two-sided war.

Conceptually, the $m$-pire graph is obtained from the ordinary graph by identifying all the nodes in a given empire — drawing them in exactly the same place. This construction often leads to multiple edges — two nodes joined by several edges instead of just one. Superfluous edges of this kind are removed, to leave just one edge.

Identifying all the nodes in a given empire automatically forces them to receive the same colour, so the number of colours needed for an $m$-pire is the same as the chromatic number of its $m$-pire graph.

In 1983 Brad Jackson (San Jose State U) and Gerhard Ringel (U California, Santa Cruz) used this approach to prove that the number $6m$ in Heawood's theorem cannot be reduced. They did this by demonstrating that you can find an $m$-pire whose m-pire graph is the complete graph $K_{6m}$. Since $K_{6m}$ definitely needs $6m$ colours, there is a $m$-pire that cannot be coloured with less than $6m$ colours.

**Earth/Moon Maps**

There are connections between Earth/Moon maps and $m$-pire maps. In fact, an Earth/Moon map can be viewed as a particular kind of 2-pire, with a slightly curious underlying geometry (two spheres) which splits all the 2-pires into two pieces. Its graph consists of two disjoint planar graphs — for example, one possible arrangement is shown in **Fig.3a**. (The rounded shape has nothing to do with the Earth or Moon: recall that any graph on a sphere, or several spheres, can be deformed so that it lies in a plane. It's just easier to show the shape of the graph here using curved edges.)

Suppose that we now think of this Earth/Moon graph as a 2-pire graph, so that nodes belonging to the same empire are identified to create **Fig.3b**. We see that the resulting graph need no longer be planar. Indeed this one isn't.

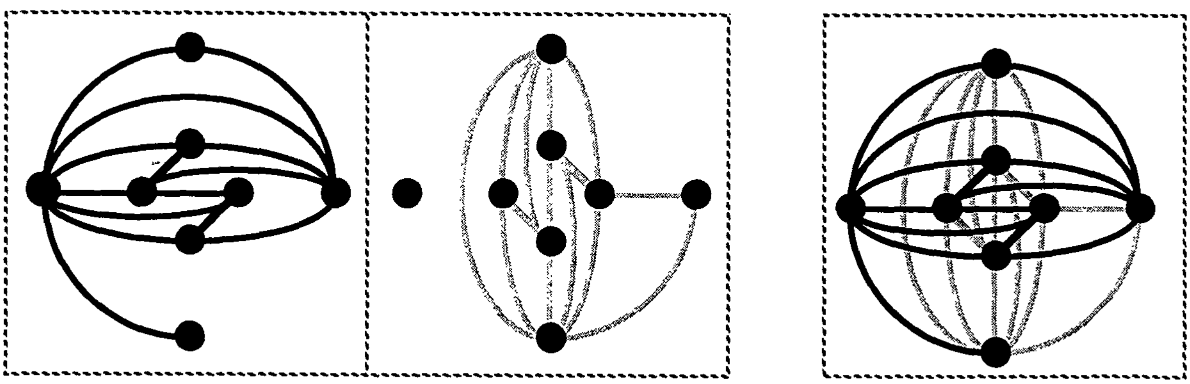

(a) (b)

**Fig.3** Earth/Moon map viewed as a 2-pire

However, the graph is 'almost planar'. The way it is constructed shows that its edges can be separated into two subsets, each of which forms a planar graph on the original set of nodes. Here the two subsets are the edges in Fig 3a and those in Fig.3b.

Such a graph is said to have *thickness* two. In general, a graph has thickness $t$ if its edges can be separated into $t$ subsets, and no fewer, in such a manner that each subset forms a planar graph. Now, every map graph is planar, even when the map lives on a sphere. An Earth/Moon map is made up from two separate planar maps: one on the Moon, the other on the Earth. Each empire is represented exactly once in either of these maps. So every Earth/Moon graph has thickness two: one planar bit for the Earth part, the other for the Moon part. The converse is also true: every graph of thickness two corresponds to an Earth/Moon map (although the territories involved may not completely cover the two worlds: there may be regions unclaimed by any of the empires).

Because an Earth/Moon graph is a special kind of 2-pire graph, Heawood's theorem implies that 12 colours are *sufficient* for any Earth/Moon graph. However, we can't conclude directly that 12 colours are also necessary. The reason is that not every 2-pire correesponds to an Earth/Moon map. In an Earth/Moon map, each empire has one region on the Moon and one on the Earth. If we think of this as a 2-pire, then the regions form two separate 'islands', and there is exactly one region from each empire on each island. In contrast, a 2-pire consists of a number of pairs of regions, which need not be arranged to form two islands — and even if they are, some empires might have both territories on the same island.

In fact, *none* of the known 2-pire graphs that actually require 12 colours can be turned into Earth/Moon maps. It therefore remains possible that *fewer* than 12 colours might always be enough for an Earth/Moon graph.

For instance, the complete graphs $K_9$, $K_{10}$, $K_{11}$ and $K_{12}$ are all 2-pire graphs, but they have thickness 3, and so cannot be Earth/Moon graphs (because those have thickness two). In fact, the thickness of $K_n$ is 3 if $n$ = 9 or 10, and is the greatest integer not exceeding $(n+7)/6$, otherwise.

Fig.3b is in fact the complete graph $K_8$, so $K_8$ has thickness 2. This means that it can be represented as an Earth/Moon graph. This proves that at least 8 colours are needed in the Earth/Moon problem. Sulanke has increased this lower limit to 9 by showing that the graph of **Fig.4** has thickness 2 and chromatic number 9.

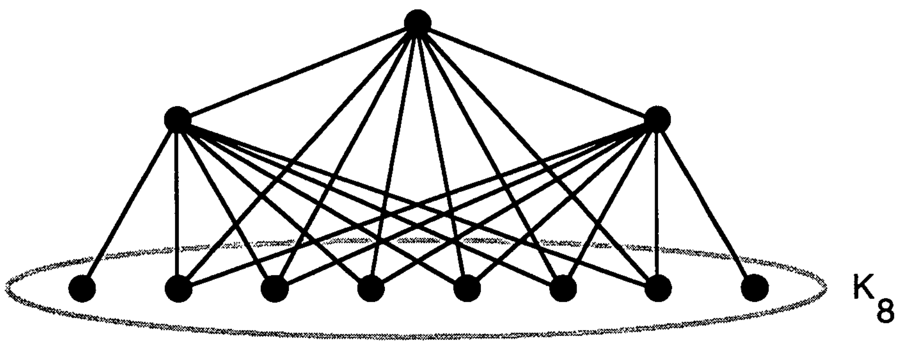


**Fig.4** Sulanke's graph has thickness 2 and chromatic number 9.

The concept of thickness, then, is the deep mathematical idea that underlies the recreational puzzle of Earth/Moon maps. You might like to think about Earth/Moon/Mars maps, where every emperor has *three* territories, one on each world. These maps are particular kinds of 3-pire map, and their 3-pire graph always has thickness three. In general a graph of thickness $t$ can be thought of as the $t$-pire graph of a system of galactic empires on a collection of $t$ planets.

**Application to Electronics**

Map-colouring problems of this kind are great fun — but they have little obvious practical significance. Even if we had planetary empires, the geographers could always colour their maps by trial and error — and in any case they might not want to follow our colouring rules. However, there *are* applications of the concept of thickness; however, they are not literal translations of the 'map' image. Instead, they apply to the testing of electronic circuits. In October 1993 Joan P. Hutchinson of Macalester College, St. Paul, Minnesota published a thorough survey of such questions: 'colouring ordinary maps, maps of empires, and maps of the Moon', *Mathematics Magazine* vol. 66 No.4 pp.211-226. In one section of the article she described an application of Earth/Moon colouring to the testing of printed circuit boards, discovered by researchers at AT&T Bell Laboratories, Murray Hill.

You can think of a graph of thickness two as a kind of 'sandwich'. On one slice of bread we draw the edges in the first set, none of them crossing; on the second slice, we draw the rest of the edges, again with none crossing. The nodes form the filling (**Fig.5**). A graph that needs $t$ layers of bread has thickness $t$.

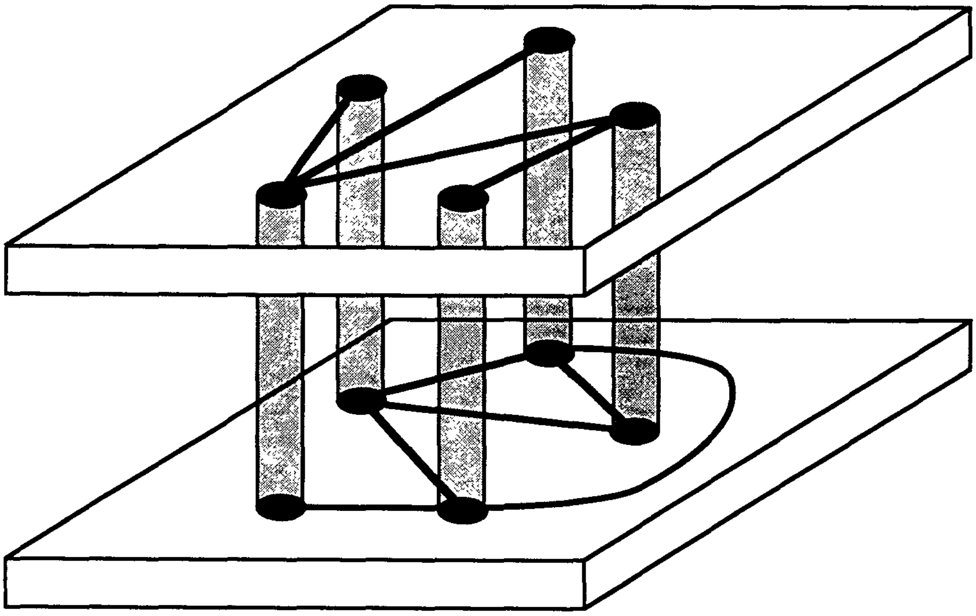

**Fig.5** Thickness-two graph as sliced bread.

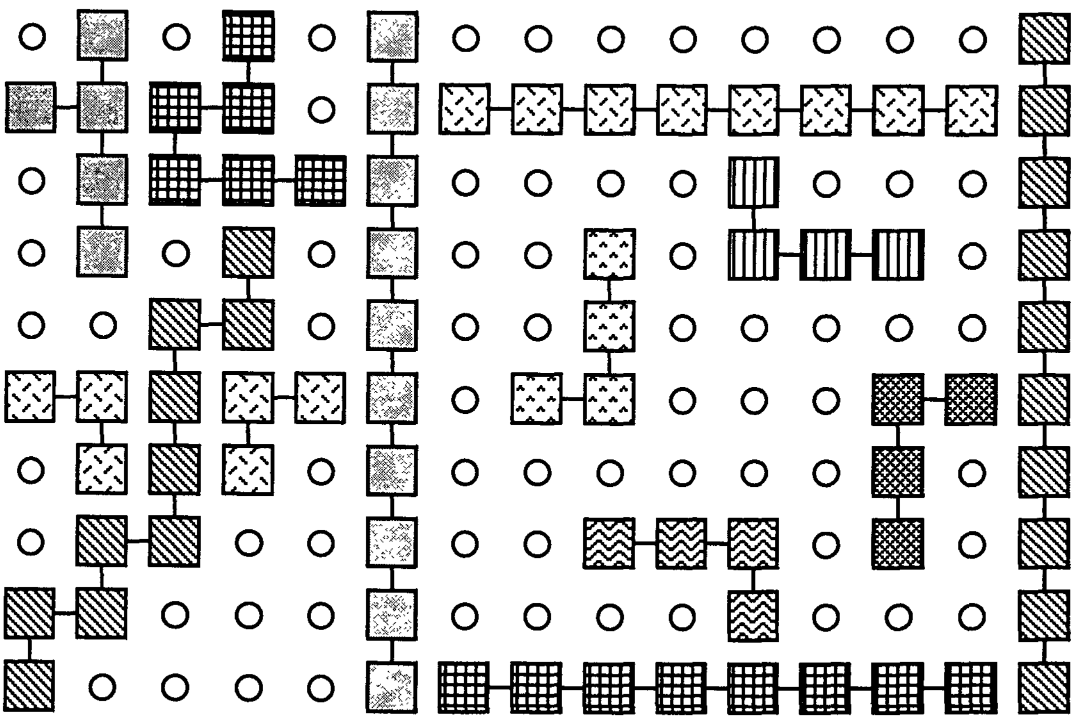

**Fig.6** A simple PCB circuit. Circles are holes for components, squares are components. Linked sets of squares are nets.

This image makes it clear why graphs and their thicknesses are relevant to electronic circuits. To begin with, think of an electronic circuit as a graph in its own right. The nodes are the electronic components, and the edges are electrical connections. If the circuit is to be constructed on one side of a printed circuit board (PCB) then it must be planar to avoid short circuits. By using two sides of the board — like the two slices of bread in the sandwich — graphs of thickness 2 become available. By using several boards, the thickness of the graph can be increased. Similar considerations apply in the more hi-tech world of silicon chips, too, because VLSI (Very Large Scale Integrated) circuits have to be built in layers.

A typical PCB is a 100×100 array of holes, where components can be attached, joined by horizontal and vertical lines that can be plated with 'tracks' of a conducting material, connecting the components together. An important problem for manufacturers of PCBs is to detect boards with spurious connections — extra bits of track that result in components being joined together electrically when they should actually be isolated from each other.

For practical reasons, manufacturers arrange the components on a PCB into 'nets'. A net is a collection of components, connected by tracks, so that the tracks contain no closed loops (**Fig.6**). In a well-made PCB, distinct nets should not be electrically connected. The problem that concerns us here is to determine, in an efficient manner, whether two distinct nets have inadvertently been linked together by an unwanted bit of track — a 'short circuit'.

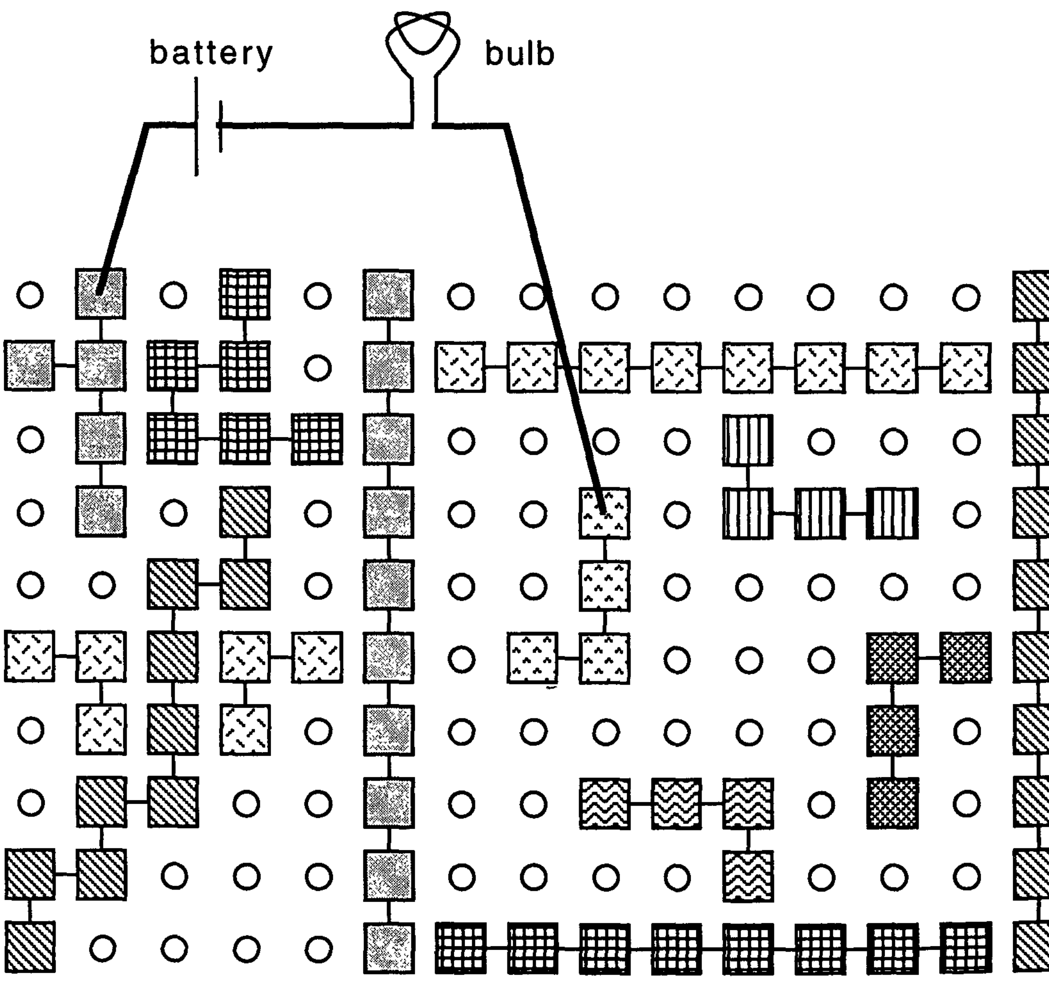


**Fig.7** Testing for a short circuit between a red net and a green net.

The most obvious way to do this is to check all pairs of nets to see whether they are connected. The simplest method is to make a 'test device' that creates a circuit that runs from one net to the positive pole of a battery, and from the negative pole through a light bulb to the second net (**Fig.7**). If the two nets are inadvertently connected by the PCB's tracks, then current will flow and the bulb will light. If not, it won't. Of course, a practical test device would use more sophisticated electronics — such as a computer attached to a robot that automatically discards a faulty board, instead of a light bulb — but that's the basic idea.

The practical snag is that with $n$ nets this method requires $n(n-1)/2$ tests — the number of pairs of nets. Since 500 nets is typical, that means 125,000 tests per board, which is much too big to be feasible. I will now convince you that applying the concept of the thickness of a graph quickly reduces the number of tests to a mere 11. In fact, a little extra thought reduces that number to just 4. This means that every board that is manufactured can be tested quickly and efficiently, so that those with unintended short circuits can be discarded.

The starting point is to turn the PCB design into a graph. The idea is to define the *simplest* graph that conveys information about short circuits between different nets: let me call this the *net graph* of the circuit design. The criterion of simplicity makes the construction of the net graph a little bit subtle. For example, because we are trying to find out whether or not there exist short circuits between different *nets*, there is no point in taking the nodes of the net graph to be the individual circuit components. Instead, we assign one node to each net. The edges of the net graph represent potential short circuits, not actual ones — because if we knew where the actual short circuits were, we wouldn't need to test the circuit. To be precise, two nodes of the net graph will be joined by an edge whenever the corresponding nets are 'adjacent' — meaning that they can be connected by a horizontal or vertical straight line that passes through no intermediate net (**Fig.8**).

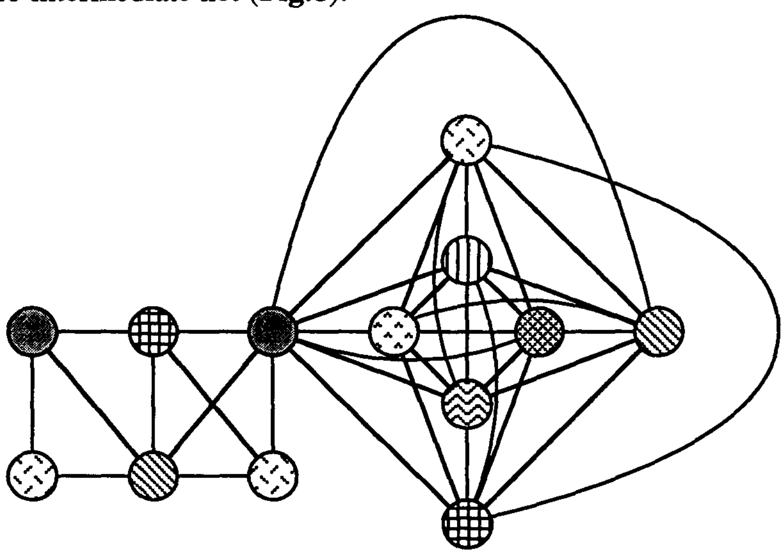

**Fig.8** The net graph for the circuit of Fig.6. colours of blobs correspond to colours of nets. The graph has been 8-coloured so that no two adjacent nodes have the same colour. Heawood's theorem guarantees a similar colouring for any net graph, but perhaps needing up to 12 colours.

Of course in principle a short circuit might exist that connects non-adjacent nets. However, nearly all such short circuits must *also* connect adjacent nets, because of the way the circuits are built. In a typical manufacturing process, the fabrication device makes two passes over the board: one to create the horizontal connections, the other to create the vertical ones. Errors arise when it lays down too much conducting material, inadvertently linking two nets that should remain disconnected: I'll call such an error a 'fabrication fault'. There are other possible ways to create a short circuit and produce a faulty board, but they are far rarer than fabrication faults, and we can ignore them.

Because the connections are laid down as horizontal or vertical lines, any fabrication fault must create an unwanted link between two adjacent nets. The extra line of conducting material may run across several further nets, but the first two that it links will necessarily be adjacent (**Fig.9**). In other words, we can detect fabrication faults by looking for short circuits between *adjacent* nets. In this sense, the edges of the net graph correspond to the possible mistakes in fabrication. The condition about there being no intermediate nets simplifies the graph, but does not lose sight of any possible mistake: instead of looking for all short circuits, it just looks for the 'minimal' ones.

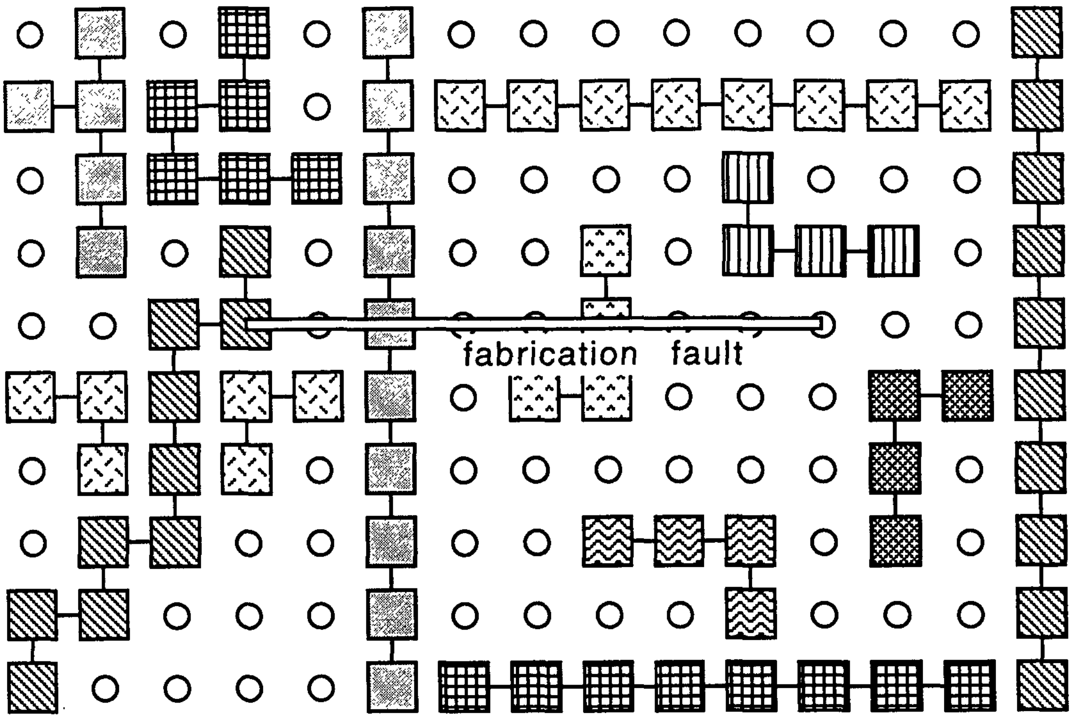


**Fig.9** Any short circuit caused by a fabrication fault must connect adjacent nets, even if it also connects others.

I said earlier that the graph whose nodes consists of the PCB's *components* has thickness two — one for each side of the PCB. The net graph also has thickness two, for the same reason. Last month I mentioned a theorem proved by Percy Heawood: any graph of thickness two can be 12-coloured. That is, each node can be assigned one of 12 colours so that nodes that are joined by an edge always have different colours. Heawood's theorem implies that the net graph of any PCB can be 12-coloured. We can transfer this colouring (conceptually) to the nets on the PCB. So the nets can each be assigned one of twelve colours, in such a way that nets of the same colour are never adjacent to each other.

Since we are seeking short circuits that link adjacent nets, we know that we can restrict our search to short circuits between nets of different colours. To discover whether such a short circuit exists, we can lump all the nets of each colour together, in the following sense. For each of the 12 colours we construct a 'probe'. This is a treelike structure made from conducting material that connects all the nets of a given colour together when it is brought into contact with the board (**Fig.10**) by the test device. Suppose that we choose two colours — say red and green. We attach both the red and green probes to the PCB, keeping them separate so that no electrical current can pass from the red probe to the green one *except* perhaps along the conducting tracks of the PCB. Now we connect a battery and a light bulb across the the two probes, and see whether any current flows.

If the PCB has been correctly made, no current will flow, because the red probe connects only to red nets, the green probe connects only to green nets, and on the PCB no red net should connect to any green net. However, if the PCB contains a fabrication fault that links a red net to a green one, then current will flow between the two probes. Now, any fabrication fault in the PCB necessarily connects two adjacent nets, and these must have different colours. So when the PCB is tested using the corresponding two probes, a current will flow in the test device.

Notice that this test doesn't tell us where the error is. Since we are discarding all faulty PCBs, not repairing them, we don't need to know. The upshot is that in order to detect the presence of a fabrication fault, it is enough to check for the existence of electrical connections — through the conducting material of the board — between all possible pairs of probes. Since

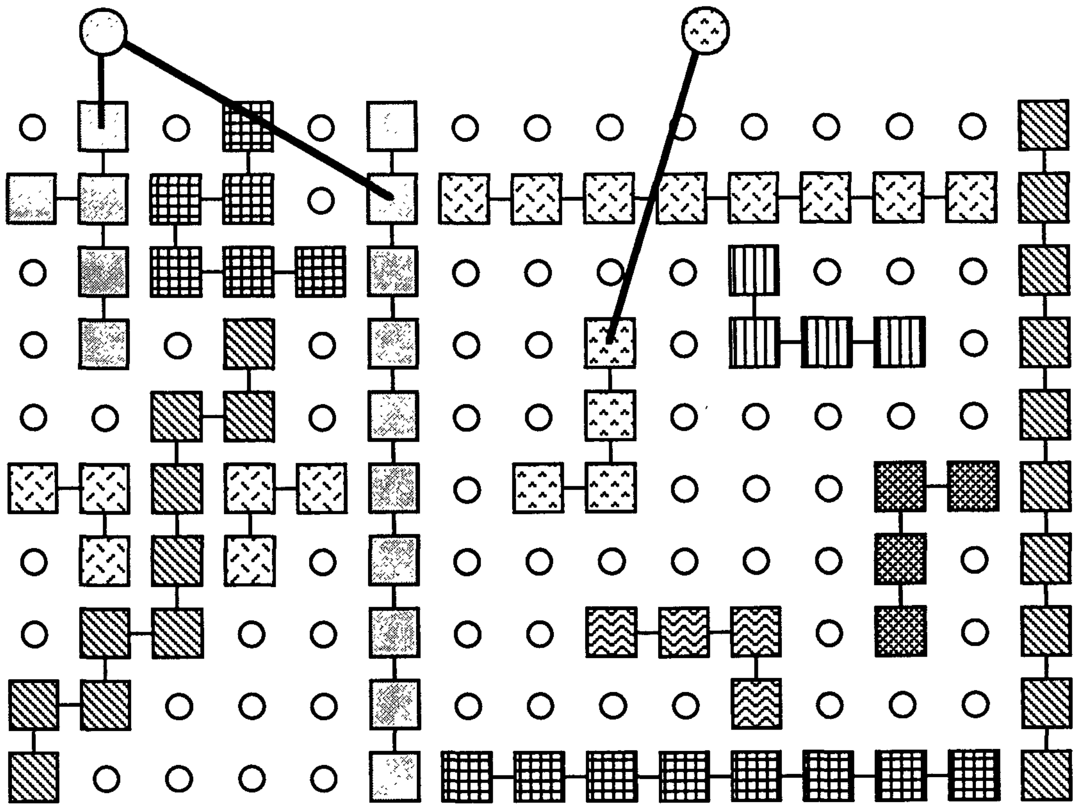

**Fig.10** Attaching two probes: one to all the red nets, one to all the green nets (here just one).

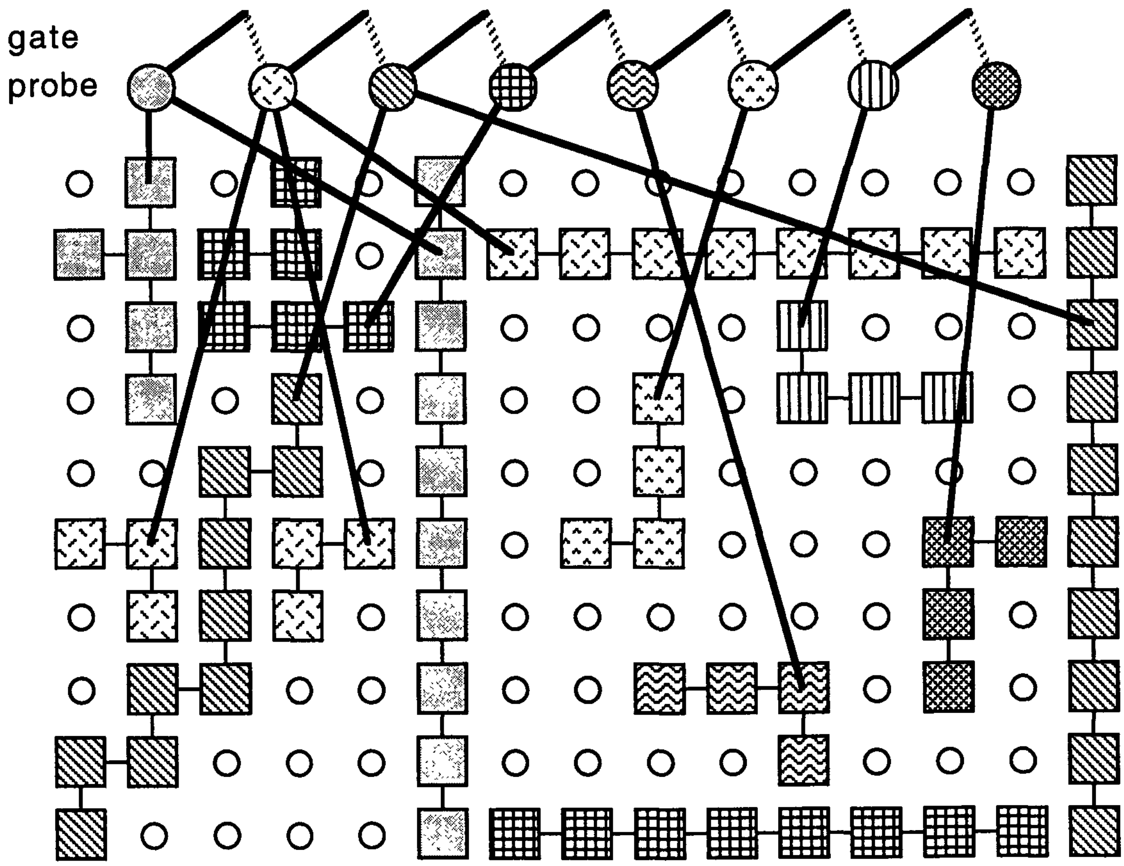


**Fig.11** Joining a complete system of probes with switchable gates.

there are only 12 probes, the number of such pairs is 12×11/2 = 66. So instead of 125,000 or more tests, we need only 66 — already a major improvement.

However, we can easily do better (**Fig.11**). Test probe 1 against probe 2; throw out any PCBs with connections between them. Now add a 'gate' to connect probes 1 and 2. Test probe 3 to see if it connects to the circuit formed by probes 1, 2, and the gate. If so, then probe 3 connects either to probe 1 or to probe 2. Either eventuality is a mistake, so we don't care which one occurs: we just throw the PCB out. Now add a second gate connecting probe 3 to the previous two, and continue in this manner. That gets the number of checks down to 11.

Allen Schwenk (West Michigan U, Kalamazoo) realised that a further reduction can be made. Write the numbers 1, ... , 12 in binary: 0001 up to 1100. Make a 'superprobe' that connects all probes that start with 0; make another that connects those starting with 1. Test whether these two superprobes are connected. If so, throw out the PCB. If not, create two more superprobes connecting probes that have the same binary digit in the second place. Check whether these are connected. Do the same for the third place and the fourth place in the binary expression. That's it. To see why it works, note that if two distinct probes are connected by a short circuit, then their binary expressions must differ in at least one of the four places, so one or other of the four tests will detect the mistake.

Of course there may be other errors in the PCB, but the ones eliminated by this method are much the most common. And a reduction from 125,000 tests *per board* to only four is well worth having as soon as the production run becomes reasonably big — because you only need build those complicated probes and superprobes *once* for each design of PCB. Indeed, a suitable 'programmable' probe/superprobe unit could cover all eventualities.



**Further Reading**

Joan P. Hutchinson, Colouring ordinary maps, maps of empires, and maps of the Moon, *Mathematics Magazine* October 1993 vol. **66** No.4 pp.211-226.

Ian Stewart, *From Here to Infinity*, Oxford University Press, Oxford 1996.