Exceptional simplicity

Professor Robert A. Wilson

2nd April 2008

Introduction

Thank you, Rob, for that introduction.

Ladies and gentlemen, colleagues, family and friends.

The purpose of this lecture is indeed for me to introduce myself and my work to you. Of course, as I am now coming to the end of my fourth year at QM, you might consider this a trifle belated.

There will be many of you who are wondering what on earth it is that a Professor of Pure Mathematics actually *does*, and I hope to enlighten you a little bit on this subject. There is a famous remark attributed to Paul Erdös to the effect that "A mathematician is a machine for turning coffee into theorems."

But perhaps more useful is this description by M. Egrafov. He said "If you ask mathematicians what they do, you always get the same answer. They think. They think about difficult and unusual problems. They do not think about ordinary problems: they just write down the answers."

I work in an area of mathematics called group theory, so first I ought to try to explain what group theory is, and why. I first became interested in group theory at the age of sixteen, when we were taught the basics at school, in order to help us with the Cambridge Entrance Examination. I found it so much easier than Calculus (which required knowledge), because it all seemed so obvious and you didn't need to know anything. Although I did not realise it at the time, I think that was the moment when my future career was determined.

What is group theory?

Group theory begins in the study of symmetry. Of course, mathematicians make it very abstract, mainly because this brings out the underlying structure more clearly, and shows how the concepts are much more general than you might at first have thought. This generality is what makes the subject so useful. Mathematics is really the ultimate transferable skill. As E. T. Bell (a famous historian of mathematics) has said:

> Abstractness, sometimes hurled as a reproach at mathematics, is its chief glory and its surest title to practical usefulness.
>
> It is also the source of such beauty as may spring from mathematics.

He also said:

> Wherever groups disclosed themselves, or could be introduced, simplicity crystallized out of comparative chaos.

The simplest form of symmetry is bilateral (two-sided, or mirror) symmetry. Here there is only one 'symmetry', namely reflection in the vertical line down the middle. To put it another way, there are two indistinguishable versions of the picture:

- the original version
- after reflection

Abstractly we say the *group* of *symmetries* has *order* 2, meaning there are two *elements* of symmetry (i.e. "leave it alone", and "reflect in the vertical axis").

A more interesting example

An equilateral triangle has mirror symmetry, but now in three different axes: there are three different places you can put your mirror.

What happens if you reflect first in the mirror m_1 and then in the mirror m_2?

Well, you see that we get a *rotation*. How far do we rotate? One third of a revolution *clockwise*.

Similarly, by combining different reflections we get other rotations. There are six symmetries in all: three reflections, two rotations, and the "do nothing" symmetry. Let us agree to call the "do nothing" symmetry a rotation (through an angle of zero degrees).

This is the symmetry group of the *kaleidoscope*, which you may have played with as children: it is made with two mirrors set at 60° to each other, so that by repeated reflections you get round the full cirlce. A kaleidoscope therefore gives you a huge variety of patterns, all with the same symmetry group.

In this example we have already seen most of the important properties of such a group of symmetries. As P. R. Halmos has said:

> The source of all great mathematics is the special case, the concrete example. It is frequent in mathematics that every instance of a concept of seemingly great generality is in essence the same as a small and concrete special case.

Or as David Hilbert put it:

> The art of doing mathematics consists in finding that special case which contains all the germs of generality.

The essential properties of symmetry are these:

- you can follow one symmetry r by another s, to form the composite symmetry, called $r + s$
- "leave it alone" is a symmetry, called 0 (because it does *nothing*), so $0+r = r + 0 = r$
- "Newton's law": every symmetry r has an equal and opposite symmetry s, which "undoes" it. Thus $r + s = s + r = 0$.
- " r then s then t" is the same whichever way you read it: $(r = s) + t = r + (s + t)$.

These are all the properties you need to define a *group*. Everything there is to say about symmetries in the abstract, follows from these simple rules. That is probably why, at the age of 16 I felt that group theory was exceptionally simple.

Reflection groups

The groups we've seen so far are *generated by* reflections, in the sense that all the symmetries can be made by combining (some number of) reflections. These are the *kaleidoscope groups*, or as we group theorists say, *reflection groups*.

What are the possible kaleidoscope groups? Let's look first at the case where there are just two mirrors, set at some angle A. We need to know that when we repeatedly reflect the mirrors in each other, we eventually get back to where we started. Since two successive reflections combine to give a rotation through an angle $2A$, this tells us that after some number n of these rotations, we get through the whole 360°, so $2An = 360°$ and the angle A is $(360/2n)°$.

What about the case where there are three mirrors? You could have three vertical mirrors, forming a small room which you could stand in. What could the angles A, B, C between the mirrors be if you want to get a nice regular pattern? Well, you obviously want to have a kaleidoscope effect at each of the three corners, so you need the angles to be $180/a$ and $180/b$ and $180/c$ for some whole numbers a, b and c. But the angles in a triangle add up to 180 so $1/a + 1/b + 1/c = 1$ and with a little bit of thought you can deduce that the only possibilities are $(3, 3, 3)$ (an equilateral triangle), $(2, 3, 6)$ (half an equilateral triangle) and $(2, 4, 4)$ (a right-angled isosceles triangle).

Let's draw some pictures of these kaleidoscope patterns: take the *fundamental region* surrounded by the three mirrors, and reflect it in all the mirrors. And then reflect the reflections in the other mirrors, and so on.

These patterns show another type of symmetry, namely *translation symmetry*. If you imagine the pattern extending to infinity in all directions, then you can move the whole pattern sideways by a fixed amount, and the pattern looks the same as it did before. The symmetries are different from the ones we've seen before, in that if you keep doing them again and again, you never get back to the starting position. We say they have *infinite order*. (A reflection has order 2, and a rotation through an angle of $360/n$ degrees has order n, because after n such rotations you get back to where you started.)

Now let us look at some examples from the Al-Hambra.

There is another way to get a three-mirror kaleidoscope: imagine two vertical mirrors, and one slanting one, so that the three mirrors meet at a point. What sort of patterns can you get then? I won't prove this now, but it turns out again that there are exactly three possibilities. The patterns formed by each pair of mirrors again have to be kaleidoscopes, and the only combinations which work are $(2, 3, 3)$, $(2, 3, 4)$ and $(2, 3, 5)$. These give you 3-dimensional kaleidoscopes with symmetries of (a) a tetrahedron, (b) a cube or octahedron, and (c) a dodecahedron or icosahedron.

The last of these is also the symmetry group of a football. You will find that there are exactly 120 copies of the fundamental region, of which 60 are 'right-handed' and 60 are 'left-handed'. This means the reflection group has order 120, and the rotation subgroup has order 60.

If you had four mirrors you could make an infinite repeating pattern in three dimensions, or a finite pattern in four dimensions (whatever that means). The most interesting of all the reflection groups occurs only in *eight* dimensions, which is a little hard to visualise. We put together 8 mirrors, at angles of either $60°$ or $90°$ to each other, according the rule given by this diagram:

if there is a line between two mirrors, they are at 60°, and if not, they are at 90°. This group is called E_8. The total number of mirror symmetries (counting all the reflections of the mirrors as well as the real mirrors) is 120. [It is not entirely a coincidence that this is the same as the number of fundamental regions of the football group.] The total number of symmetries of this 8-dimensional crystal is much larger, at 696729600. [This is the same as the number of reflected copies of the fundamental region.]
Here are some pictures of E_8 projected into 2 or 3 dimensions: thes pictures have 240 points: for each of the 120 mirrors we have one pair of points at opposite sides of the mirror.

Quotient groups

We need only one more concept to conclude our crash course in group theory, and that is the concept of a *quotient group*. This is not really a difficult concept, although it can be confusing on first acquaintance.

In case any of you get confused by what I am about to say, let me reassure you with this quote from John von Neumann, a famous mathematician who was one of the pioneers of the computing age in the 1940s.

> In mathematics you don't understand things, you just get used to them.

And although this sounds like a flippant remark, it contains a lot of truth. I think many people make the mistake, in trying to learn mathematics, of trying to *understand* something which cannot in fact be understood. (It's a bit like trying to understand God, I suppose.) I remember as a student going to a course on Quantum Mechanics, and getting into quite a state of frustration because I could not understand where Schrödinger's Equation came from. It was years before I understood that it didn't come from anywhere—it was just an *assumption*. And years later still I realised that actually *nobody* understands quantum mechanics. But quite a lot of people have got used to it and can use it effectively.

So, don't try to understand what a quotient group is—just try to get used to the idea!

A quotient group can always be obtained by *forgetting some structure*. Take as an example the Star of David, consisting of two equilateral triangles superimposed on each other thus:

Or think of it as a hexagon with six triangles stuck on the edges. Its symmetry group is just the same as the symmetry group of the hexagon on its own: consisting of 6 rotations and 6 reflections, making a group of order 12 altogether.

Let us now forget the sides of the triangles, and just look at the two triangles as objects themselves. Each symmetry of the star *either* takes one triangle to the other one, *or* it leaves the two triangles where they were (but possibly reflected or rotated in itself). Now there are only two elements of symmetry: the only important question is, have I swapped the triangles or not?

This gives us back the original mirror symmetry group that I started with. But "leave it alone" now means "don't swap the triangles". No-one said anything about rotating or reflecting it—so if we only rotate or reflect the two triangles individually, without swapping them with each other, that now counts as leaving it alone.

The upshot of all this is that we have factorized our group (of order 12) into a subgroup of order 6 (the reflection group of one triangle) and a quotient group of order 2 (the 'swapping' group). Notice that we also factorize the *order* of the group as $12 = 6 \times 2$.

And now we can do it again: the symmetry group of the triangle can be factored into a subgroup of order 3 (the rotation subgroup) and a quotient of order 2 (the 'flipping' group—flip over vs. don't flip over).

The crucial question now is, how far can we factorize a group? This is a bit like factorizing numbers: you can keep on factorizing numbers until you get down to *prime numbers*, where you can't go any further.

The prime numbers are 2, 3, 5, 7, 11, 13, 17, ...

The corresponding things in group theory are called *simple* groups. Reflection groups can never be simple (unless there's only one reflection), because there is always a quotient of order 2: reflections "flip" and rotations "leave alone" the handedness. But if we just take the rotation subgroup, that can be, and often is, simple. For example, the rotation group of a regular n-gon is simple whenever n is a prime number. The rotation group of the football is also simple. We call it A_5.

Here is another example of a quotient group. In this case we obtain a finite quotient of an infinite group. Take the hexagonal tessellation of the plane, and imagine rolling it up into a cylinder. Now the horizontal translations, which used to have infinite order, now have finite order (depending on how tightly we have rolled the cylinder). Next imagine the cylinder is made of rubber, so we can roll it up again into a doughnut shape. Thus we can make the vertical translations have finite order as well, and the whole group is now finite.

Here is an example, where the full translation group now has order 7, and the rotation group has order 6. There are no longer any reflection symmetries. This is what it looks like when rolled up.

The complex world of simple groups

In 1980, it was announced that one of the biggest problems in mathematics had been solved, namely the Classification of the Finite Simple Groups. In other words, after the work of many hundreds of mathematicians over several decades, we knew at last *all* of the finite simple groups.

Every (finite) simple group is one of:

- a group of prime order
- $A_5, A_6, A_7, \ldots$, i.e. A_n for n at least 5
- in various other infinite families of groups
- 26 *sporadic* simple groups

The smallest sporadic simple group is M_{11}, discovered by Emil Mathieu in 1860, which has only (!) 7920 elements (i.e. symmetries).

The largest sporadic simple group is the Monster, discovered in 1973, which has order 808017424794512875886459904961710757005754368000000000

Remember this is the number of symmetries of the object—approximately equal to the number of elementary particles in the sun.

The existence of the Monster is one of the strongest arguments I know in favour of the Platonist view of mathematics (that mathematics is discovered, not invented). For who could invent the Monster if it did not exist? (Plato apparently was led to this conclusion by the existence of the dodecahedron)
The groups of prime order are the only *abelian* simple groups—named after Abel, a Norwegian mathematician of the early 19th century. A few years ago the Norwegian government founded the annual Abel prize in his honour. Last week the prize was awarded jointly to John Thompson and Jacques Tits for their ground-breaking work in group theory.

Moonshine

Look at these two slides: the first one is a table of information about representations of the Monster. The second one is a completely different piece of mathematics: it is a modular function, extensively studied in the nineteenth century. Indeed, it is the only 'simple' modular function, in the sense that it cannot be broken down into 'smaller' ones.

Now do you notice a remarkable coincidence? $1 + 196883 = 196884$. When John McKay pointed out this coincidence in the 1970s, he was politely told he was barking mad. No-one could seriously believe that nineteenth century analysis had anything to do with twentieth century group theory.

And even when he pointed out the next remarkable coincidence: $1 + 196883 + 21296876 = 21493760$, no-one took any notice. But eventually the message got through, and Conway and Norton investigated these coincidences in detail. They found many more coincidences which they developed in a series of conjectures, called 'Monstrous moonshine'.

(Monstrous, both because it was to do with the Monster, and because it seemed so outrageous. Moonshine, because no-one had any idea what was going on.)

You will not think McKay was so mad when I tell you that many of these conjectures have now been proved by Richard Borcherds—who was awarded the Fields Medal (the highest prize in mathematics) for this work.

Interlude

I am conscious of the fact that by this stage, many of you will be finding the mathematics a bit daunting, despite my best efforts to keep it simple. So I thought we should have a little interlude, inspired by Monstrous Moonshine. I shall play an extract from the well-known Moonshine Sonata by Beethoven. (This is usually known in English as the Moonlight Sonata, but in German it is called Mondschein, which translates as both moonlight and moonshine.)
(Since QMUL could not provide me with a piano in the lecture room, and the electronic keyboard they provided had a compass of five octaves C to C, as opposed to the classical compass of five octaves F to F, I played instead a recording I made on my piano at home.)

The finite simple groups

Well, I think that gives an idea of the mysteriousness of moonshine. Let us return now to studying the simple groups. What do they look like? How can we study them?

The groups of prime order are just the rotation groups of the regular polygons which have a prime number of sides: the triangle, the pentagon, etc.

The *alternating groups* can also be described as rotation groups: if in n-dimensional space, we put $n + 1$ points equally spaced from each other (like the equilateral triangle in 2 dimensions, and the regular tetrahedron in 3 dimensions, and so on), then the reflection group permutes these points in all possible ways. The rotation group has exactly half the permutations, and is simple provided n is at least 4.

The *classical groups* can be described by 'rolling up' n-dimensional space into a finite space. In this rolling up process, the whole numbers go round and round in a circle, just as in the polygon example: and just as in that example, the circle has to have a *prime* number of points on it. If this prime number is called p, we describe this as working *modulo* p. There are 6 families of these groups.

The *exceptional groups* fall into ten families, related in a complicated way to the exceptional reflection groups. The most complicated groups in this family are the groups of type E_8, so-called because they can be constructed from the reflection group of E_8 in 8-dimensions. However, we need a space of 248 dimensions to represent these groups. (8 for the reflection group itself, plus 240 for the 120 pairs of mirros points.)

The 26 *sporadic groups* are even more exceptional.

Representation theory

A group always arises in nature as the symmetry group of some object, and group theory in large part consists of studying in detail the symmetry group of some object, in order to throw light on the structure of the object itself (which in some sense is the "real" object of study).

But if you look carefully at how groups are used in other areas such as physics and chemistry, you will see that the real power of the method comes from turning the whole procedure round: instead of starting from an object and abstracting its group of symmetries, we start from a group and ask for *all possible objects that it can be the symmetry group of.*

This is essentially what we call *Representation theory.* We think of it as taking a group, and representing it concretely in terms of a symmetrical object.

Now imagine what you can do if you combine the two processes: we start with a symmetrical object, and find its group of symmetries. We now look this group up in a work of reference, such as our big red book (The ATLAS of Finite Groups), and find out about all (well, perhaps not all) other objects that have the same group as their group of symmetries.

We now have lots of objects all looking completely different, but all with the same symmetry group. By translating from the first object to the group, and then to the second object, we can use everything we know about the first object to tell us things about the second, and vice versa.

As Poincaré said,

> Mathematicians do not study objects, but relations between objects. Thus they are free to replace some objects by others, so long as the relations remain unchanged.

Some examples

Now I hope you remember our little group of order 6 that we have been studying. We took this group as the group of symmetries of an equilateral triangle, which is a *two-dimensional representation* of the group.

But this group has other representations. For example, each element (of symmetry) will move the three vertices around: it will permute them, so that for example the reflection m_1 changes the order from ABC to ACB. Every symmetry of the triangle corresponds to such a permutation of the three vertices, or of the letters ABC.

We call this a *permutation representation* of the group on three letters.

The geometrical representations that we looked at before are called *matrix representations*, because in higher dimensions it becomes harder and harder to visualize them geometrically. (You might manage three or four dimensions, but by the time it gets to 10 you'll give up, and we need to look in thousands and millions of dimensions.)

These are the two main types of representation theory, and in both cases there is a concept of a "simple" representation (traditionally called *primitive* for permutation representations, and *irreducible* for matrix representations). This is similar to the way we had simple groups, and made more complicated ones by sticking them together. (In all three cases, though, the "sticking together" is a very sticky subject.)

My research: part 1: strategies

In the rest of this talk I want to tell you something about my own work in group theory, representation theory etc. Of course, it is a tall order to summarise 25 years' work in 25 minutes, so I shall just pick a few highlights.

My strategy from my Ph.D. onwards has been to divide my work between tackling easy problems and tackling impossible problems. The former, because the bean-counters require a steady stream of 'outputs' which they can measure (or at least count). The latter, because recognition comes only from solving hard problems, preferably ones which everyone else has given up on as they consider them essentially impossible. And as Richard Parker never tires of telling me—never discount the possibility of success. [I remember he gave me that advice when I applied for this job!]

On the subject of bean-counting, I came across a quote from David Hilbert (who is most famous for his list of great mathematical problems for the 20th century, posed in 1900) which might be a useful addition to the current debate on how to measure research (which is of course unmeasurable).

> One can measure the importance of a scientific work by the number of earlier publications made superfluous by it.

On this measure my most important paper is a little sideline which I wrote with John Conway and Peter Kleidman more than 20 years ago, in which we used reflection groups (specifically E_8) to construct huge numbers of finite projective planes. (Infinitely many, in fact, but that's not the point—we constructed lots which looked very similar but were actually different.)

A *projective plane* is a geometrical construct of things which we shall call points and lines, with the property that through every pair of points there is exactly one line, and also every pair of lines meets in exactly one point. This is almost like Euclidean geometry, except that in Euclidean geometry you also have the possibility that two lines are parallel, and do not meet at all.

Here is an example of a projective plane with 7 points and 7 lines.

Ernie Shult described our paper as "wiping out an entire area of research", that is, the classification of projective planes, because we had convincingly demonstrated that there are just so many projective planes that a classification was essentially impossible.

My research: part 2: my PhD

To return to my PhD thesis. I was working on maximal subgroups of sporadic simple groups (this is equivalent to the problem of finding the simple permutation representations), and was making slow but steady progress on a couple of groups of order around 10^{11}, when my supervisor went away for a month without telling me. I was forced to talk to other people, in particular Simon Norton. With his help, I looked at another sporadic group (in fact Co_2) of order around 4×10^{13}, and finished it in a week. So I decided to go for broke, and look at Co_1, of order around 4×10^{18}. (Obviously an impossible problem: Rob Curtis had looked at part of this problem in his PhD thesis, but—correct me if I'm wrong—he never considered the whole problem to be tractable.) Anyway, this was a bit harder, and took me three weeks. So that was my PhD essentially sorted.

So this is my first example of tackling an obviously impossible problem. It also nicely illustrates another theme I want to emphasise, which is to simplify everything as much as possible. As Einstein said:

> Everything should be made as simple as possible, but not simpler.

The heart of my PhD was an observation so simple it is easy to miss:

> If a group fixes a point in space, then it still fixes a point when the space is rolled up (modulo 2, in my case).

My research: part 3: the Monster

The next impossible problem I want to describe is the construction of the Monster. Now Griess had originally constructed the Monster in a space with 196884 dimensions. His 102-page paper in Inventiones Mathematicae is a magnificent achievement. But I wanted the Monster more explicitly than that—I wanted it in my computer in a form where I could do calculations (ultimately to classify its maximal subgroups). I knew that the most compact form was in a space of 196882 dimensions over integers modulo 2 (in other words, using only 0 and 1, with the convention that $1 + 1 = 0$), and I knew that each group element would require 5GB of space. Clearly impossible, when I tackled this problem in the early 1990s.

In fact I'd already built several such representations (in much lower dimensions, of course), when in June 1991, Klaus Lux sent me an email from Australia, asking me for explicit representations of several sporadic groups. Some of these I had available, built either by me or, in most cases, by Richard Parker. But two I did not have and could not find. One was called the Harada–Norton group, and the corresponding representation was in 133-dimensional space. The other was the largest Fischer group Fi_{24}, and the representation was even bigger, in dimension 781.

So naturally I started on the smaller one, on Friday 14th June. I didn't have much time to work on it, as I had an orchestra rehearsal in the evening, a concert in Malvern on the Saturday, another concert in Whitchurch on the Sunday, and a lecture on Monday morning. However, my lab notebook from that time tells me that by 10.45pm on Monday it was completed.

The next one was finished by the following Friday, so I looked around for another problem of the same sort. There was only one that immediately leapt out at me, and that was the Baby Monster group, which has a representation in dimension 4370. I was not at all sure I could do this: I had never really worked in more than 100 dimensions before, and to suddenly start work in 4000 dimensions was rather daunting. But to cut a long story short, this construction was finished within 6 days (interspersed with numerous lectures, rehearsals, concerts and other interruptions).

And that was where I stopped for the time being. These matrices formed the start of what has now become the WWW Atlas of Group representations, which now has many thousands of representations of hundreds of groups.

The Monster tamed

So what of the Monster? Much worse than the 5GB per matrix storage problem, is the problem of multiplying these matrices together (that is, combining two symmetries into one): for the Baby Monster, on the IBM3090 mainframe that I used in 1991, this took about 10 minutes if I had the whole machine to myself. On my new desktop machine in 1992, it took about the same. The next year I got a much faster machine, and improved software, which brought the time down to about a minute and a half. Using the same algorithms, on the same machine, for the Monster, a single multiplication would take about three months.

And to do anything useful in the group, we'd need to do thousands of these multiplications, if not millions.
So a naive extension of earlier ideas certainly would not work. But by 1993 I was convinced that we had a plan that would work. Richard Parker and I discussed this problem on many occasions over the years. He came up with all sorts of good ideas, most of which I discarded. I started putting the plan into operation. I gave the job to my Ph.D. student Peter Walsh.

Naturally enough, he found that my plan had a few holes in it, which we patched up as we went along. When Peter finished his Ph.D., in January 1997, I set to work to complete what he had done. I had some discussions also with Steve Linton, so we ended up writing a four-author paper on this construction.

The calculations were actually completed during the Oberwolfach Computational Group Theory Conference in the first week of June, and I gave a talk on the construction a couple of days later.

The result was, not 5 gigabytes for each matrix, but about half a megabyte. (i.e. a factor of 10000 improvement) However, the downside was that we could only write down some of the symmetries: it is as though in our group of order 12 (symmetries of the Star of David), we could write down all the rotations, but only one reflection. We know we can make all the symmetries by combining these together, but it would be easier if we could write down all the symmetries in the same way.

And instead of taking three months to combine two symmetries, we take a fraction of a second to apply a symmetry to a point. You see, it is completely unnecessary to write down the whole matrix: that tells you what the symmetry does to *every* point. If you choose your points carefully enough (what I really mean is *carelessly* enough), it is good enough to see where the symmetry takes *one* point.

For example, if I draw a point with no symmetry on the star, then it goes to twelve different places under the symmetries. So each symmetry of the star corresponds to a different point in the picture. In other words, I know everything about the symmetry, by seeing where this one point goes.

So the Monster was tamed.

The Monster reveals its secrets

When I embarked on the construction, I had no ambitions to do anything with it. It seemed just so enormous that it would be impossible to actually use the matrices we had constructed.

However, it gradually became clear that *some* calculations were possible, and over time, we became more ambitious.

We showed that M is a Hurwitz group—an open question since the discovery of the group a quarter of a century earlier. A Hurwitz group is a bit like the rotation part of a reflection group, with mirrors at angles of 90°, 60° and $25\frac{5}{7}^{\circ}$. Of course you will complain that these angles add up to less than 180°. But do you remember the football group? Those angles were 90°, 60° and 36°, which add up to *more* than 180°. This is because triangles drawn on a sphere have their sides bulging out. Similarly we can invent a 'hyperbolic space' in which triangles have their sides squashed in.

My student Beth Holmes completed another construction of the Monster which I had started, which was designed to be more friendly for the maximal subgroup problem. Using this second construction, we found new maximal subgroups $L_2(59)$ and $L_2(29){:}2$—also longstanding open questions, which apparently cannot be answered any other way. And we have nearly, but not quite, completed the whole maximal subgroup question.

Moonshine and nets...

My research: part 4: recent results

To come further up to date. For about 5 years now I have been writing a book on the finite simple groups, in which I have been trying to give simple constructions of all the finite simple groups. Let me be clear, that this is another impossible task (!), and I have not succeeded, and will not succeed. Nevertheless, I have discovered some really nice simple constructions of some of the groups. Let me give you one example—the Ree groups. They were discovered in about 1960 by Rimhak Ree, who gave an algebraic construction of them. There are two families, one living in 7-dimensional space, the other living in 26-dimensional space. Very soon afterwards, Jacques Tits gave a more geometrical construction of the 7-dimensional case. But all these constructions are too complicated to go in a graduate-level textbook, so I wanted something simpler.

Another quotation from David Hilbert will set the scene:

> That I have been able to accomplish anything in mathematics is really due to the fact that I have always found it so difficult. When I read, or when I am told about something, it nearly always seems so difficult, and practically impossible to understand, and then I cannot help wondering if it might not be simpler. And on several occasions it has turned out that it really was more simple!

This accurately describes my relationship with the Ree groups. They seemed so difficult, and almost impossible to understand. The only solution was to build them myself from scratch (well, not really from scratch—I had the benefit of as much hindsight as I wanted to use).

We start with the projective plane of order 2 (which we have already seen) and build a space of 7 dimensions out of it: each blob represents one of the 7 dimensions, all perpendicular to each other. Now roll it up modulo 3, so that no matter what direction you walk in, after three steps you get back home.

Now we build some symmetries: each line in the picture represents a (generalised) mirror—it fixes anything in the three directions it represents, and reflects (negates) everything in the other four perpendicular directions. Then there is the symmetry $t \mapsto t+1$ which rotates the 7 mirrors, each to the next (the next one after 6 is 0!). And rotating the picture gives us the symmetry $t \mapsto 2t$.

Next, we build another 7-space according to the rule given in this picture. [In more mathematical terms, each arrow represents a wedge product, and the rule is

$$t^* = (t{+}1)\wedge(t{+}3) - (t{+}2)\wedge(t{+}6) \quad (\mathrm{mod}\ (t{+}1)\wedge(t{+}3) + (t{+}2)\wedge(t{+}6) + (t{+}4)\wedge(t{+}5))]$$

And the Ree group is the group of symmetries which act the same in both 7-spaces simultaneously. Actually, that is not true. There is a small technicality which is too technical to give in this lecture, but which is nevertheless relatively elementary. We have to introduce finite fields and field automorphisms, and 'twist' the construction using a field automorphism.

Conclusion

Norbert Wiener said

> A professor is one who can speak on any subject—for precisely fifty minutes

Since I have already run over that time, I had better stop here.